

USING SOCIAL MEDIA CONTENT TO INFORM AGENT-BASED MODELS
FOR HUMANITARIAN CRISIS RESPONSE

by

Sarah Wise
A Dissertation
Submitted to the
Graduate Faculty
of
George Mason University
In Partial fulfillment of
The Requirements for the Degree
of
Doctor of Philosophy
Computational Social Science

Committee:

_____ Chair of Committee

_____ Director of Graduate Studies

_____ Director, Krasnow Institute
for Advanced Study

Date: _____ Spring Semester 2014
George Mason University
Fairfax, VA

Using Social Media Content to Inform Agent-based Models for Humanitarian Crisis
Response

A dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy at George Mason University

By

Sarah Wise
Bachelor of Science
The University of Chicago, 2009

Director: Andrew Crooks, Professor
Department of Computational Social Science

Spring Semester 2014
George Mason University
Fairfax, VA

Copyright © 2014 by Sarah Wise
All Rights Reserved

Dedication

To John Christiansen, my mentor and friend.

Acknowledgments

This work would not have been possible with Dr. Andrew Crooks encouraging me, supporting me, and guiding me every step in the way. His effort and attention over the course of my graduate studies have shaped both this thesis and myself as a researcher, and I am the richer for it. It is my hope that with the completion of this work, he will again experience weekends. Andrew, I can never thank you enough.

I am also deeply grateful to the other members of my committee: Dr. Robert Axtell, Dr. Anthony Stefanidis, Dr. Arie Croitoru, and Dr. William Rand. Your insights were invaluable, and your patience and flexibility in dealing with the problems of a long-distance defense were superhuman. Thank you for your guidance, your knowledge, and your graciousness at every step.

The thesis effort was further supported by my entire family, in particular my mom/Number One Cheerleader. It would never have come off without the emotional and material support of Aurora Thome, John Paul Gonzales, and David Masad, each of whom I'm lucky to know. Karen Underwood of the Computational Social Science Department has my undying gratitude for getting me through the thesis process (and graduate school in general) in one piece. Further thanks are due to Dr. Tao Cheng of University College London for her understanding and generosity during my final months of writing. Finally, I would never have been in a position to do this research at all were it not for John Livingston and Joshua Strong, who told me not to quit coding at a crucial moment in high school, and for John Christiansen, who defined my undergraduate years and set me on this path.

To each of you, I extend my deepest thanks and affection. WE DID IT!

Table of Contents

	Page
List of Tables	x
List of Figures	xii
Abstract	xviii
I Background	1
1 Introduction	2
1.1 Motivation	2
1.2 Research Question	4
1.3 Research Contribution	5
1.4 Organization of the Dissertation	6
2 Humans in Crisis	9
2.1 Principles of Human Behavior	10
2.2 How Do People Behave in Crisis Situations in Particular?	12
2.2.1 The Role of Information	13
2.2.2 The Decision To Evacuate	13
2.2.3 Action	15
2.2.4 Emotion	16
2.2.5 Caveats	17
2.3 A Brief Overview of Modeling Human Behavior	18
2.4 Modeling a Crisis	21
2.4.1 Disease	26
2.4.2 Evacuation	26
2.4.3 Logistics	27
2.4.4 Manmade	28
2.4.5 Natural	28
2.5 Agent-Based Modeling	32
2.5.1 Agent-Based Traffic and Pedestrian Models	33
2.5.2 Agent-Based Modeling of Evacuation	33
2.6 Wildfires	34

2.7	Summary	36
3	Geospatial Information Systems and the Evolution of Data	38
3.1	The Rise of Web 2.0	39
3.1.1	Applications	39
3.1.2	Social Media	41
3.2	What is Big Data?	42
3.2.1	Gathering Big Data	43
3.2.2	Crowdsourcing	43
3.3	Big Geosocial Data	45
3.3.1	Volunteered Geographic Information	46
3.3.2	Ambient Geographic Information	48
3.4	Challenges	50
3.4.1	Validity	50
3.4.2	Theory	52
3.4.3	Other Concerns	52
3.5	Summary	53
4	Human Social Networks	55
4.1	The Size of the Network	57
4.2	Social Connections	59
4.2.1	Who Are These People, Anyway?	60
4.2.2	Making Contact	62
4.3	The Network Community	64
4.4	Social Media	66
4.4.1	Introduction to Twitter	67
4.4.2	Extracting Information Content from Social Media	69
4.4.3	Classification Problems	70
4.4.4	Measuring Social Networks	72
4.4.5	Challenges of Working with Social Media Data	74
4.4.6	Towards a Model of Twitter	75
4.5	A Worked Example	80
4.5.1	Geotagging	82
4.5.2	Hashtag Tracking	85
4.5.3	Social Communities	88
4.6	Summary	89

II	Methodology	92
5	Sentiment Detection	93
5.1	History	94
5.2	Uses of Sentiment Detection	94
5.3	Analysis Techniques	96
5.3.1	Levels of Analysis	96
5.3.2	Measures of Sentiment	97
5.3.3	Approaches to Sentiment Detection	98
5.3.4	Effective Features for Sentiment Classification	100
5.3.5	Selecting a Lexicon	101
5.4	Challenges	102
5.5	Methodology Employed in this Work	104
5.5.1	The Data	104
5.5.2	The Algorithm	106
5.5.3	Results	108
5.6	Sentiment Analysis of Real Data	110
5.7	Summary	116
6	Population Synthesis	118
6.1	Introduction	119
6.1.1	Synthetic Reconstruction	121
6.1.2	Combinatorial Optimization	122
6.1.3	Other Methodologies	123
6.1.4	Methodology Utilized in this Work	126
6.2	Data Gathering	127
6.3	House Creation	127
6.3.1	Derive Household Constraints	128
6.3.2	The Process	128
6.4	Households	129
6.4.1	Derive Population-Level Constraints	129
6.4.2	The Population of Individuals	131
6.4.3	Household Creation	131
6.4.4	Results	133
6.5	Workplace Creation	141
6.5.1	Data	141
6.5.2	The Process	142
6.6	Intimate Social Network Generation	143

6.6.1	Data	144
6.6.2	Process	145
6.6.3	Results	146
6.7	Social Media Network Generation	149
6.7.1	Data	149
6.7.2	Process	149
6.7.3	Results	152
6.8	Agent Synthesis	154
6.9	Summary	155
7	Modeling The 2012 Colorado Wildfires	156
7.1	Introduction	156
7.2	Incorporating Data	157
7.3	Model Overview	162
7.3.1	Purpose	165
7.3.2	Entities, state variables, and scales	165
7.3.3	Process overview and scheduling	169
7.4	Design Concepts	171
7.5	Details	172
7.5.1	Initialization	173
7.5.2	Input	173
7.5.3	Submodels	174
7.6	Summary	194
III Results		195
8	Verification, Validation, and Results	196
8.1	Control Case	198
8.2	Wildfire	198
8.3	Parameter Sweep	199
8.3.1	Communication Success Probability	199
8.3.2	Contact Failure Probability	200
8.3.3	Tweet and Retweet Probability	212
8.3.4	Comfort Distance	219
8.3.5	Observation Distance	219
8.3.6	Decay Parameter	226
8.3.7	Maximum Speed	226
8.4	Results	233

8.4.1	Sentiment	233
8.4.2	Road Usage	233
8.5	Summary	237
9	Discussion	238
9.1	Overview	238
9.2	Research Contributions	240
9.3	Limitations	242
9.4	Further Work	244
9.5	Conclusion	247
	References	248

List of Tables

Table	Page
2.1 Overview of different theoretical perspectives of behavior and the operational-ization of the perspective	11
2.2 A survey of models of crisis scenarios	22
2.3 A survey of agent-based models of crisis scenarios	30
4.1 An overview of the network terminology used in this thesis	58
4.3 Average number of people known for relation types (Source: McCarty et al., 2001)	61
4.4 Frequency of contact in terms of instances of contact (Source: Berg et al., 2010)	63
4.5 Properties of human social networks in general	65
4.6 Properties of samples taken from the Twitter social network	76
5.1 AFINN results comparison: total success scores	108
5.2 SentiWordNet results comparison: total success scores	108
5.3 ANEW results comparison: total success scores	109
5.4 AFINN results comparison: breakdown by classification scores	109
5.5 Comparison of valence classifications of real texts	113
6.1 Types of data required by the simulation	127
6.2 A comparison of the generated and real populations and households	135
6.3 A sample of the generated households showing individual members, where each line represents a household and each row member represents an individual as Age(Sex)	140
6.4 Parameters of the generated intimate social network	146
6.5 The percentage of different demographic groups which utilize Twitter. (Source: Duggan & Brenner, 2013)	150
6.6 Parameters of the generated social media network	153
7.1 Default parameters for Agents.	168
7.2 Parameters utilized by the Wildfire object.	170

8.1	The range of parameters utilized during the parameter sweep. Default values are shown in bold.	199
-----	--	-----

List of Figures

Figure	Page
3.1 The GalaxyZoo interface, with its user-friendly set of classification questions and simple interface	44
3.2 The OpenStreetMap iD editor, displaying a well-labeled area	47
3.3 A Facebook status with embedded geospatial information: an example of blended geospatial and social content	48
4.1 Comparison of rate of generation of all (red) versus specifically geotagged (blue) tweets, in linear (A) and log (B) scales	81
4.2 The set of geotagged tweets generated during the peak of the Waldo Canyon wildfire (from June 23 until July 5) mapped onto the Colorado Springs wider area road network to give a sense of Twitter user location during the peak of the wildfire. Individual tweets are designated by opaque red dots, so that multiple tweets from the same location are brighter red.	83
4.3 A closeup of the set of tweets generated during the peak of the Waldo Canyon wildfire mapped onto the Colorado Springs study area	84
4.4 Comparison of use of the “#waldocanyonwildfire” hashtag in linear (A) and log (B) scales. The lines indicate the rate of hashtag usage in all (red) or exclusively in geotagged (blue) tweets	87
4.5 The four largest communities detected by modularity in the dataset. The color and size of the nodes reflect their degrees, with larger and darker nodes having relatively higher degrees. (A) consists largely of Colorado social and news organizations; (B) centers around the username of the Denver Broncos and (C) around the Denver 9 News Channel; (D) contains a large number of Colorado Springs official organizations, including emergency personnel and local media sources	90
5.1 Tweet valences over the entire sample period	111
5.2 Tweet valences over the entire sample period, log-scaled to give a sense in the change of rank	112

5.3	Tweet valences over the peak period of the Waldo Canyon wildfire evacuations	114
5.4	Tweet valences over the peak period of the Waldo Canyon wildfire evacuations, log-scaled to give a sense in the change of rank	115
6.1	A flowchart of the population synthesis process	125
6.2	The set of all generated house locations	129
6.3	A close-up of generated house locations	130
6.4	Generated workplace locations	142
6.5	The large and very dense generated intimate social network. The color of a node indicates the community to which the node has been assigned, each community being represented by a different colour	147
6.6	The degree distribution of the generated intimate social network.	148
6.7	The generated social media network. The color of a node indicates the community to which the node has been assigned, each community being represented by a different colour	152
6.8	The degree distribution of the generated social media network	154
7.1	Map of study area with overlaid fire progression	158
7.2	Geographic data representing A) elevation, B) land use type, and C) permeability of surfaces	163
7.3	Map of study area with overlaid evacuation zones	164
7.4	Overview of the model structure, detailing how data is used within subprocesses which underlie the simulation	166
7.5	Overview of entities, their relationship to on another, and the parameters which effect them	167
7.6	Close-up of study area with the point at which the fire is initialized in the model versus the true extent of the wildfire on July 24 and June 4	176
7.7	Overview of an agent's activation of its subprocesses	178
7.8	The default behavioral subprocess of an Agent	179
7.9	GoogleMap's recorded average traffic data for Colorado Springs at A) 8am and B) 5pm on a Monday. Red areas indicate heavy traffic while green areas indicate relatively free-flowing traffic	181
7.10	GoogleMap's recorded average traffic data for Colorado Springs at 12am on a Monday. Red areas indicate heavy traffic while green areas indicate relatively free-flowing traffic	183
7.11	The observation subprocess of an Agent	184

7.12	The default behavioral subprocess of an Agent	185
7.13	A Fast and Frugal decision tree for Agents	192
8.1	The normalized heatmap (A), stress profile (B), and speed profile (C) generated by the default parameter settings of the complete model. Time is measured in 5 minute intervals, or 288 to a day	202
8.2	The normalized heatmap differential (A), stress profile profile (B), and speed profile (C) generated by the default parameter settings of the model with the wildfire submodel disabled. Time is measured in 5 minute intervals, or 288 to a day.	203
8.3	The development of the wildfire submodel: the yellow area indicates the extent of the wildfire at the end of the third day (A) and at the end of the peak wildfire period, after 11 days (B)	204
8.4	The real Waldo Canyon wildfire as it existed in the real-world at the end of the third day: the yellow area indicates the extent of the wildfire	205
8.5	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Communication Success Probability set at 10% (B) and 50% (C).	206
8.6	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Communication Success Probability set at 10% (B) and 50% (C).	207
8.7	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Communication Success Probability set at 10% (B) and 50% (C).	208
8.8	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Contact Failure Probability set at 10% (B) and 90% (C).	209
8.9	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Contact Failure Probability set at 10% (B) and 90% (C).	210
8.10	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Contact Failure Probability set at 10% (B) and 90% (C).	211

8.11	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Tweet Probability set at 50% (B) and 90% (C).	213
8.12	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Tweet Probability set at 50% (B) and 90% (C).	214
8.13	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Tweet Probability set at 50% (B) and 90% (C).	215
8.14	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Retweet Probability set at 50% (B) and 90% (C).	216
8.15	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Retweet Probability set at 50% (B) and 90% (C).	217
8.16	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Retweet Probability set at 50% (B) and 90% (C).	218
8.17	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Comfort Distance set at 100 m (B) and 100000 m (C).	220
8.18	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Comfort Distance set at 100 m (B) and 100000 m (C).	221
8.19	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Comfort Distance set at 100 m (B) and 100000 m (C).	222
8.20	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Observation Distance set at 100 m (B) and 10000 m (C).	223

8.21	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Observation Distance set at 100 m (B) and 10000 m (C).	224
8.22	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Observation Distance set at 100 m (B) and 10000 m (C).	225
8.23	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Decay Parameter set at 10% (B) and 90% (C).	227
8.24	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Decay Parameter set at 10% (B) and 90% (C).	228
8.25	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Decay Parameter set at 10% (B) and 90% (C).	229
8.26	Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Max Speed set at 1000 m/5 min (B) and 8000 m/5 min (C).	230
8.27	Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Max Speed set at 1000 m/5 min (B) and 8000 m/5 min (C).	231
8.28	Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Max Speed set at 1000 m/5 min (B) and 8000 m/5 min (C).	232
8.29	Visualization of the closed and heavy-traffic roads during the evacuation period. All images taken from the Denver Post (2012a); all road closure data taken from the Colorado Springs After Action Report (City of Colorado Springs, 2013)	235
8.30	A collection of images taken on June 26 in Colorado Springs, documenting the relative traffic levels in various parts of the city. All images taken from the Denver Post (2012a)	236

Abbreviations Used in this Thesis

Agent Based Modeling	ABM
Ambient Geographic Information	AGI
Belief, Desires, and Intentions	BDI
Emergency Notification System	ENS
Exchangeable Image File Format	Exif
Geographic Information Systems	GIS
Iterative Proportional Fitting	IPF
Location-based Services	LBS
OpenStreetMap	OSM
Part of Speech	POS
Physical, Emotional, Cognitive, and Social Factors	PECS
Support Vector Machines	SVM
Volunteered Geographic Information	VGI

Abstract

USING SOCIAL MEDIA CONTENT TO INFORM AGENT-BASED MODELS FOR HUMANITARIAN CRISIS RESPONSE

Sarah Wise, PhD

George Mason University, 2014

Dissertation Director: Andrew Crooks

Crisis response is a time-sensitive problem with multiple concurrent and interacting subprocesses, applied around the world in a wide range of contexts and with access to varying levels of resources. The movement of individuals with their shifting patterns of need and, frequently, disrupted normal support systems pose challenges to responders trying to understand what is needed, where, and when. Unfortunately, crises frequently occur in parts of the world that lack the infrastructure to respond to them and the information to inform responders where to target their efforts. In light of these challenges, researchers can make use of new data sources and technologies, combining the information products with simulation techniques to gain knowledge of the situation and to explore the various ways in which a crisis may develop. These new data sources - including social media such as Twitter and volunteered geographic information (VGI) from groups such as OpenStreetMap - can be combined with authoritative data sources in order to create rich, synthetic datasets, which may in turn be subjected to processes such as sentiment analysis and social network analysis. Further, these datasets can be transformed into information which supports powerful agent-based models (ABM). Such models can capture the behavior of heterogeneous individuals and their decision-making process, allowing researchers to explore the emergent dynamics

of crisis situations. To that end, this research explores the gathering, cleaning, and synthesis of diverse data sources as well as the information which can be extracted from such synthetic data sources. Further, the work presents a rich, behaviorally complex agent-based model of an evacuation effort. The case study deals with the 2012 Colorado Wildfires, which threatened the city of Colorado Springs and prompted the evacuation of over 28,000 persons over the course of four days. The model itself explores how a synthetic population with automatically generated synthetic social networks communicates about and responds to the developing crisis, utilizing real evacuation order information as well as a model of wildfire development to which the individual agents respond. This research contributes to the study of data synthesis, agent-based modeling, and crisis development.

Part I

Background

Chapter 1: Introduction

1.1 Motivation

A crisis has been defined as a serious threat to the basic structures or values of a social system, characterized by time sensitivity, highly uncertain circumstances, and the need for critical decisions to be made (Rosenthal et al., 1989). If handled poorly, a crisis transforms into a disaster. Crisis response is a time-sensitive problem with multiple concurrent and interacting subprocesses, applied around the world in a wide range of contexts and with access to varying levels of resources. The movement of individuals with their shifting patterns of need and, frequently, disrupted normal support systems pose challenges to responders trying to understand what is needed, where, and when. Unfortunately, crises frequently occur in parts of the world that lack the infrastructure to respond to them and the information to inform responders where to target their efforts: the earthquake in Haiti in 2010 (Heinselman & Waters, 2010) and extensive flooding in Pakistan in 2010 (Khan & Salman, 2012) occurred in regions with outdated or nonexistent geographical information and geospatial data records incapable of supporting response efforts. The spatial reality of a crisis situation can determine which individuals have been effected and how they respond and move, yet responders lack the information about road structures or the population distribution to structure their efforts.

Compounding the problem of inadequate information in general is the fact that the situation is, by its very nature, rapidly changing as a result of a number of potentially interacting subprocesses. As individuals move around within a crisis area, they may block the paths of responders, swamp aid areas as they concentrate upon them, or over-utilize resources such as clean water or sanitation systems to the point that the facilities are damaged and unserviceable, compounding the existing problems. The notion of complexity as

understood in complexity science is useful here: a complex system is one in which the overall system is not neatly decomposable into separate subprocesses, such that the whole is greater than the sum of its individual parts (Epstein & Axtell, 1996). Many social phenomena are complex, and crisis situations are no exception, as individuals react to the situation in which they find themselves and produce emergent, higher-level phenomena such as traffic, panic, information-spreading, disease-spreading, spontaneous aid points or refugee camps, riots, and so forth. A number of factors therefore influence crisis development, and it is important to understand their impact when trying to conceptualize how the situation progresses. Questions of spatiality matter in crisis situations - the distance to the nearest aid center or whether one's home was located within the part of the city that was destroyed will obviously influence an individual's behavior (Min & Hong, 2011). Human factors are also important: not all individuals will have access to the same resources or information, have the same mobility patterns or make the same decisions (O'Donnell et al., 2009). This heterogeneity is an important aspect of individual response to crisis, as are the social relationships that can both oblige and inform (Lindell, 2011; Sarcevic et al., 2012). Human behavior as it emerges from the combination of communication, individual knowledge, personal attributes, spatial location, and human relationships is an important factor in crisis situations, but one that has historically been extremely difficult to capture in all its complexity.

The combination of lack of information and lack of ability to use such information prevents responders from having a complete picture of the situation, but new technologies can address these problems. During the Mumbai terrorist attacks of 2008, there was widespread rapid generation of Flickr imagery, Twitter messages, map mashups, and wiki articles to provide real-time coverage and emerging analysis of the crisis event, creating an aggregate picture of the situation that would have been unavailable by any other means (Croitoru et al., 2014). The Haitian earthquake in particular saw the birth of a new kind of response, with widespread participation in relief activities by individuals who never set foot in the crisis area. Volunteers with access to satellite imagery of the disaster-struck areas helped create maps of the road networks (Zook et al., 2010; Harvard Humanitarian Initiative,

2011). Others set up a central phone number to which affected individuals could text requests for help (Heinzelman & Waters, 2010), at which point the texts were translated by more volunteers (Caragea et al., 2011), then geocoded and recorded so that a map of need could be created for responders. These efforts represented a new front in humanitarian response, one that promises powerful new tools to come (Biewald & Janah, 2010).

Computational social science, and in particular agent-based modeling, can be brought to bear upon the remaining questions of how the system will develop. The methodological power of agent-based models is that fundamental social structures and group behaviors are emergent from the interactions of individual agents (Epstein & Axtell, 1996). Fueled by these new sources of information, models of human behavior can be developed and used to project how individuals will learn, move, and interact with one another. Accurately capturing the dynamics of human behavior is one of the greatest challenges facing complexity and computational social science, but with the information that can be extracted from new data sources it is possible to build increasingly rich and nuanced models (Weinberger, 2011).

1.2 Research Question

The goal of this research is to develop a conceptual framework for crisis exploration. The effort includes the processing of an incoming data stream in an online fashion as well as structuring the data so that it can support an agent-based model of citizen response. Geospatial data from a variety of sources and microsimulation methods are combined with this information and used to generate synthetic population data. Once processed and cleaned, the data is fed into a spatially explicit agent-based modeling (ABM) framework. The ABM utilizes the incoming data as well as a variety of geographical information system (GIS) sources to dynamically develop an understanding of the situation and how it may develop in the short-term future.

In designing a framework that deals with such a wide range of technologies, data sources, and combinations thereof, it would be more correct to say that this thesis addresses not

one but a number of research questions. In exploring how responders and researchers can use new sources of information, the work speaks to questions such as how social network analysis, sentiment detection, crowdsourced information, and behavioral models can help in crisis response. The framework represents a body of research, developing methodologies which could potentially be implemented into a platform for use by responders in the future. At its core, the research question this project will address is this:

Can data from a wide range of sources be synthesized within an agent-based model in order to reliably and quickly project how crisis situations might evolve?

1.3 Research Contribution

The major contribution of this work is the synthesis of a range of techniques and their application to this time-sensitive, uncleaned-data-rich, complex situation. The processes of data source synthesis, social network analysis, and ABM have previously been handled separately, but there is a growing realization that the approaches are more powerful in conjunction (Crooks & Wise, 2013). Usually, ABMs of humans lack meaningful social networks; if they have social networks which parallel real-world structures, the models are usually not spatial. In this model, the spatially-explicit ABM is populated by individuals contextualized within a variety of social networks, choosing to engage with their networks in a data-derived fashion and to behave accordingly. Watts (2013) argues that this combined approach is the way forward toward answering questions of complexity, celebrating the “emerging intersection of the social and computational sciences, an intersection that includes analysis of web-scale observational data, virtual labstyle experiments, and computational modeling.”

The behaviors of the agents are not simply probabilistically determined, but dependent on a heuristic informed by their personal attributes, location, and ever-evolving set of knowledge. Thus, agents plan paths and select destinations based on the information they

have, acting depending on their personal assessments of the situation rather than flowing down roads as if they were electrons. Movement is a function of complex behaviors, allowing the model to capture previously inaccessible dynamics. In short, agents behave based on networks that resemble real social networks, informed by their meaningful knowledge, and contextualized within explicit spaces. Perhaps the key challenge in ABM at this point in its development is this understanding and operationalization of human behavior. The focus on developing high-quality, meaningful representations of human behavior in a modeling context is one of the central problems of the field (Weinberger, 2011), and this thesis attempts to contribute to that discussion.

1.4 Organization of the Dissertation

This thesis can be grouped into three overarching parts: the background material, methodology, and analysis parts provide a structure for the thesis. Broadly, the background material is presented in Chapters 2 through 4, while Chapters 5 through 7 form the methodology section. The analysis section contains Chapter 8 and Chapter 9. The chapters are explored further here.

Chapter 2 presents an overview of the principles of human behavior, with specific focus given to behavior in crisis situations. The chapter explores models of human behavior, before reviewing how others have simulated humans in crisis situations. Particular attention is given to other ABMs of crisis scenarios, with a brief discussion of some of the models that will inform and underlie the ABM developed later in Chapter 7. Overall, the chapter contributes information about the processes being explored and simulated in the rest of the thesis, contextualizing the understanding of the behaviors that are investigated in other chapters.

The combination and synthesis of new kinds of data into information to support further analysis is explored in Chapter 3. Beginning with a review of the emergence of new forms of interaction and the data these processes generate, the chapter outlines some of the ways that researchers have attempted to extract meaning from them. Building upon this, Chapter 4

explores human social networks in depth, considering the way individuals interact with others and how those relationships can be conceptualized. The chapter moves on to an in-depth study of social media networks, presenting a range of research dedicated to this emerging field, and demonstrates a worked example of the way that knowledge can be extracted from data derived from such networks, specifically a set of posts (“tweets”) on Twitter. These chapters contextualize the rest of the thesis, grounding it in the literature devoted to these various fields of research.

Moving from the background section into methodology, Chapter 5 reviews how researchers have understood and studied sentiment across a range of contexts. A simple method for extracting sentiment is developed and then tested on the set of tweets analyzed in Chapter 4. Chapter 6 presents a brief history of methods for population synthesis before presenting the method utilized in this work. Each of the datasets and subprocesses is explained and reviewed for effectiveness before ultimately creating the population of agents that supports Chapter 7.

Chapter 7 explains the case study being explored in this work before specifying the ABM that is the heart of the thesis. The specific application that is developed deals with the 2012 Colorado wildfires, specifically the Waldo Canyon wildfire, which burned from June through July of 2012 and forced more than 28,000 residents of Colorado Springs to evacuate their homes (City of Colorado Springs, 2013). The chapter specifies the data utilized to support the model and gives a full description of the processes and structures utilized by the model itself. Exploring the integration of the physical process of the wildfire with the social and behavioral processes of the agents, the chapter rounds out the methodology section of the thesis.

Chapter 8 begins the analysis section, exploring the efforts made toward the verification and validation of the model. Once it is established that the model has been constructed as designed in Chapter 7 and that the parameters of the model have the influence expected based on Chapters 2 - 4, the chapter proceeds to present the results of the model run to simulate the real-world situation, paralleling the real-life scenario that played out during the

course of the Waldo Canyon wildfire. The results of the model are compared to information drawn from a variety of sources, investigating the effectiveness of the model at capturing the emergent dynamics of the real-world system. Chapter 9 reviews the work presented, summarizing the findings and their research contributions as well as discussing potential future avenues of research.

Chapter 2: Humans in Crisis

The Nobel-award winning physicist Murray Gell-Mann is quoted as having once exclaimed “imagine how difficult physics would be if electrons could think!” (Miller & Page, 2007). In dealing with the behavior of humans it is necessary to consider not only how individuals think, but how they feel, perceive, make decisions, and so forth, with all of the complexities that limited or imperfect information and high-stress situations can bring to the question. Indeed, much of the complexity inherent in trying to understand crisis situations revolves around not only the potential threat, but the ways in which individuals might respond to it. Understanding the behavior of the individual in a particular context is a key part of understanding how a situation will develop. But what do researchers mean by “human behavior”? How does one understand behavior versus decision-making or planning or action, and how can it be modeled? How does behavior change in a crisis, and how has this specific subset of circumstances been modeled?

This chapter gives an overview of how human behavior is conceptualized and operationalized in Section 2.1 before exploring the specifics of behavior in a crisis setting in Section 2.2. Toward the end of modeling the phenomena described in the first two sections, Section 2.3 reviews current approaches to modeling human behavior. Section 2.4 briefly presents a range of models which project the development of crisis situations themselves in order to give synthetic contexts to the individuals being simulated, while also presenting a variety of methodologies that have been employed in the study of such situations. Section 2.5 specifically reviews agent-based models which bridge the gap between crisis situation and human behavior. Section 2.6 provides context for the development of the model presented in Chapter 7. Overall, the chapter contextualizes the work done in the rest of the thesis in the literature and body of work reviewed here.

2.1 Principles of Human Behavior

Understanding how humans behave is inherently difficult. What distinction exists between, for example, decision-making and behavior? On what timescale is behavior being described? To what degree is one interested in the social environment, the biomechanics of cognition, the influence of emotion, the individual’s life stage, and so forth? This question of focus depends on the purposes of the asker. But it is not enough to question how to understand the behavior of an individual; researchers describe phenomena like stampedes and riots as examples of “human behavior”, although both of those terms necessarily describe the concerted interactions of a group of people. At what level, then, of social organization are discussions of human behavior targeted?

Defining what a researcher means by human behavior is a task in and of itself. Many different theoretical perspectives on human behavior exist, and each understands behavior from behind its own lens (Hutchison, 2010). Popular perspectives and their operationalizations are summarized in Table 2.1. These definitions and theories of human behavior are drawn from and applicable to a wide range of fields, and different fields highlight different aspects of these theories to support their own purposes. Thus, for example, an economist studying human behavior might find a rational choice perspective to be more useful than a humanistic perspective for her purposes, while a social worker might care more about psychodynamic or developmental perspectives in understanding behavior. Given the importance of the specific discipline to the chosen highlighted features of behavior, it makes sense for a researcher to select a perspective which highlights the necessary components of his work. Kennedy (2012) proposes the following set of principles of human behavior for use in an agent-based modeling context, which are adopted in this work:

1. humans as processors: humans process sensory information about the environment, their current internal status, and remembered history in order to decide upon a course of action. They are influenced by personality traits and circumscribed by limited information input, memory, and processing capability. This is related to Simon’s

Table 2.1: Overview of different theoretical perspectives of behavior and the operationalization of the perspective

Perspective	Focus
Systems	outcome of reciprocal interactions of persons within linked social systems
Conflict	focus on conflict, inequality, dominance, and oppression as drivers for behavior: look at economics, politics
Rational choice	behavior based on self-interest and goal-oriented behavior; interaction is the exchange of resources
Social constructionist	focus on learning, interaction, and social, shared meaning/understandings
Psychodynamic	focus on internal processes like needs, drives, and emotions
Developmental	focus on development of behavior over a lifetime
Social behavioral (aka Social learning)	focus on behavior as it is learned by individuals in interacting with their environment
Humanistic	focus on the individual's agency, search for meaning

(1996) idea of internal and external validity.

2. motivations: humans are motivated to fulfill their needs in decreasing order of basicness, as defined by Maslow (1943).
3. rationality: a rationally behaving system must be able to represent knowledge, learn, remember new knowledge, and apply knowledge to determining the behavior of the agent (Axelrod, 1997).
4. emotional/intuitive/unconscious behavior: the basic emotions (interest, joy, happiness, sadness, anger, disgust, and fear: Izard, 2007) may either modify rational decision-making or be completely separate mental processes, and are the subject of intense debate (Scherer, 1999; Loewenstein & Lerner, 2003).

5. social behavior: human behavior is shaped by others. Dunbar (2004) notes that because humans can imagine the goals, thoughts, and feelings of others, they incorporate these projections into their own plans. Behaviors on the part of one person influence and combine with the behavior of others (Latané, 1981; Friedkin & Johnsen, 1999; Surowiecki, 2005; Kennedy & Eberhart, 2001).

These principles are incorporated into the simulation’s needs and assumptions, as discussed in Chapter 7. In this thesis, a distinction is made between decision and behavior in that the decision process is the way an individual makes choices among behavioral options, following Jager and Mosler (2007). This process can be impacted by knowledge, emotions, and norms, but the set of available behavioral responses remains the same.

It is also important to be clear on the role of emotion in human behavior. As with behavior, it is difficult to define a word as basic yet profound as “emotion”; Mulligan and Scherer (2012) note that no consensus definition exists among the various disciplines that study emotion. They suggest defining emotions as affective processes, with different individual emotions being different processes. Ortony et al. (1987) classify emotion along at least two dimensions, namely evaluation (good-vs-bad) and potency (powerful-vs-powerless). Regardless of the specific operational definition, research and personal experience both indicate that human behavior is affected by emotion. The mechanism is straightforward: emotions impact cognition, and through cognition the decision-making process (Cohn et al., 2000). This impact can be manifested, for example, through the distortion of an individual’s priorities and the production of stress (Cohn et al., 2000). Emotion, then, plays an important role in determining behavior, especially in situations like crises in which heightened emotions distort the weighting of the individual’s needs and available options.

2.2 How Do People Behave in Crisis Situations in Particular?

Human behavior in crisis situations is subject to extreme stimuli, and can differ from behavior in less extreme contexts. However, the way people move and the choices they make

are subject to specific rules even in these extremely high-stress situations. It is crucially important to capture these dynamics if a model is to accurately simulate the behavior of a population in crisis. The following is a discussion of the decision-making process individuals undergo as well as the actions individuals in such situations have been known to take.

2.2.1 The Role of Information

One of the most important drivers of an individual's behavior in a crisis situation is the information he has at his disposal. Knowing about the existence, location, and nature of a threat allows an individual to respond appropriately, while the emotional impact of the crisis varies depending on the knowledge one has of the situation. The quality of information one has seriously impacts the success of an individual's response to the crisis. Proux and Sime (1991) show that broadcasting information rather than an alarm noise results in faster evacuation times overall. Whether individuals seem to prefer more familiar exits because of their knowledge of them or the emotional feeling of safety they experience while using them is open to debate, but the trend to utilize known escape routes has been documented (Frantzich, 2001; Sime, 1985). This tendency contrasts with situations where the environment is unfamiliar and people tend to utilize the closest exist (Nilsson et al., 2009). While situational information is tremendously useful, having information about or training in dealing with a crisis is also helpful. Nilsson et al. (2009) demonstrate the significant difference in speed of response between uninformed, informed, and trained individuals in a disaster mock-up setting, a finding supported by the work of Kinateder et al. (2013). Thus, information gathering is a time-sensitive but necessary step, and one that has major ramifications for the individual's handling of the situation over time.

2.2.2 The Decision To Evacuate

Once an individual has information about the situation, she may fail to respond to the crisis and continue about her business. If, on the other hand, she recognizes the extent of the problem, she realizes that she must take action. Ripley (2008) asserts that there are three

stages through which an individual passes in understanding a disaster: denial, deliberation, and the decisive moment in which the individual makes a choice about what to do. Based on an individual's life history and experiences, one might pass through them more or less rapidly, but Ripley notes that the denial phase can extend far beyond a timeframe that bystanders find imaginable, citing cases like the diners who died at their dinner tables without ever trying to escape a 1977 fire at the Beverly Hills Supper Club. Nilsson et al. (2009) give a similar example of a fire in a road tunnel after which most of the casualties were found inside or around their escapable cars. In many cases, individuals fail to adjust their set of beliefs to realistically understand what is happening. This bifurcation can be treated as a decision individuals make about whether or not to respond to the situation.

There are two important facets of the making the choice to response to a crisis: the use of simple rules to make decisions and the importance of an individual's belief or opinion about the threat posed by the situation. The former is well-supported by the recent flourishing of heuristics. The study of heuristics has enjoyed a great deal of discussion since Simon (1957) introduced the idea that individuals use simple rules of thumb to guide their actions, a notion upon which Kahneman and Tversky (1996) built extensively. Heuristics are frequently employed in crisis simulations, and show themselves to be extremely effective (Afshar & Haghani, 2008; Liu et al., 2007). Heuristics are a realistic approximation of how humans address problems, and capture the irrationality of human behavior by making simple choices based on emotion- or stress-weighted factors and values.

While the mechanism by which the decision is made shapes the behavioral response, an individual's belief in or assessment of a crisis is perhaps the most important aspect of the decision-making process, more than any other demographic characteristic. Surveys of individuals who chose not to evacuate the New Orleans area before Hurricane Katrina show that access to transportation was a much-cited factor in their failure to evacuate. However, far more common among the population that failed to evacuate was the conviction that the storm would not be as bad as it was (Simerman et al., 2005). Half of these individuals stated that they had the option to evacuate, but that they had actively chosen not to do so.

Blendon et al. (2007) interviewed individuals living in high-risk hurricane-vulnerable areas and found that income did not predict the quarter of respondents who told the researchers that they would ignore government orders to evacuate, making situational assessment a more important factor than income in this particular case. Thompson (1986) and Aptekar (1990) describe numerous other similar cases of individuals fatally mis-assessing crisis situations. The success of an evacuation, then, depends partially on the choices made by potential evacuees, and capturing the way individuals assess a developing crisis is important for modeling a potential disaster.

2.2.3 Action

The actions an individual takes include information gathering, group interaction, and goal-seeking (Fischer, 1996). Even before an individual has made a decision about whether and how to evacuate, their actions and movement are influenced by their goals, new information, and a number of social factors.

Information gathering in crises is itself an important behavior. Drabek (1992) states that after being told to evacuate, people will check with four or more sources before making a decision about what to do. Individuals consult family, coworkers, news reports, and even strangers in the immediate vicinity to inform their decision (see Gershorn, 2007). The process is called milling, and it can significantly delay the individual's evacuation (Drabek, 1992). Individuals also glean information from their environment, although Nilsson et al. (2009) note that in their study stressed individuals failed to observe evacuation-relevant information (such as exit signs) that less-stressed individuals did observe. Thus, emotion plays into the information-gathering process as well. As mentioned above, the information one has about a situation can significantly impact the decisions one makes and the resulting behaviors one adopts, so this stage is time-consuming but important.

Another important aspect of crisis behavior is the impact of one's social group on the choices the individual members make. Vaught et al. (2000) note the tendency of miners in disaster situations to stick together, making choices as a group, and Drury (2009) notes

that even former strangers quickly band together and develop a strong sense of solidarity in the face of a disaster. Group behavior and the desire of the individual to stay with the group is a crucially important factor in crisis behavior. Ripley (2008) cites examples of individuals crawling over the seats in a crashed airplane to reunite with their families or traveling companions, blocking others in the process and at the cost of actually evacuating the plane themselves. Groups can also have a positive influence on behavior, however: there are multiple real-life examples of individuals following others who they see evacuating, therefore successfully evacuating without even becoming fully aware of the nature of the emergency (Norén & Winér, 2003). Nilsson et al. (2009) found that the most frequent reason for evacuation given by the subjects in their emergency mock-up was that the subject had seen others evacuating, and also suggested that this social influence resulted in lower levels of stress in those who were merely following without having observed the emergency themselves.

One of the most interesting aspects of behavior in crisis situations is not the action people do take but the actions they don't take. Surprisingly, many survivors report the relative calmness, politeness, and overall lack of chaos they experienced in real-life crisis scenarios. Proulx (2002) describes a "lethargic response" of individuals to a fire, even when the fire is in the same room as they are. Ripley (2008) cites multiple examples of people stopping to allow others to enter the stairway in front of them during the evacuation of the Twin Towers on 9/11, and carrying others to safety or slowing to allow firefighters to pass them more comfortably. In many documented cases, the scrambling and clawing or even rapid egress one might assume of evacuees is simply not the case and queuing up for the opportunity to escape is the order of business.

2.2.4 Emotion

In real crisis situations, emotions can understandably run high. Drabek (1986) reports that people show signs of emotional disturbance as an immediate response, beginning with anxiety and confused thinking (Aptekar, 1990) as well as fear (Cohn et al., 2000). Despite

or perhaps even because of this initial mix of emotions, reports drawn from real crisis scenarios show that people adapt and respond relatively well to the extreme stresses upon them (Drabek, 1986). Because responders typically deal with an affected population for an extended period and cannot be present at the moment of the crisis, researchers often discuss how survivors pass through emotional stages on the scale of days, weeks, or even years post-crisis (Shore, 1986; Aptekar, 1990). However, during the period of the crisis itself, emotions like anger, guilt, defensive behavior, anxiety, shame, aggressiveness, and abnormal behavior due to medical issues have been noted (Cohn et al., 2000; George, 1986).

The results of these emotions in crisis situations vary: Cohn et al. (2000) note that in crisis situations, the emotional need for safety can drive individuals to form groups, which can either help or hinder an individual’s evacuation efforts. On the other hand, as mentioned above, emotionally-derived stress can result in lower attention paid to surroundings, so that individuals fail to observe information that would be helpful to them (Nilsson et al., 2009).

As with all aspects of human experience, individual cases vary. Frederick (1980) points out that different crisis situations impact victims differently, so that the emotional response necessarily varies. This should be kept in mind, even if incorporating meaningful variation in a population into a simulation is difficult.

2.2.5 Caveats

This section has focused on a number of factors, and to highlight their importance others have been downplayed. While calm evacuations are a fascinating and frequent occurrence, panic is a fact: stampedes are widely documented everywhere from street fairs to the Hajj (Batty et al., 2003; Ripley, 2008). Quarantelli (1954) suggested that panic is a function of people feeling that they may be trapped if they do not act, combined with a sense of helplessness and isolation. These factors should also be considered, but it is important to understand that they are not necessarily a “default” behavior. Helbing et al. (2000), for example, break down the actions undertaken by individuals and create an emergent crowd/jamming behavior based on no more than the force exerted by individuals attempting

to rush through an opening, yet rush hour commuting rarely results in injuries. Thus it is important to consider the suite of behaviors and attributes which characterize an agent, and to recognize that one agent’s default might be another’s exception.

On the subject of default behavior, George (1986) notes that demographic characteristics such as sex, age, personality, ethnicity, and so forth may shape an individual’s response to a particular crisis. While some researchers report variation along, for example, gender lines (Kinatader et al., 2013) other studies belie any strong influence based upon the same characteristics in a very similar scenario (Nilsson et al., 2009). In the real-world, the apparent impact of demographics may be the result of dependent variables a child fails to evacuate a plane because of his size, not his age, or a higher percentage of business passengers are male and traveling alone so the population of survivors skews male because groups are less successful at evacuating, and so forth. It can be difficult to disentangle these dynamics from the real-world data, but disingenuous to suggest that two individuals who differ only in ethnicity or age would necessarily experience different emotions in the event of a plane crash. By attempting to capture the underlying dynamics of why an individual acts the way she does, we can more accurately capture the nuance of emotion and behavior.

Additionally, this discussion focuses on the importance of belief, behavior, and information gathering to the actions of individuals. Researchers have pointed out that in many cases, the impact of a crisis has less to do with the behavior of the individuals and more to do with the system in which they find themselves. Without systems of infrastructure in place, behavior can influence only so much. Thus, it is important to understand the interface of structure and behavior, which the model presented here strives to do.

2.3 A Brief Overview of Modeling Human Behavior

Given all these factors that play into behavior, how does one simulate behavior in general? Modeling humans is a problem with obvious applications, and has been attempted often. Researchers often break their work down by the level of decision-making being considered, specifying whether they are attempting to model behavior at the level of the individual,

the group, or the society. The former two are typically modeled as individuals - that is, the group itself is modeled in terms of its behaviors and qualities - while the last the frequently dealt with statistically (Kennedy, 2012). Kennedy (2012) breaks down the different kinds of approaches toward modeling human behavior into three categories: mathematical, cognitive frameworks, and cognitive architectures.

Mathematical approaches are the most lightweight and simple way to incorporate human behavior into models, although they have their drawbacks. Such approaches use direct and custom coding of behavior into the simulation to capture the choices, deliberation, and actions of humans. Granovetter's (1978) threshold-based rules for individuals to decide whether or not they feel comfortable participating in a riot, given the number of other individuals who are already rioting, is a simple example of this kind of behavior implementation. To give a sample case of the simplest form of mathematical approaches, Gode and Sunder's (1993) random number generator selecting among a set of predefined choices, while Hannon and Ruth's (1994) dynamic modeling is a more complicated example of the same class. The drawback of this method, of course, is that it so often depends on unrealistic randomness. Additionally, the problem space must be exceptionally narrow for the full range of behaviors to be incorporated this way, to the point that all the relevant behaviors usually cannot be represented in such a system. Simplicity is both their virtue and their flaw.

At the next level of complexity are conceptual cognitive frameworks, or frameworks of behavior implemented within a given target system. These frameworks incorporate abstract concepts such as beliefs, desires, intentions, and emotional factors, among other features, to motivate and explain agent behavior. Popular cognitive frameworks include BDI (Belief, Desires, and Intentions), PECS (Physical, Emotional, Cognitive, and Social factors), and the Fast and Frugal framework. BDI was developed by Rao and Georgeff (1991), and is implemented by transforming a decision tree into a possible worlds model, from which a deliberation process determines the best course of action. PECS (Schmidt, 2002) includes a self-model, an environmental model, memory for behavior protocols, planning, and reflection, as well as motives represented by state variables. PECS is powerful, but difficult

to train because of all of its associated variable settings. Fast and Frugal (Gigerenzer & Goldstein, 1996) harnesses the power of heuristics to develop a tree of rules based on data. It is efficient and fast, but the heuristics employed must be appropriate for the situation at hand or it is ineffective.

Perhaps the most complex approaches to capturing behavior are cognitive architectures. Cognitive architectures model the cognitive functioning of an individual at the millisecond scale (Pirolli, 1999) and are unlike the other approaches described here in that they are not rule-based. The focus is on abstract or theoretical cognition, and typically the unit of study is the individual, not groups of individuals. Two of the most prominent cognitive architectures are ACT-R and Soar. ACT-R (Anderson & Lebiere, 1998, Anderson et al., 2004) deals with symbolic and sub-symbolic level representations of knowledge, and focuses on low-level, short-term cognitive phenomena. It does not address higher-level concepts like intention or desire. Soar, developed by Lehman et al. (2006) handles symbolic-level representations of human problem-solving tasks, and is similar to a framework like BDI with regard to its internal representation of the world, state variables, and goals. Other architectures exist, but few enjoy such large support communities as ACT-R and Soar.

Kennedy (2012) asserts that two of the greatest overarching challenges in building models of humans are getting the data to support the model and eventually performing verification and validation on the resulting product. With frameworks like PECS or architectures such as ACT-R, which incorporates so much very specific data and must be calibrated precisely, it can be difficult to produce robust results. Similarly, testing the effectiveness of a model in representing behavior is extremely challenging. This is never more the case than in situations where little validation data exists, because it is difficult or unethical to test how humans behave *en vivo*. Especially in the case of crisis scenarios, it is unacceptable to put human subjects through really accurate recreations of the scenarios against which we might want to validate. Thus, it is particularly important to look for patterns in these circumstances, and to consider the behaviors described by survivors of disasters and crises in general.

2.4 Modeling a Crisis

Different models have used different techniques to address crisis situations, specified or general. A review of these models is presented in Table 2.2, while this section will present a brief overview of how the phenomena have been modeled in the past, using discrete event modeling, mathematical models, microsimulation techniques, geographic information system (GIS) techniques, and system dynamics models. A representative sample of models is presented here, as a review of all existing models is beyond the scope of this thesis. Given the different systems and structures which characterize different crisis scenarios, the models are broken down into categories in order to explore how they have collectively addressed such questions, and how successfully they have captured the necessary dynamics. All types of modeling are characterized by certain limitations as well as strengths, and these are addressed in their respective sections. Agent-based interpretations of some of these problems will be addressed in Section 2.5.

Table 2.2: A survey of models of crisis scenarios

Disaster	Type	Authors	Description
Disease	Discrete Event	Hupert et al., 2002	discrete event simulation to test varying configurations and staffing patterns of bioterrorism-response aid points
	Discrete Event	Aaby et al., 2006	geospatial optimization of points of delivery for influenza epidemic
	Mathematical	Lee et al., 2009	geospatial optimization/simulation of points of delivery of healthcare services for emergency response (epidemiology etc)
	Mathematical	Wein et al., 2003	system of differential equations describing an anthrax attack and different kinds of emergency response
	Microsimulation	Brouwers, 2005a	microsimulation model of smallpox transmission - using anonymized data and explicit spatiality
	Microsimulation	Brouwers, 2005b	microsimulation models of flood management and disease transmission - explicitly spatial with economic and social heterogeneity

Continued on next page...

Table 2.2: A survey of models of crisis scenarios

Disaster	Type	Authors	Description
Evacuation	GIS-based	Cole et al., 2005	GIS-based approaches to support evacuation of volcanic activities
	GIS-based	Cova & Church, 1997	GIS-based model of community evacuation vulnerability
	GIS-based	Zepeda & Sol, 2007	GIS-based method seeking to design evacuation routes with incomplete GIS data
	GIS-based	Zepeda et al., 2005	GIS-based approach with road network analysis for evacuation planning
	Mathematical	Cova & Johnson, 2003	optimization of network flow for lane-based evacuation of a complex road network
	Mathematical	Kim et al., 2007	network analysis of evacuation routes
	Mathematical	Kulshrestha et al., 2011	optimization of shelter locations for evacuation of unspecified disaster
	Mathematical	Saadatseresht et al., 2009	optimization of spatially-informed evacuation combining GIS with a variety of equation optimizations

Continued on next page...

Table 2.2: A survey of models of crisis scenarios

Disaster	Type	Authors	Description
Logistics	Mathematical	Balcik & Beamon, 2008	optimization (linear and dynamic programming) of quantities and locations of stockpiled distribution centers for relief items
	Mathematical	Ozdamar et al., 2004	optimization of vehicle routing for efficiently dispatching relief supplies to a community with a rapid-onset disaster
	System Dynamics	Cuervo et al., 2010	system dynamics model of supply chain in disaster response
	System Dynamics	Hoard et al., 2005	system dynamics model of hospital surge and mass casualties in rural settings
	System Dynamics	Min & Hong, 2011	system dynamics model of transportation of relief goods to disaster area
Manmade	Mathematical	Dombroski et al., 2006	dispersion model of response to a “dirty bomb”
	Mathematical	Feng & Keller, 2006	optimization of distribution of iodine tablets after nuclear accident

Continued on next page...

Table 2.2: A survey of models of crisis scenarios

Disaster	Type	Authors	Description
	Mathematical	Georgopoulos et al., 2004	simulation of exposure of health care workers to hazardous materials from victims of chemical disaster
	Mathematical	Papazoglou & Christou, 1997	optimization of emergency response policies for nuclear accidents given health effects and costs
Natural	Discrete Event	Paul et al., 2006	simulation (discrete event) of post-earthquake patient surge at regional hospitals
	Mathematical	Barbarosoglu & Arda, 2002	optimization (stochastic programming) of flow of relief supplies in a post-earthquake transportation system
	Mathematical	Regnier, 2008	simulation (Markov model) of hurricane prediction accuracy and lead time for evacuations
	Mathematical	Shim et al., 2002	spatial decision support system of flood behavior with integrated model of water levels
	System Dynamics	Fawcett & Oliveira, 2000	system-dynamics approach to post-earthquake casualty treatment
	System Dynamics	Su et al., 2009	discrete event simulation and system dynamics approaches to disaster response effort post-quake

2.4.1 Disease

Modeling disease has been a focus of research efforts for some time: Bailey presented a review of different mathematical models of the spread of disease in his book in 1975. Lee et al. (2009) and Wein et al. (2003) adopt mathematical approaches, using integer programming and optimization methods to attempt to simulate the spread of a generalized bioattack and anthrax respectively. Hupert et al. (2002) and Aaby et al. (2006) attack similar problems with discrete event simulation, studying the flow of patients in either a bioattack or an influenza situation. In some of these cases, efforts are made to incorporate spatiality and heterogeneity - for example, Wein et al. (2003) attempt to determine the exposure dose based on location, and to modulate the probability of infection based on individual characteristics - but features such as movement through space, emotion, and decision-making are completely absent. Brouwers (2005a) considers the question of disease transmission with a microsimulation model, improving upon the heterogeneity and spatiality questions - her MicroPox model represents individuals with personal attributes who exist in different locations. However, the state of being in transit is represented as being a place in and of itself, and their behavioral range lacks any consideration of knowledge, emotion, or reactivity beyond being prevented by illness from going to work. These models represent important steps toward understanding the dynamics of the spread of disease under a range of circumstances, but the difficulty they have in capturing space or behavior prevents them from explaining much of what makes disease such a complex system (Epstein, 2009).

2.4.2 Evacuation

Researchers who model evacuation tend to utilize either mathematical modeling or geographic information systems approaches to the system to attempt to understand the flow of people and vehicles through the environment. Network and optimization-oriented approaches are typical in the population of mathematical efforts (see Cova & Johnson, 2003; Kim et al., 2007; Kulshrestha et al., 2011, Saadatseresht et al., 2009). Examples of the latter category of GIS models include: Cole et al. (2005), Cova and Church (1997), Zepeda

et al., (2005), and Zepeda and Sol (2007), all of whom seek to use GIS to explore vulnerable populations and road networks. Cole et al. (2005) compare the location of threats with population centers and the road network, suggesting that vulnerability can be easily explored as a function of distance. Zepeda and Sol (2005) investigate path planning along realistic road networks, while Cova and Church (1997) attempt to use the road network to identify areas of vulnerability. The mathematical approaches tend to focus on the structure of the road network (e.g. Cova & Johnson, 2003; Kim et al., 2007) and sometimes on spatial questions such as the optimal position of shelters (see Kulshrestha et al., 2011). None of these models are able to incorporate the imperfect information to which individuals might have access, nor can they capture the communication dynamics that might influence the time at which individuals decide to evacuate or their evacuation destination selection. Zepeda and Sol (2005) do allow for a heuristic-based route selection process, but even then there is no heterogeneity among individuals. Behavior remains an elusive factor in these models.

2.4.3 Logistics

In many cases, researchers have developed models to deal generally with crises of an unspecified nature, focusing on the logistics of the response effort. Balcik and Beamon (2008) utilize linear and dynamic programming to attempt to optimize the positioning of distribution centers, while Ozdamar et al. (2004) take a similarly mathematical approach to the question of routing relief vehicles in crisis situations. These methods are obviously inherently concerned with spatial questions, but they rely on datasets which may not exist or may have grown quite inaccurate in the face of the rapid-onset disasters the models seek to ameliorate. Incorporating the dynamic traffic patterns that might emerge in a crisis situation into such systems would be impossible, robbing them of some of their power. System dynamics models are another popular approach to logistical planning, with Cuervo et al. (2010) and Min and Hong (2011) both modeling relief supply chains. Again, space is a problem, as is the heterogeneity of the individuals affected by the crisis - system dynamics

models can account for neither, making efforts to simulate the transport of goods through space and the emerging demand of the population for relief effectively impossible. Hoard et al. (2005) utilize system dynamics models to explore post-disaster surges in the number of patients admitted to hospitals, modeling the different parts of the hospital as sinks and the patients as flows. Again, the heterogeneity of individuals poses a challenge to the effectiveness of such models at capturing the dynamics of patient hospital usage. These methods provide guidelines, but their very construction prevents them from capturing important spatial dynamics.

2.4.4 Manmade

As with logistical models, models of manmade disasters have also frequently focused on questions of distribution and the optimization of response efforts. Some of the models are extremely abstract: Papazoglou and Christou (1997) carry out an optimization based on a variety of goal metrics which include little heterogeneity or spatiality. Some of these models attempt to incorporate human behavior into the systems they study: Dombroski et al. (2006) focus on behavioral factors in their simulation of a post-terror attack scenario. In general, incorporating the way different attributes respond individual response to crisis and how those responses translate into the situation being managed hampers the ability of such models to project the dynamics of such situations.

2.4.5 Natural

One of the most popular types of crisis modeling involves simulating specific natural disasters. Many efforts to model natural disasters have dealt with attempting to optimize the distribution of resources to relief efforts (e.g. Paul et al., 2006; Barbarosuglu & Arda, 2009). These efforts tend not to include behavior or meaningful spatiality, features which are also missing from the system dynamics-type models (e.g. Fawcett & Oliverira, 2000; Su et al., 2009). While these models have the benefit of being able to tailor their models based on their better understanding the processes by which individuals or infrastructures

are threatened, they frequently are forced to omit important dimensions of the crisis. Given the importance of emotion and information to the behavior of individuals in the face of a crisis, the inability of these methods to incorporate those aspects of experience is limiting.

Table 2.3: A survey of agent-based models of crisis scenarios

Disaster	Authors	Description
Disease	Eubank et al., 2004	the EpiSims platform, exploring graphs showing contact patterns between individuals
	Muller et al. 2004	model of different types of disease-bearing agents interacting
Evacuation	de Silva & Eglese, 2000	network evacuation model underlying a spatial decision support system (also in de Silva, 2001)
	Elmitiny et al., 2007	model at the individual driver level of evacuation of a city, with emphasis on studying best evacuation policies
	Epstein et al., 2011	combination ABM/fluid dynamics model deals with hypothetical aerosol release to determine effectiveness of different response policies
	Jha et al., 2004	used to simulate low-level evacuation traffic dynamics
General	Chen et al., 2010	model of emergency behavior based on GIS and spatial information

Continued on next page...

Table 2.3: A survey of agent-based models of crisis scenarios

Disaster	Authors	Description
Natural	Chen, 2008	focuses on different evacuation strategies before a hurricane
	Crooks & Wise, 2013	ABM at the individual level, simulating the distribution of aid post-quake in Haiti
	Dawson et al., 2011	ABM of flood and evacuation behavior used to estimate vulnerability
	Sabino et al., 2008	ABM with explicit spatial data for validating emergency plans
	Zhang et al., 2009	ABM modeling evacuation of heterogeneous households before a hurricane

2.5 Agent-Based Modeling

Agent-based modeling as a discipline has the ability to incorporate a variety of dimensions of experience, contextualizing heterogeneous agents in environments that are both physical and social (Heppenstall et al., 2012). Agents can have attributes which influence their appraisal of risk, and their decision-making processes can produce a rich suite of behaviors which produce a range of emergent phenomena such as traffic and higher-level information-sharing. This section will present a few agentized simulations which deal with the range of research areas explored in Section 2.4. Section 2.5.1 will present how agent-based models have deal with pedestrian and vehicular movement in general, providing context for the exploration of agent-based models which deal with evacuation in Section 2.5.2.

Table 2.3 presents a selection of the existing agent-based models which deal with crisis situations. Simulating disease is a task to which ABM is particularly suited, given its spatiality and easy implementation of heterogeneity - EpiSims (Eubank et al., 2004) is an ABM which models synthetic individuals carrying out daily courses of action within diverse regions. The framework supports the simulation of a variety of kinds of outbreak, and allows for different intervention efforts to be explored and compared. Muller et al. (2004) present a model of sleeping sickness, focusing on the spatial dimension of the spread of disease and the interactions between different kinds of disease-bearing and -susceptible agents. In addition to disease, ABM methods have been applied to natural disaster situations as well (e.g. Chen, 2008; Dawson et al, 2011; Sabino et al., 2008; Zhang et al., 2009). Some of these focus on the behavior of individuals: Chen (2008) considers how different individuals pursue evacuation strategies in the face of an oncoming hurricane. Zhang et al. (2009) simulate roughly the same process at a household level, while Sabino et al. (2008) simulate at the level of the individual but assess their plans at aggregate levels of the population. Generalizing from specific natural disasters, many researchers have studied the question of emergency behavior (see Chen et al., 2010). Because agents have the ability to respond to their environments - and to one another - the models present a much more powerful range

of behaviors and interactions, enriching the resulting dynamics compared to the models presented in Section 2.4. Evacuation in particular has been a source of interest, and a wide range of models address the process specifically.

2.5.1 Agent-Based Traffic and Pedestrian Models

Models of traffic and pedestrians have existed for a long period of time: Fruin (1971) introduced the idea of the flow of individuals through space based on the density of people. Historically, questions of speed of movement have been addressed through aggregate-level methods - network flow models (e.g. Cova & Johnson, 2003) and fluid-dynamics models (e.g. Helbing, 1996) which treat individuals essentially as if they were particles. Many researchers have developed ABMs of traffic and pedestrian movement, including the simple version built into the NetLogo suite of examples (“Traffic Grid” - Wilensky, 2003; see also Banos et al., 2005) and extensions upon it (the StarLogo platform: see Batty et al., 1998). The SWARM modeling system is also popular (see Batty, 2003; Batty et al., 2003). For a comprehensive review of pedestrian models, readers are referred to Johansson and Kretz (2012).

2.5.2 Agent-Based Modeling of Evacuation

Many ABMs have dealt with the question of evacuation in general (e.g. de Silva & Eglese, 2000; Elmitiny et al, 2007; Epstein et al., 2011; Jha et al., 2004). There is a range in the level at which the individual is modeled - Jha et al. (2004) focus on very low-level traffic dynamics while Epstein et al. (2011) adopt a hybrid approach which draws from fluid dynamics to explore evacuation questions. Spatiality is key in these situations, as demonstrated by Sabino et al. (2008) in their investigation on evacuation when a dam is broken. It is also important to draw the distinction between evacuation models which simulate individual people and models which simulate vehicles, as there can be very different dynamics.

Researchers who deal with pedestrians frequently deal with trying to predict pedestrian

evacuation from a building (e.g. Castle, 2007; Castle & Longley, 2008; Okazaki & Matsushita, 1993; Kerridge et al., 2001). Other researchers specify their pedestrian evacuation spaces even further, specifically dealing with stadiums (Samuelson et al., 2008), the insides of aircraft (e.g. Sharma et al., 2008) or subway platforms (e.g. Hoffmann et al., 1998). The scale at which behavior and movement are simulated in these model contrast with simulations that specify the movement of vehicles over kilometres: the RedfishGroup explores how cars interact to cause traffic congestion during a wildfire-driven evacuation (Throp et al., 2006). Similar models deal not with evacuation per se but mass pedestrian movement in a crisis context - Crooks and Wise (2013) present an agent-based investigation of utilization of aid centers in the aftermath of a disaster. The model in Chapter 7 draws from all of these works, especially the evacuation models, and uses them to inform its structure. The spatiality, heterogeneity, and interactions among individuals implicit in these models allows them to capture dynamics that other methodologies cannot. The feedbacks and path dependency of crisis systems make these processes crucial to the development of the situation, so that ABM's ability to explain these dynamics is powerful and important indeed.

2.6 Wildfires

This section will explore a sample of agent-based models which simulate wildfires, which will inform the implementation of a wildfire model presented in Chapter 7. Forest fires have been the subject of a great deal of research, as their rapid development and unpredictability make them a formidable threat to responders. To that end, a number of fire modeling systems have been developed. A range of approaches to modeling fires exist, depending on the specific goal of the researcher and which aspects of a fire she wants to model. Indeed, the question deals with a number of complex processes occurring at a range of scales (Séro-Guillaume et al, 2008). Stratton (2006) distinguishes between fire models and fire modeling systems, the latter of which he classifies as interconnected sets of empirical and deterministic models equations which predict fire growth and behavior. As an example of a fire system

model, the FARSITE simulator (Finney, 1998) integrates the output of a surface fire spread model, a crown fire initiation model, a crown fire spread model, and a dead fuel moisture model (Stratton, 2006). It does not include measures of fire effects, however, nor does it include a smoke model. Certain combinations of submodels are very popular: frequently fire spread models are paired with atmospheric models (e.g. Filippi et al, 2011) to capture the feedback processes which can drive local fire behavior. Some popular fire modeling tools include BehavePlus (Andrews et al., 2005), NEXUS (Scott, 1999), FVS/FFE (Reinhardt & Crookston, 2003), FIRETEC (Linn et al., 2002), WFDS (Mell et al., 2007), Prometheus (Tymstra et al., 2010), and the Forest Service Fire Behavior Predictor (Hirsch, 2003).

The simulations presented above can be classified into a number of categories. Margerit and Séro-Guillaume (2002) and Pastor et al. (2003) both classify models as addressing either the propagation of the fire front or on the spreading process itself. Margerit and Séro-Guillaume (2002) further distinguish within the categories: among physical models, general diffusion models are different from simulations that deal with the combustion of specific types of vegetal matter, while propagation models can take a geometrical “envelop [sic]” approach, a semi-empirical approach which approximates local energy dynamics, or a cellular-automata based model. Balbi et al. (2009) break down propagation models into five categories which are largely similar, for example omitting envelope models and adding their own proposed category. Mell et al. (2005) and Pastor et al. (2003) agree that there are only empirical, semi-empirical, and physics-based models. Filippi et al. (2011) note that that realistic physical models are extraordinarily computationally expensive.

Frequently with more advanced fire models, a range of submodels are combined to form the fire modeling system. Pastor et al. (2003) identify surface fire, crown fire, spotting, and ground fire models as being specific fire models, while other submodels addressing things like atmospheric propagation or fuel moisture are frequently also incorporated. Some prominent submodels documented in the literature include models of surface fire spread (e.g. Albini, 1979), crown fire spread (e.g. van Wagner, 1977), spotting (the process by which burning material is transferred by wind to other areas, as in Albini, 1979), point-source fire

acceleration (Forestry Canada Fire Danger Group, 1992), and fuel moisture (Nelson, 2000). Cellular automata models in particular have been a very popular way to combine these models and capture the interactions among the various dynamics (see Albinet et al., 1986; Alexandridis et al., 2008; Berjak & Hearne, 2002; Hernández Encinas et al., 2007a; 2007b; Sullivan & Knight, 2008; Yassemi et al., 2008). In determining the most important factors in the spread of fires, Alexandridis et al. (2008) cite the work of Fons (1946), who claims that the factors that most affect the rate of spread and shape of a forest fire front are the type of vegetation, humidity, wind speed and direction, physical topography (e.g. slope and natural barriers), vegetation thickness, and spotting. Given the effectiveness and simplicity of the model of Alexandridis et al. (2008), it is selected and adapted to the purposes of this framework in a process detailed in Chapter 7.

Given Filippi et al.’s (2011) observation about the difficulties involved in adequately capturing nuanced atmospheric interaction, cellular automata offer a commonly accepted, light-weight, reasonable approximation of the processes in question. Further, in situations where the framework presented here could be validated, the data representing the real-world position of the wildfire could ideally be updated in a dynamic fashion, making it less necessary to focus on capturing the development of the fire in its every detail. The increasing availability of rich datasets allows researchers both to create rich models and to validate existing models against them, easing the process of development significantly.

2.7 Summary

The literature presented in this chapter informs human behavior in crisis situations and the modeling thereof. A wide range of crisis simulations was reviewed to provide a sense of the context in which this work exists, with specific care given to the way movement and behavior is understood. In general, few of the existing models of crisis situations even have the capacity to include features of human behavior - Section 2.4 details how network flow models, system dynamics models, and mathematical models fail to distinguish between individuals,

let alone incorporate heterogeneous behavior or needs into their construction. Section 2.5 shows that even among agent-based models that have the capacity to modulate behavior by personal attributes or information, many fail to do so, making their primary advantage over other methodologies their inclusion of spatiality into the system. However, given the importance of these features of the individual to her choices and actions, they significantly influence her movement through space, as detailed in Section 2.2. These considerations inform the framework developed here, as will be discussed further in Chapter 7.

Chapter 3: Geospatial Information Systems and the Evolution of Data

Recent technological developments and the increasingly accessible internet have resulted in the flowering of new forms of data. These types of data provide us with a rich body of materials which can be studied, independently or in concert, and mined for forms of information that not only were inaccessible to researchers but completely non-existent a decade ago. The emergence of new forms of interaction and data generation follows the development of the internet from what researchers have dubbed Web 1.0 to Web 2.0, a construction and distinction that is expounded upon in Section 3.1 below. Section 3.2 lays out how, as a part of this transition, the study of big data has come into its own, requiring the development of new paradigms and techniques to sift through and structure unprecedented amounts of information. These new ways of interaction have also made possible other emergent trends such as the use and study of volunteered geographic information (VGI: Goodchild, 2007) and ambient geographic information (AGI: Stefanidis et al., 2013), which will be described in Section 3.3. The challenges associated with ensuring data quality and gathering these new sources of information are elaborated upon in Section 3.4. Overall, the synthesis of these new kinds of information can be used to inform simulations and to understand aspects of geography and social interaction that would otherwise be opaque to researchers (e.g. Crooks & Wise, 2013). They represent a promising new field of study, one that will only continue to expand and develop in time. Section 3.5 summarizes some of the ways that this information can currently be used and how researchers can profit from doing so.

3.1 The Rise of Web 2.0

Officially, the term “Web 2.0” was coined in 2004 by Dale Dougherty, a vice-president of O’Reilly Media Inc. (O’Reilly, 2005). The term was meant to distinguish the new incarnation of the web and the internet as it had been during the dot-com boom of the 1990s and its subsequent bust. Although the differences in the approach to connectivity were amorphous and pluralistic, a few key aspects of the web that emerged from the bust help to explain what is meant by Web 2.0. Anderson (2007) defines the core functions of Web 2.0 as follows: individual production and user generated content; harnessing the power of the crowd; data on an epic scale; architectures of participation; network effects; and openness. The trend toward providing online services and using the web as a platform itself has also allowed for greater accessibility, a development which has had serious consequences when it comes to the use of social media, the feasibility of crowdsourcing, and the use and collection of all kinds of geosocial data, phenomena which will be discussed extensively in the following sections. Geosocial data here is information which is both spatial and social, situated at the intersection of the two spaces. In general, social media and an ever-expanding range of applications which particularly support interactivity and user-generated content will be addressed in this section, to give a sense of the kinds of services being offered and the ways that they are being utilized.

3.1.1 Applications

Some of the major applications that have arisen out of the Web 2.0 movement are Facebook, Flickr, MySpace, YouTube, Twitter, and Wikipedia (Crooks et al., 2014). In general, services including blogging platforms of various types, wikis, multimedia sharing services, RSS feeds and other tagging-based syndications, and podcasting are all examples of this new kind of engagement (Anderson, 2007). Such applications allow users to explore options when it comes to everything from the available eateries in their general vicinity (UrbanSpoon - <http://www.urbanspoon.com/>, Zagat - <http://www.zagat.com/>, or OpenTable - <http://www.opentable.com/>, to name a few), to organizing ride-sharing for commuters

(Carma: <http://car.ma/>), to submitting requests for non-emergency city maintenance work (SeeClickFix: <http://seeclickfix.com/>). These examples reflect only a fraction of the growing and diversifying set of available applications.

A development of particular interest and relevance to this work is the intersection between applications and geospatial information. There has recently been an emergence of location-aware apps which tailor the services they provide to the user's location. These apps, known as location-based services (LBS; Küpper, 2005), can be used to find information about nearby businesses, amenities of various types, and deals associated with nearby facilities. While these are explicitly spatially aware, other sources of information incorporate spatiality without making it the focus of the application: for example, Flickr and Facebook allow users to tag their locations when posting content. These questions will be explored more fully in Section 3.3, but they reflect an important growing trend: the personalization and tailoring of information.

This trend is particularly apparent in the realm of online mapping platforms, which increasingly provide services rather than software. This packaging of information into specialized formats has led to the creation of all kinds of customizable maps, many of which build upon other projects which are themselves services as well as platforms (Crooks et al., 2014). For example, the OpenStreetMap (OSM) platform synthesizes and makes accessible user-provided information about roads, footpaths, cycle paths, rail lines, buildings, and many other geographic features, all in an open format (<http://www.openstreetmap.org/>). This information is frequently utilized as a part of other platforms like CartoDB (<http://cartodb.com/>) or MapBox (<https://www.mapbox.com/>), the latter of which is used to power mapping features on applications including foursquare, Pinterest, and GitHub. Frequently the information that is mapped has social or cultural significance, synthesizing space and society in new ways - MapTales (<http://maptales.es/>) allows users to embed personal stories into their maps, while Soundcities (<http://www.soundcities.com/>) allows users to record and upload sounds from different areas and associate them with mapped locations, creating a kind of auditory map. Such custom maps are a radical departure from the first generation

of mapping technologies made available on the internet, whose functionality was limited to zooming in and out of static tiles (see Crooks et al., 2014, for a discussion of these early mapping facilities).

3.1.2 Social Media

As interactivity has led to more complex, customized, and diverse forms of geospatial information sharing on the internet, it has led to the explosion of new kinds of social interaction and sharing. Croitoru et al. (2014) specifically identify blogs and microblogs such as Blogger, WordPress, Twitter, Tumblr, and Weibo as examples of the kind of virtual communities and sharing technologies known as social media. They also include social network services such as Facebook, Google+, and LinkedIn as well as multimedia content sharing services like Flickr, YouTube, Vine, Vimeo, and others into the category of social media platforms. In general, the purpose of a social media platform is to enable nonspecialist members of the general public to contribute, upload, disseminate, and exchange information via the platform in question (Kaplan and Haenlein, 2010).

As a result of its explosion in popularity and influence, social media has been the subject of a great deal of research and exploration. Chapter 4 will address some of the research done into how individuals interact and share information over these platforms; social media has been used to try to track sentiment (see Kouloumpis et al., 2011), the emergence and modification of human social networks (see Glasgow et al., 2012), the emergence of news and social phenomena (see Stefanidis et al., 2013b; Sakaki et al., 2010; Brownstein et al. 2008), and spatiality and movement in general (see Crooks et al., 2012; Vieweg et al., 2010). The study of social media data poses challenges to researchers attempting to address technical questions of storage, processing, and data harvesting, as well as theoretical questions: these issues will be addressed in the remainder of this chapter.

3.2 What is Big Data?

As with the term Web 2.0, a precise definition of the term “big data” has been the subject of some debate (TechAmerica, 2012). The consensus has generally settled around a few properties which distinguish big data as a field, namely data volume, velocity, and variety (Croitoru et al., 2014). While the size of the datasets is one way in which the data is clearly “big”, velocity and variety are less obvious qualities. Velocity is used to indicate the rate at which data is produced as well as the speed with which an analytical process must operate to generate meaningful information about that data stream (TechAmerica, 2012). Thus, a methodology which identifies clusters of conversations within social media posts but does so months after the product of analysis would have been useful does not fall within the constraints of big data. Likewise, variety implies the range of sources and types of data which are processed. Croitoru et al. (2014) note that many big data efforts combine information from across a range of platforms; researchers might synthesize texts, images, videos, and raw geospatial information to gain a sense of the development of a protest in near-real time, for example.

Big data methodologies have the potential to mine rich combinations of information sources, but they also face a number of challenges specific to the field. The challenges associated with warehousing and processing historically unprecedented amounts of information are inherent to the problem, but frequently the way the information itself is formatted complicates the matter - social media data tends to be unstructured or ill-defined (Sahito et al., 2011). Developing tools toward automation requires that researchers consider and deal with these complications, often by treating the problem as one of data-cleaning (Rahm & Do, 2000). A number of researchers have addressed these questions of unifying different data sources and analyzing the resulting information (e.g. Shimojo et al., 2010; Kalnikaite et al., 2010; Lyons & Lessard, 2012).

3.2.1 Gathering Big Data

Before the analysis can begin, however, there is an even more fundamental question: how can researchers gather big data in the first place? Social media services frequently offer diverse platforms via a range of technologies - Facebook has versions of its service tailored for computers, mobile phones, and iPads and other tablets. Twitter users can utilize a streamlined, highly-tailored dashboard application to manage multiple accounts at once (see TweetDeck, <https://about.twitter.com/products/tweetdeck>), but they also have the capability to post messages from their non-smartphone mobile devices. Trying to harvest information in this context is obviously something of a challenge (see Croitoru et al., 2014; Stefanidis et al, 2013). Some services provide application programming interfaces (APIs) which provide users with the ability to query the content of the services - Twitter and Flickr both have this capability. However, even these formal means of accessing the information stored within a service have their problems: Croitoru et al. (2014) point out that the Twitter API can return the same information in response to different queries, yet return that content in different formats. Transforming this information into knowledge requires a tremendous amount of automation and a rigorous conceptual data model.

3.2.2 Crowdsourcing

One fascinating source of big data is crowdsourcing, a method of data creation wherein a large group of users without central organization generate content which is accessible and shareable as a web-based service (Howe, 2006). The efficacy of this technique relates to the idea of the Wisdom of the Crowd, suggested by Surowiecki (2005), which suggests that the aggregated result of pieces of information may be more powerful than the sum of the individual pieces. Crowdsourcing has been utilized in a variety of contexts, to do everything from help to classify galaxies (Galaxy Zoo: <http://www.galaxyzoo.org/>) to select the next t-shirt design a website will produce (Threadless: <https://www.threadless.com/>). Platforms to facilitate the crowdsourcing of specific tasks exist: Amazon's Mechanical Turk (<https://www.mturk.com/>) is one such example. These platforms harness the capabilities

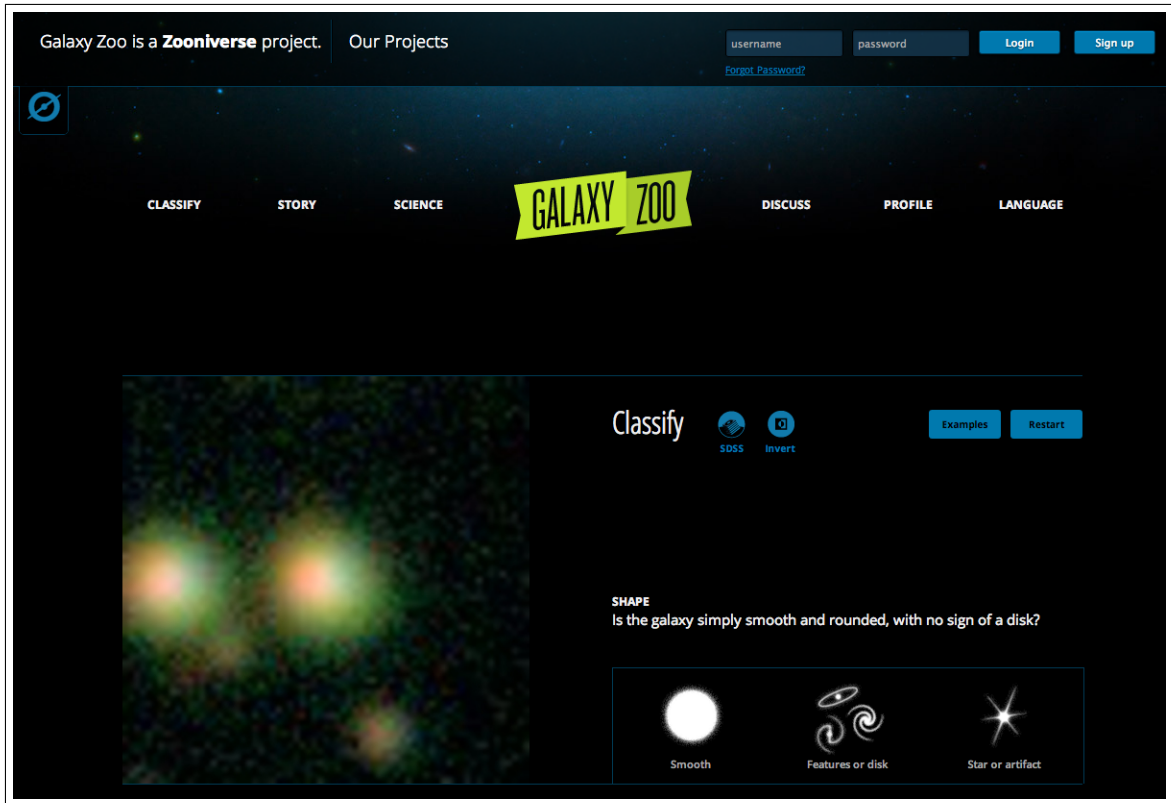


Figure 3.1: The GalaxyZoo interface, with its user-friendly set of classification questions and simple interface

of a huge range of individuals in order to achieve complicated tasks, and represent a form of data production - and analysis - that was previously prohibitively complicated. A particularly interesting result of this flowering of crowdsourcing platforms and technologies has been the emergence of crowdsourced geospatial data generation (see Hudson-Smith et al., 2009a), a topic which will be discussed in Section 3.3. To give a sense of the feel of crowdsourcing sites, Figure 3.1 shows an example of a crowdsourcing website interface, with its simple interface and an easily partitioned task for the user to complete.

3.3 Big Geosocial Data

As mentioned in Section 3.1, the blossoming of Web 2.0 technologies has led to a renaissance in the way geospatial information is created, shared, and distributed. Haklay et al. (2008) term the emerging body of geospatial information the GeoWeb, an appellation Turner and Forrest (2008) use to mean an “interconnected, online digital network of discoverable geospatial documents, databases and services”. Geospatial datasets have been large in terms of volume from the beginning of their existence (Crooks et al., 2014), and recent trends in data collection have taken this maxim to an extreme. Digital Globe generates 1 to 2 PB of data each year, on top of their 30 PB of archived data, while NASA generates 5 TB of data daily (Vatsavai and Bhaduri, 2013). Google Earth currently maintains 20 PB of imagery at a wide range of resolutions, and their store is only growing as they seek to expand into offering Street View of new areas (up to and including the Great Barrier Reef (<https://www.google.com/maps/views/streetview/oceans>)). Geospatial datasets are therefore “big” in the most obvious sense.

However, not all geospatial datasets are exclusive spatial: many can productively be subjected to a mix of spatial and social analysis (Croitoru et al., 2014). As geospatial information has moved away from being the domain of the specialist and toward the toolbox of the amateur, individuals have been able to incorporate geospatial information into their daily lives in rich, textured ways, allowing ordinary citizens to communicate and share information about their lived spaces (Sui & Goodchild, 2011). Crooks et al. (2014) present a history of the development of the GeoWeb, beginning with static map interfaces with little user interaction and ending with “Digital Earths” such as Google Earth, NASA World Wind, Bing Maps, and ESRI’s ArcGIS Explorer. They call this final generation of geospatial data service a “geo-browser”, and emphasize its continued development alongside and in parallel with the preceding generation, that of map mashups. While Digital Earths and geo-browsers seek to present accurate, current information regarding spatial relationships and the physical world, mashups utilize user-friendly APIs such as that of Google Maps to

produce highly specialized maps of consumable content without requiring the mapmaker to be highly geospatially trained (Haklay et al., 2008).

Because of both the availability of rich, extensive geospatial information from Digital Earths and the power of tools like the web mapping platform OpenLayers combined with map servers like GeoServer and MapServer, mapping is increasingly accessible to the general public and to researchers alike (see Longley & Singleton, 2009). In the context of this thesis, it is important to note that mashups are increasingly generated to help responders process information in the aftermath of hurricanes (Miller, 2006), floods (e.g. Hudson-Smith et al., 2009b), or earthquakes (Zook et al., 2010). Organizations like MapAction now deploy specialist mappers to disaster areas in order to create customized geospatial maps for responders of all types. These cartographers, and others, increasingly make use of datasets generated not by traditional methods of top-down managed digitization as practiced by the US Geological Survey or the UK’s Ordnance Survey, but by volunteered geographic information.

3.3.1 Volunteered Geographic Information

Volunteered geographic information (VGI) is the process of citizens actively collecting and contributing geospatial information (Goodchild, 2007) using Web 2.0 technologies. VGI has been made possible by increasingly cheap tools (including GPS units and open-source platforms for digitization) and has resulted in a tremendous reduction in the cost of collection and compilation of data (Crooks et al., 2014). Much has been made of the change in the way information is used now that big data processing is affordable (TechAmerica, 2012), and VGI is a perfect example of this. Longley et al. (2010) point out that in the past, data capture costs often accounted for 85% of a project’s overall costs; thanks to citizens engaging with spatial data as active creators and contributors, this cost shrinks away.

As a result of the increasing availability of user-contributed information, many organizations have harnessed the power of volunteers to rapidly create huge quantities of geospatial data. By crowdsourcing the digitization of the road network in Haiti after the

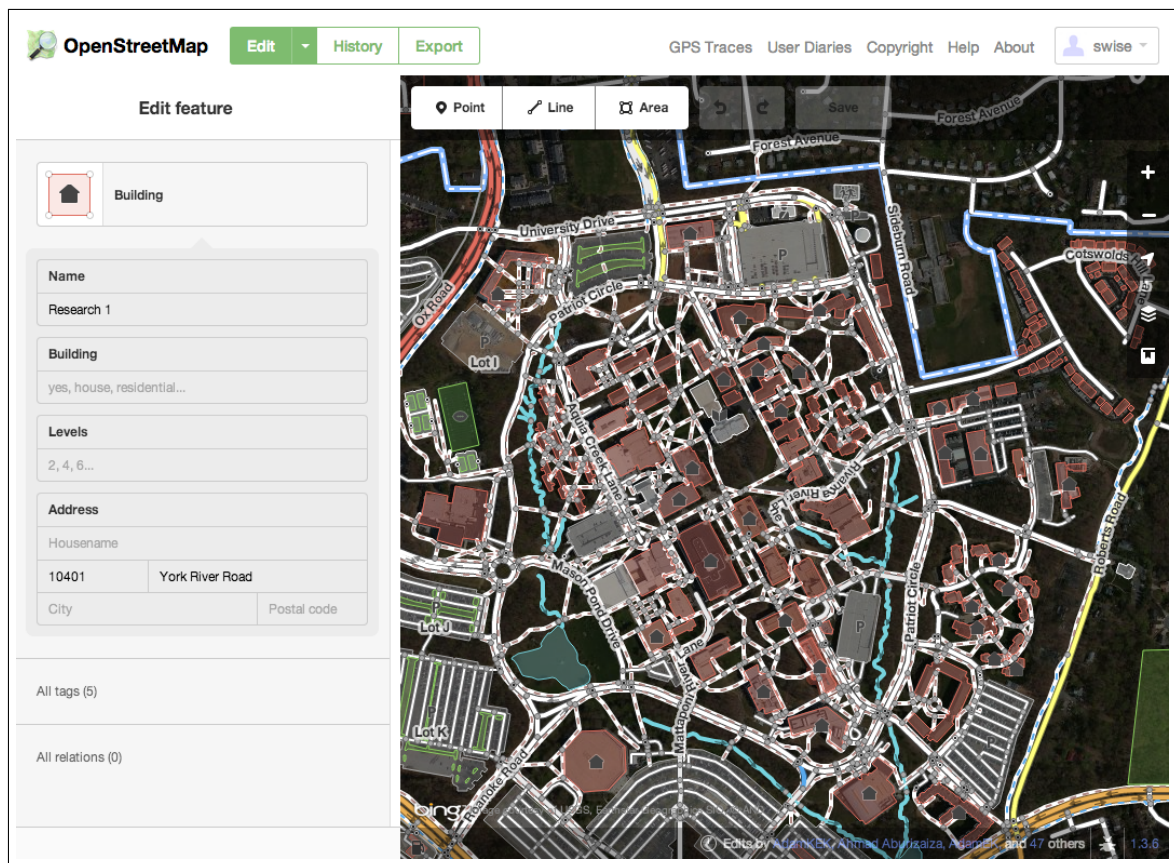


Figure 3.2: The OpenStreetMap iD editor, displaying a well-labeled area

2010 earthquake, the Humanitarian OpenStreetMap Team was able to arm responders with geographic information to guide their efforts (ITO World, 2010). Groups such as Occupy Sandy utilized citizen-tagging efforts to guide others attempting to provide relief after Hurricane Sandy (<http://occupysandy.net/map/>), while the group CrisisCommons synthesized a variety of geosocial information in the service of their Haiti response efforts (http://wiki.crisiscommons.eu/wiki/Haiti/2010_Earthquake#Maps). This type of participatory humanitarian mapping is a new form of engagement, and it relies on the provision of Web 2.0-style tools which can users can learn to use quickly and effectively, and without being professionally trained to do so (Crooks et al., 2014; Hudson-Smith et al., 2009a). Figure 3.2 shows the interface to the OpenStreetMap editing window, which opens in-browser

and allows users to tag roads, buildings, areas, and so forth, or to add metadata to the existing geometries. The role this information has played in post-disaster response and recovery has been increasingly appreciated, and the demand for this kind of data is likely to grow in the near future (Norheim-Hagtun & Meier, 2010; Zook et al., 2010).

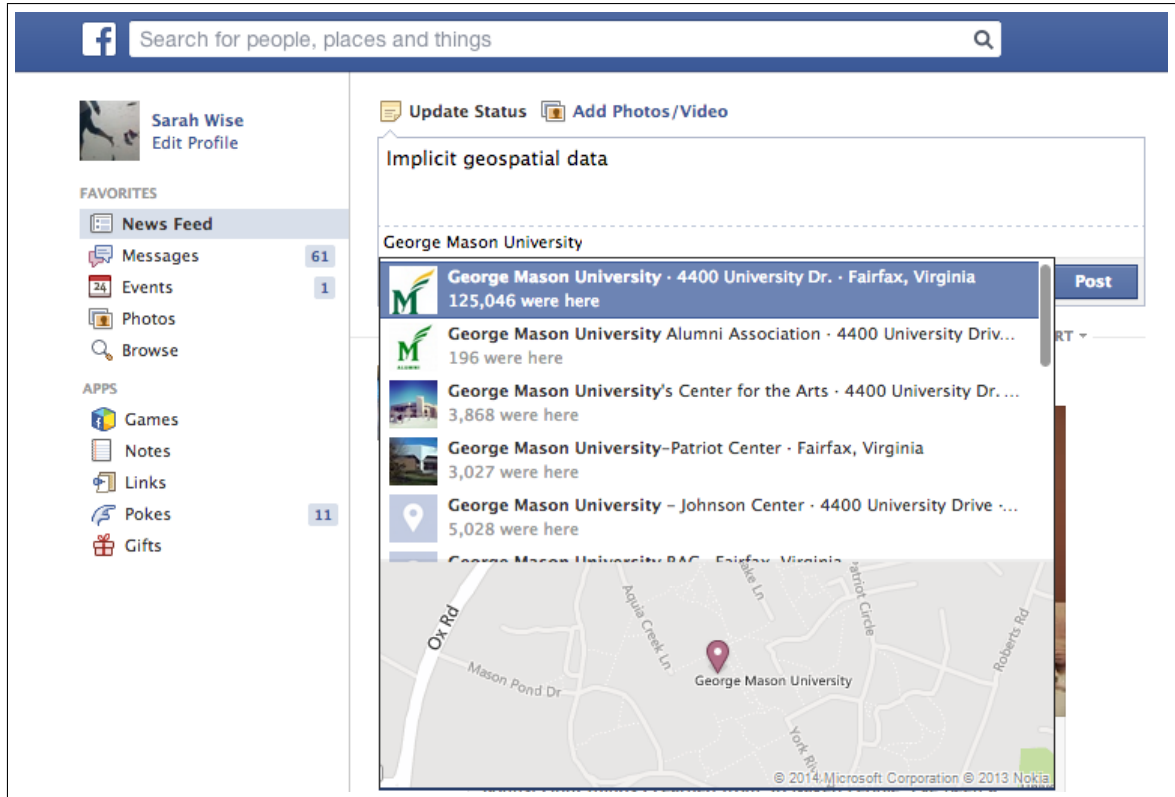


Figure 3.3: A Facebook status with embedded geospatial information: an example of blended geospatial and social content

3.3.2 Ambient Geographic Information

Distinct from yet related to VGI is the concept of ambient geographic information (Stefanidis et al., 2013). AGI refers to the information that is not explicitly or intentionally geospatial, but instead has geographic information embedded into the generated content.

A Flickr photo tagged with a location, for example, or a series of Tweets and Facebook posts describing the progress of a flood, have geographic footprints which can be extracted and melded together with other geospatial data to create a narrative. Figure 3.3 shows an example of this blended type of data in the form of a Facebook status post being tagged with a specific location. Stefanidis et al. (2013) compare the individual users who generate ambient information to sensors, “detecting” emergent phenomena and consequently generating information about them. Because the study of AGI is inherently concerned with the intersection between geography and social media, it is interdisciplinary by nature, drawing upon the disciplines of geography, computational social sciences, linguistics, and computer science (Crooks et al., 2014). Despite the challenges, the potential insight AGI offers into social dynamics across a range of cultural, societal, and human dimensions is immense, and only growing (Scharl & Tochtermann, 2007).

The ways that information can be extracted, of course, vary over space, time, and social dimension. Croitoru et al. (2012) note that AGI frequently is extracted from either the precise coordinates associated with GPS-enabled tagging, as from a cell phone, or from toponyms included in the text. Stefanidis et al. (2013) report that somewhere from .5% to 3% of all Twitter posts have explicit coordinates included in their metadata, but that the frequency with which tweets contained such data varied; during the Fukushima disaster in Japan, 16% of posts were geoenabled, largely because citizens of an extremely technologically advanced country were moving around and therefore not updating social media from home or work computers. Beyond precise geotagging, Croitoru et al. (2014) found that 40% to 70% of tweets included a descriptive toponym which gave insight into the location of the user; Friedland and Sommer (2010) estimated that 4.5% of Flickr images and 3% of YouTube content are geolocated. Croitoru et al. (2014) argue that as services like Twitter and Facebook offer more and more options for users to tag content with geographic information, the richness and potential for the field only expand.

3.4 Challenges

Along with the obvious power offered to researchers by developments like crowdsourcing, VGI, and AGI come a suite of concerns. Primary among these is the question of validity, as knowledge moves ever further from being the product of a single, central, expert authority. Leaving questions of accuracy and trustworthiness aside, developing the conceptual tools to deal with data at a scale researchers have never before experienced presents a series of challenges at a theoretical level, and these questions push the boundaries of research. Finally, questions of privacy and other less tangible concerns are an important part of any research effort. This section will address these questions and suggest ways that they can be mitigated.

3.4.1 Validity

Perhaps the most obvious question researchers, responders, and private individuals have asked when considering the usage of these new sources of information is whether the sources themselves can be trusted (Friedman, 2006; Tapscott & Williams, 2006). Most of the volunteers in such efforts are not professionals and do not follow standards of data collection and verification (Keen, 2007). How can users of the data be sure that a largely independent group of amateurs with little coordination and varied motivations for participation will produce something worth using? Haklay (2010) carried out a study of the quality of OpenStreetMap data, validating it against canonical, authoritative datasets from the UK's Ordnance Survey. His findings indicate that the structure of the collection effort itself strongly impacts the quality of the data, and stresses the idea of “forgoing completeness” in mapping in general. Further, he notes that while OpenStreetMap data varies in quality over different areas, it can be rapidly corrected without need for an extensive, formal review process; authoritative data sources sometimes suffer from similar incongruities as a result of update cycles and cartographic limitations, but it takes much longer to update them (Goodchild, 2008). Jackson et al. (2013) also stress the high quality of VGI, even in areas of relative deprivation. The variation in quality of geographic data over different areas

within the same VGI platform gives rise to the suggestion that it is best to treat different mapped areas as different datasets entirely (Haklay et al., 2010). Both Haklay (2010) and Crooks et al. (2014) argue for the concept of data that is “good enough”, or suited to the particular task at hand. This idea is particularly important in crisis situations, where (as discussed in Chapter 2) time is a precious resource and courses of action must be decided, with or without thrice-checked geospatial datasets.

In situations where a curated, authority-sanction datastore was never available, validity concerns take on a different form. How can researchers be certain that the datasets they extract are representative, and that the insights drawn are reflective of the system overall rather than a particularly vocal minority? Duggan and Brenner (2013) clearly show that the usage of Twitter usage varies with age, sex, urbanity, and income, without these factors being independent of one another. Thus, there is an inherent bias in who uses social networks and, as a result, the content that is generated as a part of such platforms. That being said, social media is growing in influence, and ever more individuals participate across a wide range of platforms, meaning that existing techniques will only become more and more useful as more and more individuals join up (see Smith, 2011; Nielsen, 2012). Active members of social media platforms generate a disproportionate amount of the content (Kwak et al., 2010), and this leaves aside the possibility that individuals might maliciously inject information into the system in order to confound attempts at analysis. Bremmer (2010) cite the example of organizations that are paid to support the Chinese government by blogging or posting on message boards pro-Communist Party opinions in an effort to sway public opinion, while Gupta et al. (2013) describe the difficulty of dealing with spammers and other purposefully fraudulent posters in the context of information extraction. Even when the signal is relatively clear, it can be unclear what individuals are discussing without greater context, making information extraction challenging (Vieweg et al., 2010). To a certain extent, many of these issues can be addressed through big data and crowd-sourcing techniques - Ford (2012) suggests waiting for a crowd to come to a consensus and verify information before that information is treated as trustworthy, and Okolloh (2009)

emphasizes that independent reports can be mutually supportive, lending credence to the report of any one individual. These efforts are promising, and speak to the emerging need for a conceptual understanding of big data and the rigorous synthesis of diverse sources of information.

3.4.2 Theory

As new sources of information and new ways to combine that information become available, affordable, and otherwise feasible for researchers to explore, there develops a corresponding need for a rigorous understanding of how to understand the new data. Croitoru et al. (2014) note that traditionally, geospatial analysis of human systems has focused on large-scale studies of populations at the aggregate level, individuals or groups at a smaller scale, and on remote sensed imagery. This focus was the result of data availability limitations and practical questions about how to warehouse and process these large datasets. They compare the emergence of big data and its associated techniques to the invention of the microscope, prompting a similar shift in the paradigm of how socio-spatial phenomena are observed and analyzed. West (2013) argues that the coming years will prompt the development of such methodologies, which he suggests will draw inspiration from complexity science. As mentioned during the discussion of validity, the emergence of volunteered information in particular has prompted groups and individuals to try to create new technical and social processes that will help verify information. These methodologies are evolving in real time, and to an extent being stress-tested by organisations like the open-source information collection, visualization, and mapping group Ushahidi (<http://www.ushahidi.com/>). The interplay between research and practice is underway, and promises to advance both substantially.

3.4.3 Other Concerns

Other ethical concerns include questions of privacy, exploitation, and social justice. Obviously many private citizens do not intend for their messages to be harvested by data

trawling efforts, however public the settings of their profiles. Even more problematically, by gathering and synthesizing ambient geospatial information, researchers have the potential to reveal information that social media users did not intentionally or explicitly make public (Friedland and Sommer, 2010). These new kinds of privacy and public interaction can be confusing for individual users, as the online uproar every time Facebook changes its privacy settings suggest (see Hoadley et al., 2010). Even if individuals are aware of their contributions, as in the case of VGI, some researchers have voiced concerns about the ethicalness of the means of production. Haklay (2010) questions whether VGI can be exploitative, noting that volunteers may be misled into believing that their efforts support some greater goal when in fact they support a particular enterprise which profits from their volunteered efforts - he gives the example of Google Map Maker, wherein users create the map but Google maintains the copyright. Perhaps more troubling is his concern regarding the quality of data coverage in different areas: Haklay (2010) notes that, while government-funded organizations such as the United States Geological Survey or the United Kingdom's Ordnance Survey are required to map all areas regardless of remoteness or socioeconomic status, VGI is subject to no such regulations. He reports that in the United Kingdom, wealthy areas are better mapped than poorer areas. However, the findings of Jackson et al. (2013) do not reflect this systematic discrepancy, suggesting that earlier issues may have been the result of the extreme newness of the phenomenon. These concerns are unlikely to prevent researchers from bringing these tools to bear on acute crisis situations - data production targeting the effected areas is unlikely to be be exploitative, nor are response organizations likely taken to task for insufficiently anonymizing individual messages in the pursuit of saving lives - they are important considerations for those designing tools for longer-term projects and other humanitarian situations.

3.5 Summary

Together, the technologies introduced in this chapter form a powerful toolbox of techniques for the exploration of new data, old data, and the synthesis between them. Some of the most

interesting work has dealt with combining authoritative and non-authoritative data sources in order to gain more insight into the situation (see Sui & Goodchild, 2011). Already, these new tools are being used to give citizens insight into the development of forest fires (see Goodchild & Glennon, 2010; Liu & Palen, 2010; Roche et al., 2013), to detect and track earthquakes (see Crooks et al., 2013), to detect hotspot emergence during political events (see Stefanidis et al., 2013), to watch the development of riots (see Tonkin et al., 2012) and epidemics (Achrekar et al., 2011), or to follow flooding (see Graham et al., 2012). By linking data sets with geographic and social information together, disparate data points turn into a narrative, reflecting structures which have never before been noticed, let alone documented. By drawing on these new kinds of information, researchers can power ever more rigorous forms of analysis. Big data, volunteered and ambient geography can be combined and synthesized to produce new forms of knowledge, and this knowledge can in turn be fed into simulations, as it will be in Chapters 6 and 7.

Chapter 4: Human Social Networks

Human experience is shaped by one’s knowledge of and interactions with other people. Given the importance of information in a crisis situation, as highlighted in Chapter 2, understanding the role of one’s social network in transmitting information and sentiment is extremely important. Toward that end, this chapter deals with the question of human social networks, exploring their composition, structure, and function, as well as how to automatically reconstruct them. Section 4.1 gives a sense of the scale of the problem, discussing the number of individuals involved in the discussion that follows. Section 4.2 delves into the connections between individuals, exploring the kinds of relationships between individuals and the frequency with which these individuals contact one another. Zooming out of the individual-level view of a network, Section 4.3 discusses the social network overall in terms of a number of metrics which characterize it. This general understanding of social networks is operationalised in Section 4.4, which explores both existing social media networks and efforts to simulate or emulate these networks and the way they influence the spread of information. The information discussed here supports the social network generation model presented in Chapter 6 and the social network utilization patterns described in Chapter 7.

It is important to note here that in the literature dedicated towards social networks, many specific definitions or conventions exist. Some terms are simple and pervasive, such as the use of the word “node” to indicate the vertices of the graph, which in the case of human social networks are obviously individual humans. Likewise, the terms “link” or “edge” refer to a connection between two nodes, such that the set of all edges make up the network. Other network-specific terminology will be introduced below when appropriate, but a number of other important concepts should be highlighted from the beginning. For example, one of the most important distinctions in talking about networks is that networks may be approached as either **egocentric** or **whole**, where whole networks take a given

population and explore the relationships among the individual members while egocentric networks focus on an individual and try to explore all of their social influences (Berg et al., 2010). A detailed review of social network concepts is beyond the scope of this chapter; only the elements which relate to this thesis will be discussed here, but many detailed texts on the subject are available (e.g. Wasserman & Faust, 1994).

The precise relationship between the real-world analogs and in-model representations of a network's links can vary depending on the goals of the researcher. In network terminology, links can be directed or undirected, or weighted or unweighted, among other options. These different kinds of relationships can describe an individuals in an unequal relationship (A is a fan of B, or C is D's professor) or relationships of varying strength (A and B are best friends, while A and C are merely casual acquaintances). And the relationships indicated by the links can vary: in different works, links may indicate intimacy, the provision of aid, or simply frequency of contact (McCarty et al., 2001; Wellmen & Wortley, 1990; Berg et al., 2010). The idea of homophily, or the tendency of individual to associate and bond with similar others, contextualizes many relationships (Lazarsfeld & Merton, 1954). For example, Wellmen and Wortley (1990) describe relationships along different members of a community, noting that one might help a neighbor out without really caring for him, or be closely linked with a family member without showing frequency of contact or provision of support; thus, links indicating the intimacy between individuals would form a very different network than the links which indicate frequency of contact. Even when the nature of a link seems clear, it can be more complicated than it appears: online social networks allow for relationships to be formally defined between individual users, but such connections may be meaningless depending on the context. Facebook users increasingly receive friend requests from complete strangers, which are usually accepted, casting doubt on to the meaningfulness of personal networks in certain contexts (Leow, 2009). Thus, it is extremely important to be precise in the terminology one uses to describe networks and the links within them, as well as how those links translate into the operational structures that characterize real life interactions.

In addition to these concepts, it is useful to define a number of network properties which are used to characterize the structures described in this chapter and the rest of the thesis. This work will use the following terms as defined in Table 4.1.

4.1 The Size of the Network

Before embarking on a study of the qualities of an individual's social network, it is useful to have a sense of roughly the scale of the networks with which the individual interacts. Within a given egocentric network, the concept of Dunbar's (1992) number can be helpful to gauge the approximate appropriate size of an individual's group of social contacts. Dunbar's number is described as roughly the cognitive limit to the number of people with whom an individual can maintain a social relationship at any given point in time. The number was derived by Dunbar (1992) based on his assumptions about the size and procreating power of primates' neocortical processing capacity, and is currently assumed to range between 100 and 230 people, with 150 being favored in the wider literature (Berg et al., 2010). The number excludes relationships that are inactive (such as childhood friends or deceased relatives) as well as non-social relationships (such as coworkers with whom one has no social contact or the barista at the local coffee shop). These exclusions highlight the fact that Dunbar's number purports to capture the number of social relationships one can maintain, rather than the number of people one can recognize and name. Dunbar views this number as the mean group size, and he validates it against data about village, community, and group sizes in a number of different areas, suggesting that his number correlates well with less-developed societies.

Other researchers have introduced variations on this theme: McCarty et al. (2001) derive a mean network size closer to 290 for modern communities in the United States. Wellman and Wortley (1990) break the networks they find in late 1970s Toronto down into egocentric networks of about 137 socially close intimates and 207 less intimate but significant contacts, noting that the median active network has four intimate ties and seven significant

Table 4.1: An overview of the network terminology used in this thesis

Network Properties	
Betweenness	a measure of a node's centrality in a network. The number of shortest paths from all vertices to all others that pass through that node.
Clustering Coefficient	a measure of the degree to which nodes tend to cluster together. This measure may be applied to the graph overall (globally) or else to the degree to which a given node is embedded within a dense cluster (locally). A more thorough description of the metric is available in Watts and Strogatz (1998).
Degree	the number of edges associated with a particular node. In dealing with directed networks, researchers may distinguish between the number of links which link to the node (that node's "in-degree") and the number of edges that originate at the node (that node's "out-degree").
Distance	the number of edges in the shortest path between two nodes. If no such path exists, the distance is said to be infinite. Frequently, a network is presented in terms of the average network distance, which is found by taking the average distance between every possible pair of nodes in the network.
Diameter	the maximum distance between any two vertices in a network.
Modularity	a measure of the strength of division of a network into modules.
Types of Network	
Scale-Free	a network in which the degree distribution of the nodes follows a power law.
Small-World	a network in which most nodes are not neighbors, yet the average distance is relatively low.
Fully Connected	a network in which every node is connected to every other node.

ties. These higher numbers may have to do with an idea Dunbar (1993) suggests. He assumes that social contact and gossip in human society has its analog in social grooming in primate society, but because grooming is time-intensive and necessarily co-located, as well as focused on a single other, it is a costly activity. Dunbar suggests that language evolved as a “cheap” form of the social grooming construct. As social interaction became so much more convenient, social groups could expand in size, and with the advent of new technologies such as social media social relationships can endure over great distances (see Hampton et al., 2000). Thus, the size of a social group should correlate to the ease of interaction. Section 4.2.2 will address some of the ways in which the frequency of communication has changed over time, giving more insight into this question, but the survey work presented by the researchers mentioned here suggests that an artificially generated egocentric network should roughly approximate the Dunbar number in terms of active relationships.

4.2 Social Connections

How are individuals linked to one another? Even within their egocentric networks, most people exist within a number of different networks simultaneously, so that an individual may belong to a group of coworkers, her family, a hobby group, a religious organization, and so forth. Wellman and Wortley (1990) work with the idea of a “personal community network”, or an individual’s set of active community ties, which they suggest tend to be socially diverse, spatially disperse, and sparse. They specifically mention the clusters of friends and family which exist within such networks, noting that the networks are low-density overall except for these smaller structures. Wellman and Wortley (1990) also investigate the different kinds of contact between individuals, noting for example that while a person frequently has contact with his coworkers, he rarely calls upon them for support, emotional help, or financial assistance. Burt (1984) breaks relationships down among the categories of friendship, work, kinship, and acquaintance, and notes that individuals were more willing to discuss stressful or personal questions with intimate relations than they were with acquaintances in general. Hampton et al. (2000) analyze the emerging modern dynamics in technologically advanced

areas, noting that individuals are much more likely to associate with others as a function of homophily than of distance. They argue that, while physical proximity remains a factor in relationships, homophily has become more important in an age of airplanes, cars, telephones, email, and so forth.

However, weak relationships are important to the spread of information. Granovetter (1973) in particular focuses on the strength of weak ties, famously noting that the majority of people find out about jobs through weak ties, not strong ones. Kleit (2001) confirms that weak social ties remain an important way for information to be transferred, also citing the instance of individuals finding jobs, with the interesting caveat that proximity matters more in these instances of weak-tie transferral. Information can spread over the light-weighted edges in a network, traveling between clusters of people and reaching between groups who otherwise rarely communicate. It has therefore been suggested that weak ties therefore have important functions for information flow in general, but that the intimate relationships that influence and motivate tend to be based more strongly on homophily (Louch, 2000). Relationships that are not strong or intimate can be very important, and it is important to consider how these relationship influence both interaction and communication.

4.2.1 Who Are These People, Anyway?

Given these smaller clusters of structure and different kinds of relationships, who exactly makes up this network and how does the individual know them? McCarty et al. (2001) report average values for the number of immediate family, coworkers, fellow hobbyists, school friends, neighbors, and so forth that individuals interviewed in phone surveys reported. Their results are shown in Table 4.3.

It is reasonable to wonder about the degree of overlap among these populations. For example, how many neighbors are also friends? Wellman and Wortley (1990) report that friends are not typically neighbors, although they also do not live too far away: “most” ties extend beyond the neighborhood, but 23% of active network members live within a mile of the respondents, even within suburban areas. When their study was published in 1990,

Table 4.3: Average number of people known for relation types (Source: McCarty et al., 2001)

Relation Types	
Immediate family	3.5
Other birth family	24.0
Family of spouse or significant other	12.3
Coworkers	35.6
People at work but don't work with directly	62.1
Best friends/confidantes	4.3
People known through hobbies/recreation	12.3
People from religious organization	43.4
People from other organization	17.1
School relations	18.3
Neighbors	12.8
Just friends	22.6
People known through others	22.6
Childhood relations	6.8
People who provide a service	7.7
Other	3.9

they found that the median distance between linked individuals was 10 miles. Wellman (1992) makes an important distinction: in North America individuals rarely consider more than one neighbor a close friend (that is, an intimate tie), and “neighborhood” relations represent only a quarter of all active social ties of any strength (Fischer, 1982). Scellato et al. (2011) present findings which better reflect modern developments in human social works, noting that almost half of all social links on social media applications are to individuals within 100km of the user. Scellato et al. (2011) further note that network size and spatial distribution seem to be correlated, with more highly connected users having more spatially distributed social networks.

Researchers have long suggested that individual who are friends with one another tend to be very similar - as previously cited, Lazarsfelt and Merton first introduced the idea of homophily in 1954, and it remains an important concept. Evidence has been collected that shows that individuals tend to have friends who are similar to themselves, and that this is particularly the case for relatively dense social networks (see Fischer, 1982; Marsden,

1988). Autant-Bernard et al. (2007) argue that in the context of individuals deciding whether to collaborate, social distance is more important than physical in the formation of cooperative partnerships, and that geography impacts interaction because of how it shapes and constrains the social network. That is, individuals will tend to know others who work in the same country because of the national structures which promote those individuals knowing one another, and will therefore be more likely to work with others in the same country. Louch (2000) also argues for the importance of homophily to strong network ties, stating that individuals who share race, education, and religion are more likely to be intimate, in descending order of influence. He also notes that the absolute difference in age between individuals influences the likelihood of connection, although there are non-linearities in the data, and that gender influences homophily more strongly for married individuals than for others. Homophily influences the development of the network in significant ways, and is an important consideration for researchers attempting to create synthetic social networks.

4.2.2 Making Contact

As noted above, the strength of the tie does not necessarily correspond to the nature of the relationship, nor the frequency of contact. Wellman and Wortley (1990) note that physically nearby network members make contact more frequently, either by phone or in person, regardless of the strength of ties. However, they note, the strength of a relationship is significantly correlated with the frequency of telephone contact: while individuals reported seeing most network members no more than twice a week, phone contact was much more frequent. Wellman and Wortley (1990) report that the median frequency of face-to-face contact for individual members of personal networks was 24 days per year, while the median frequency of telephone contact was 12 days per year.

Of course, communication has changed significantly in the past five years, let alone the last twenty. Table 4.4 shows the frequency of contact by different modes of communication in 2010, giving a sense of how newer forms of communication permeate current patterns of

Table 4.4: Frequency of contact in terms of instances of contact (Source: Berg et al., 2010)

Contact Frequency per Year		
Type	Mean	SD
Face-to-Face	48.7	70.4
Telephone	31.5	53.2
E-mail	13.7	33.6
Text Message	7.2	26.2

communication. While the telephone may have eased the process of communicating with distant friends, technologies like emailing, text messaging, Facebooking, and Tweeting, to name only a very few, have put the ability to rapidly contact huge quantities of people in the hands of nonspecialists. Socialization is possible on a scale unprecedented in human history, and this is even more evident as 91% of American adults own cell phones, with a full 56% of the adult population having smartphones (Duggan & Rainie, 2012). The United States Census Bureau (2013) reports that in 2011 71.7% of American households had internet access at home, and Duggan and Brenner (2013) report that 67% of internet users use Facebook. The pervasiveness of new forms of communication ties individuals together in ways that were unthinkable a decade ago.

One important emergent pattern in communication is the idea of the information cascade. An information cascade is the process of individuals making decisions sequentially to imitate the actions of others (Easley & Kleinberg, 2010). In social media, this may take the form of a tweet which is generated by an individual, retweeted by that individual's followers, then retweeted again by the followers of the followers until the piece of information has been massively, widely retweeted. The phenomenon of information cascades has been studied in the context of email chain letters (Wu et al., 2004; Liben-Nowell & Kleinberg, 2008), blog posts (Galuba et al., 2010; Gruhl et al., 2004; Leskovec et al. 2007), and image or video-sharing sites (Szabo & Huberman, 2010; Cha et al., 2008; Crane & Sornette, 2008). Sometimes this rapid, wide-spread propagation of information can take a humorous turn: a high school student who posted a Facebook invitation to a party at her house without

restricting the visibility of the event to her friends found that some 200,000 people had indicated on the party's webpage that they would be delighted to attend (De Zwart et al., 2011). Information cascades can help researchers understand the structure of networks, identify influential nodes, and understand what content is popular with whom (Lerman & Galstyan, 2008).

4.3 The Network Community

Given the dynamics at play in these egocentric networks, what can be said of the community of individuals, the network of everyone interacting together? Building upon the concepts introduced in the introduction to this chapter, it is possible to meaningfully discuss a number of network-construction models. Three of the most popular models are random networks (Erdos & Renyi, 1960; Bollobás, 2003), preferential attachment (Barabási & Albert, 1999), and small world rewiring (Watts, 2003). All produce networks with some, but not all, of the characteristics of real-world social networks. To briefly introduce the models, random graphs in general were first explored in the work of Erdos and Renyi (1960), with Bollobás (2003) serving as a major modern reference. Random graphs are graphs generated, simply enough, by some random process, for example (as in the model of Erdos and Renyi) by declaring all possible edges to be equiprobable and then rolling a die to determine which actually exist. This method produces networks with low average path length, but the node degrees are Poisson-distributed. Preferential attachment models, sometimes known as Barabási-Albert, generate graphs by adding nodes to the network one by one, linking them to a node with probability proportional to the number of links the node already has. The resulting graph has low average network distance, some degree of local clustering, and is scale-free. Finally, the Watts-Strogatz model of small world rewiring constructs an undirected ring of nodes where each node is connected to a fixed number of neighbors on both sides of themselves. Then, for each node, each edge is rewired with a fixed probability, where rewiring consists of reattaching the link at random among all of the nodes in the

Table 4.5: Properties of human social networks in general

Parameter	Value	Source
Avg Degree	10	Eubank et al., 2004
Network Diameter	6	Eubank et al., 2004
Avg Clustering Coeff	0.480	Eubank et al., 2004
Avg Path Length	6	Albert & Barabási, 2002

network. The resulting networks have low average network distance, some degree of local clustering, and a set of degrees which form a Poisson distribution.

To compare these generated networks with the best understandings of real social networks, researchers have found that the characteristic of real-world social networks include a low average network distance (Watts 2003; Szabó et al. 2003); a moderate clustering coefficient (Kilduff & Tsai 2003; Watts 2003); and an approximate power-law distribution of node degree (Albert & Barabási, 2002). Specifically, the average network distance should be approximately equal to $\frac{\log(n)}{\log(d)}$ (where n represents the number of people involved and d is the average degree of all individuals) (Watts 1999; Bollobás 2003; Durrett 2007). For example, the famous case of Milgram’s small world experiment suggests that the average network distance for social networks in the United States should be approximately 6, and that a small world network structure might be an appropriate model in general (1967). Java et al. (2007) report that based on their research, the Twitter social network conforms to a small world network structure with power law distributed degree distributions. They report that Twitter has a diameter of 6, high reciprocity, and a clustering coefficient of about 0.106, relatively low. Thus, Twitter conforms to the expected structure of a social network, as do parts of the blogosphere (Shi et al., 2007). The properties of a number of real-world social networks are shown in Table 4.5. Exploring this line of research further, it is informative to consider in depth the structure of real-world social networks, especially social media networks. It is to these that this chapter now turns.

4.4 Social Media

As discussed in Chapter 3, social media refers to the social interaction and virtual communities which have emerged from the technological developments of Web 2.0. Stefanidis et al. (2013) define Web 2.0 as six often overlapping concepts: individual production and user-generated content, harnessing the power of the crowd (e.g. crowdsourcing: see Howe, 2006), data on a massive scale, participation-enabling architectures, ubiquitous networking, and openness and transparency. O'Reilly (2005), Anderson (2007), and Batty et al. (2010) all discuss the impact of Web 2.0 technologies on this new kind of sharing. Wu et al. (2011) describe the notion of “masspersonal” communications, the amplification of public interpersonal communication to mass communication with a potentially enormous body of unknown others. Divol et al. (2012) assert that in the context of consumer behavior, social media's four primary functions are “to monitor, respond, amplify, and lead” opinion. These emergent, inherently participatory communities foster the spread of information and sentiment, and have exploded into major cultural relevancy in the past few years.

It is important to note that this kind of social interaction is not limited to casual chatter - many people use social media in real-life situations, for example to gather or distribute information in crisis situations. Sutton et al. (2008) describe how members of the public utilized online communities to understand and report on the development of wildfires in their local area, often producing more accurate and current information than was available through professional news organizations. In a darker turn, while citizens were able to get more accurate, current, and specific information through Twitter during the 2008 multi-incident, coordinated attack in Mumbai, the terrorists responsible for the attacks utilized those same sources of social media to help synchronize their efforts and maximize casualties (Oh et al., 2011). Vieweg et al. (2010) note that “members of the public use social media to support the gathering and dispersal of relevant, useful information, and online destinations like Twitter and other Internet forums support such disaster-related citizen participation”. They cite specific examples, including the use of Facebook to communicate during the 2007

Virginia Tech and 2008 Northern Illinois University campus shootings and various other platforms utilized during the California wildfires of 2007 and the Sichuan earthquake of 2008.

This work overall and this chapter in particular will deal primarily with the Twitter platform, both because of its ubiquity and its relative openness compared to platforms such as Facebook. A brief introduction to the platform follows, after which the discussion moves to a more general description of some of the challenges, classification efforts, and metrics associated with social media analysis.

4.4.1 Introduction to Twitter

Twitter is a microblogging platform which was launched in 2006 with the intention of being a mobile phone-oriented platform (Lotan et al., 2011). Having grown to 230+ million monthly active users as of January 2014, 76% of active users access Twitter via their phones, with the remainder interacting exclusively through computers (Twitter, 2014). The platform allows individual users to post short texts to their accounts and “follow” other users. The posts, called “tweets”, are limited in length to 140 characters and are all characters, although it is now possible for users to include URLs and links to images. The user’s dashboard includes all of the tweets generated by the users he is following, but it is important to note that following is not reciprocal: the Twitter social work is a directed graph. Users can “retweet” the post of another user, rebroadcasting the post to their own network while tagging the post as being the original work of another user (“RT @OriginalPoster: here is some information!”). If the retweeting user has modified the text in any way, it is sometimes called a “modified tweet” (“MT @OriginalPoster: ...information!”). Users also have the ability to tag other users in their posts by including their username or “handle” (e.g. “Hello, @otherUserName!”) and to use hashtags to mark a tweet as being part of a larger discussion context (e.g. “This tweet exists in a #largerContext”). It is important to note that a tweet which begins with a mention (“@OtherUser...”) will be directed to the mentioned user and be visible only to the followers of the mentioned user; in cases where

an individual wants to mention another user but make their tweets “public”, the tweet is typically prefixed with a single period to avoid the directed setting. Thus “@OtherUser” is relative private, directed communication while “.@OtherUser” is relatively public. Twitter itself reports that in January of 2014, approximately 500 million tweets were sent each day.

Twitter usage is constrained by a number of factors, some of which will be more substantially discussed in Chapter 6 as a part of the effort to construct a synthetic population of Twitter users. In general, usage varies with age, sex, location, socio-economic bracket, and education (Duggan & Brenner, 2013). Sutton et al. (2008) studied the use of Twitter in the California wildfire context, reporting that only 10% of the individuals they surveyed as a part of their research were Twitter users and many of those were only recent adopters who had come to the platform in search of wildfire information. However, Vieweg et al. (2010) suggest that those who adopt are likely to continue using the technology and even expand upon their original usage, citing the case of new Twitter users during the 2008 American Democratic and Republican national conventions. Twitter reports that 77% of accounts are from outside of the United States and that over 35 languages are supported. As a result of this diversity, the platform has been used in very different contexts for very different purposes: Twitter is utilized by fans seeking to communicate with celebrities (Marwick & boyd, 2010), vulnerable persons attempting to organize in the face of emergency events (Hughes & Palen, 2009; Sutton et al., 2008; Sarcevic et al., 2012), protesters and observers live-reporting on protests (Croitoru et al., 2012; Stefanidis et al., 2013; Lotan et al., 2011; González-Bailón et al., 2012), individuals reporting on developing situations (Crooks et al., 2014), and friends casually conversing about nothing in particular, to describe only a small subset of examples. Thus, Twitter is an extensive and diverse platform for communication, utilized by many different groups for many different reasons. The regularities, structures, and patterns of communication which characterize and shape interaction along this network are explored further in Sections 4.4.6- 4.4.6, and Section 4.5 gives a worked example of a study involving the gathering and analysis of Twitter data to give a sense of such efforts.

4.4.2 Extracting Information Content from Social Media

One of the most challenging efforts in social media analysis is the pursuit of methods for meaningful extraction of information from social media. Researchers have attempted to determine everything from consumer judgement of various products (Berger & Iyengar, 2013; Berger & Milkman, 2012a) to political attitudes toward candidates (Fink et al., 2012; O'Connor et al., 2010) to the location of persons effected by the 2010 Haitian earthquake as a part of the immediate response (Heinzelman & Waters, 2010). In some cases, the goal of the research is to capture information about how people feel toward a product or a politician - Chapter 5 is dedicated to the process of sentiment extraction from social media, and a more thorough discussion is available there. In other cases the goal of the research has been to effectively create a map from the information, as described in Crooks et al.'s work (2013) on the Twitter community's reporting of the mild 2011 earthquake along the East Coast of the United States. Stefanidis et al. (2013) in particular describe the power of harvesting ambient geospatial information, noting that "social media feeds do not aim to empower citizens to create a patchwork of geographic information: geography is not their message. Nevertheless, the message has geographic footprints, for example, in the form of locations from where the tweets originate, or references in their content to geographic entities." Vieweg et al. (2010) similarly described utilizing spatial information gleaned from textual sources.

In addition to the texts of tweets, social media posts can in some cases specify location to an even more precise extent. Facebook allows users to tag their location (Facebook, 2014), as shown in Figure 3.3. Twitter users can enable geotagging on their mobile phone, a feature which automatically embeds the location of the user at the moment the tweet is uploaded within the tweet itself. This information can be used to construct a map of user locations, and information contained within the tweets can be spatially analyzed in association with this information. Further, information from other platforms with geoenabled information like the photo-sharing site Flickr can be combined spatially in this way, allowing researchers to gain a multifaceted view of the situation through language, pictures, and location. This

form of synthesis holds a great deal of potential for future research, and the current work only scratches the surface of what is possible.

4.4.3 Classification Problems

Given the amount of information that could potentially be extracted from social media posts, it is tempting to try to impose structures onto the data gleaned from such analysis in order to make it more manageable and accessible. At the moment, many efforts still go about classifying social media posts by hand-coding them into the desired categories. Classification can occur along a number of different axes: for example, during the Ushahidi response to the 2010 Haitian earthquake, responders attempted to classify communications into categories of need such as food distribution, missing persons, requests to forward a message, water shortage, medical emergencies, trapped persons, sheltering questions, and so forth (Ushahidi, 2010). Vieweg et al. (2010) classify the data they gather into similar situational categories: “Warning, Preparatory Activity, Fire Line/Hazard Location, Flood Level, Weather, Wind, Visibility, Road Conditions, Advice (i.e. advice on how to cope with the emergency, and/or advice on other Twitter users or websites to follow), Evacuation Information, Volunteer Information, Animal Management, and Damage/Injury reports.” Other researchers have focused less on the tweets themselves and more on trying to understand who is generating the tweets - Lotan et al. (2011) group the Twitter users they study into categories including mainstream media, mainstream new media, non-media organizations, mainstream media employees, activists, digerati, political actors, celebrities, researchers, bots, and “others”. Lotan et al. (2011) describe efforts to classify different users as mainstream media, bloggers, activists, celebrities, researchers, and so forth in order to understand how different kinds of users influence the spread of information. These categories are useful to responders, to researchers, to marketing departments, and to anyone who wants to understand the different kinds of information being posted to and spreading through social media networks.

While these structures are invaluable for those trying to understand the flow of information, hand-coding is extremely expensive, difficult, and arguably subjective. Caragea et al. (2011) describe this frustration (emphasis added):

”While there is useful information in these tweets and text messages, they are not well-organized to allow critical information (e.g., water, medical supply, food) to be delivered to those who need them in a timely and efficient fashion. Relief workers from different organizations, such as NGOs, military units, and government agencies, need IT support for analyzing tweets and text messages for ease of aggregation and targeted real-time broadcasting. Hence, **the ability to classify tweets and text messages automatically**, together with the ability to deliver the relevant information to the appropriate personnel are essential for enabling the personnel to timely and efficiently work to address the most urgent needs, and to understand the emergency situation better...”

Given the broad acceptance of this wisdom, it is little surprise that extensive research has been done along these lines. Crooks et al. (2013) attempt to track the spread of an earthquake using Twitter with simple filtering mechanisms, while Croitoru et al. (2012) construct a framework to extract information from social media feeds and analyze it spatiotemporally. Machine learning efforts consist of the construction of algorithms for training computer programs to classify objects quickly and precisely, which in the case of Caragea et al. (2011) has allowed them to automatically classify tweets according to their word usage. Many research teams have attempted to identify spam messages (see Healy et al., 2004; Gómez Hidalgo et al., 2006; Cormack et al., 2007) while Gupta and Ratnov (2008) classify short online dialogs and Munro and Manning (2010) classify medical text messages. Many of these efforts have been in languages other than English - Munro and Manning (2010) extract information in the Chichewa language, Kwak et al. (2011) explore Korean-language Twitter interactions, and Munro (2010) considers the usage of automatic classification as applied to Haitian Creole. This research into other languages is encouraging because it

suggests that the techniques developed in this thesis will be extensible to languages which are historically underrepresented online, which is valuable in situations where the effected individuals do not speak English.

Analysis on the aggregate level has proven to be particularly interesting. Stefanidis et al. (2013) investigate geospatial hotspot emergence as a function of references to gazetteer entries and abnormal peaks in references to specific terminology. Petrovic et al. (2010) follow a similar strategy, looking for news stories which break via the Twitter platform by means of a locational hashtag; Sakaki et al. (2010) similarly detect Japanese earthquakes through aggregating twitter users and treating them as “sensors”. A series of studies have focused on automatically extracting and aggregating information about potential disease outbreaks from news articles (Brownstein et al., 2008), search engines (e.g. Polgreen et al., 2008), and blogs (Corley et al., 2010). It is upon this tradition of automatic identification and classification that this work builds, in particular Chapter 5.

4.4.4 Measuring Social Networks

As a part of attempts to explore a range of phenomena, researchers have sought to explore the structures which underlie Twitter by analysing a number of different metrics. Some researchers focus on network properties like degree (the number of edges associated with a node), often breaking down the analysis into consideration of in-degree (the number of followers a user has) and out-degree (the number of users a user follows). Java et al. (2007), in the context of user intentions and virtual community structure, focus their analysis on some basic network properties: the number of nodes versus links, the average degree, clustering coefficients, connected components, and reciprocity. Cha et al. (2010) consider not only the in-degree of a user, but how frequently and widely their posts are retweeted and how frequently they are mentioned in the tweets of others. Krishnamurthy et al. (2006) classify the users under study into named categories depending on discrepancy between their in and out degree: the users they term broadcasters have many more followers than they themselves are following while acquaintances show more reciprocity in their relationships

and evangelists follow many more people than follow them. In the work of Glasgow et al. (2012), the researchers consider not only the in- and out-degree of individual users, but also the connections defined by expanding the realm of analysis to users who are within two or three directed edges of one another. Many other examples of this focus exist (for more examples see Vieweg et al., 2010).

Some researchers have noted the importance of considering not only the formally defined follower/followed structure of social networks, but also the emergent network structures apparent in retweeted information. Glasgow et al. (2012) describe the importance of distinguishing between mentions (the relatively private “@User”), publicized mentioned (the less private “.@User”), and retweets (“RT”) when analyzing the social context of communication, declaring these conventions “socially and semantically different”. Rather than focusing on the users, some researchers have turned their sights to the life of tweets themselves, especially in light of information cascades. Lotan et al. (2011) study information flows, analyzing the different types of users involved in the conversations as a function of the average number of responses they received per tweet, the total number of tweets they generated across different threads, the total number of times that tweets across threads received responses, the number of tweets they generated over the number of replies all their tweets received, and the total number of actors in each category. González-Bailón et al. (2012) compared the networks emergent from data they gathered by streaming versus searching Twitter in an attempt to understand how better to create these functionally interactive networks from real data.

The rationale for using a retweet network to analyze communication is that tweeting or retweeting represents active engagement in a community; a user might create a username and never sign in again, making their follower list a hazard for researchers trying to scrape data about interaction. Cha et al. (2010) suggest that because 92% of retweets contained a URL versus the 30% of “mention” tweets containing URLs, mentions are more identity-driven and are frequently driven by the name value of the user. They argue that information that is retweeted is propagated because the information itself is valuable, rather

than deriving its interest exclusively from its source; while celebrities are frequently mentioned, mainstream news organizations tended to generate more retweets over a range of topics. If this is so, analysis of Twitter data should consider these social conventions and their implicit meanings when attempting to measure phenomena. These ideas play into concepts like the notion of influence, operationalized by different researchers in different ways. Some calculate influence as a function of how many people will in all likelihood read a user's tweet, retweet it, and so forth (Tunkelang, 2009) while others consider the difference between popularity and generated retweets/mentions (Cha et al., 2010). Weng et al. (2010) compare their influence measure with more general follower size information.

4.4.5 Challenges of Working with Social Media Data

Some specific challenges exist for researchers attempting to utilise social media in the analysis and research. Analysis which involves extracting information or network properties like interactivity or influence from communications between individuals is particularly fraught. As will be explored further in Chapter 5, any effort which involves parsing natural language faces a number of problems (see Caragea et al., 2011). Vieweg et al. (2010) note the challenge of identifying contextual information omitted in messages which are produced as part of a conversation: they give the example of attempting to analyze the flooding of a specific river, when users would refer to “the river” or “the flood level” rather than the identifying river name itself. Hughes and Palen (2009) experience similar challenges in attempting to distinguish between discussions of the Democratic and Republican National Conventions. Caragea et al. (2011) mention the concern that a single tweet could belong to multiple categories, thereby complicating classification efforts. Users, too, can fit into a variety of categories and play a multiplicity of roles: Java et al. (2007) note that users can be information sources, information seekers, or friendly acquaintances in different communities, often within the same period of time.

The difficulty in determining what to study is another fundamental problem. Whether a researcher focuses on the explicit structures of a social network (officially-defined “friends”

or “followers”) or whether he extracts implicit information (constructing a network based on interactions between users) is an ongoing question. Huberman et al. (2009) note that people interact with a significant subset of their listed friends. This is especially challenging in light of privacy settings - Huges and Palen (2009) report that approximately 30% of the tweets they attempted to collect were marked private and therefore inaccessible to them. Gathering this data is difficult under the best of circumstances, as multiple researchers describe, and it is necessary to take precautions to avoid gathering biased samples as a function of the method of collection (Krishnamurthy et al., 2006; Wu et al., 2011; Hughes & Palen, 2009). Even when the data is successfully collected and the structure of the network is successfully extracted from ambiguous language, the way individuals use and interact with the medium is constantly changing. Hughes and Palen (2009) note the rise in the inclusion of URLs in tweets relative to the reports of Java et al. (2007). Further, analyzing data over shorter and shorter timescales brings with it all of the dangers of overfitting a model, yet the phenomena being studied are often characterized by sudden, dramatic spikes, a trend which has only increased as social media has continued to explode in popularity (Leskovec et al., 2009). Taken together, all of these concerns mean that even when the structure is correctly analyzed, it is constantly changing, and the implications of its structures have ever-shifting meanings. This changeability is an important consideration for researchers.

4.4.6 Towards a Model of Twitter

Given the varying approaches to the measurement and analysis of Twitter activity, how should one approach the process of building a model of Twitter? Arguably the most important considerations are the explicit structure of the network, the activity level of the users, and the way users interact with information. Together, these reflect the information that is introduced to the network via user activity level and the way information is propagated along the network as a function of the network morphology and the decisions individuals make about pushing information through it. A correct implementation of these features of

Table 4.6: Properties of samples taken from the Twitter social network

Parameter	Value	Source
Avg Degree	18.86	Java et al., 2007
Network Diameter	6	Kwak et al., 2010
Avg Clustering Coeff	0.106	Java et al., 2007
Avg Path Length	4.12	Kwak et al., 2010

Twitter should in theory produce a resulting realistic pattern of the spread of information.

Network Structure

Java et al. (2007) suggest that the overall structure of the network should be one of hubs and authorities, a claim Hughes and Palen (2009) confirm. Despite this focus on the elite few, Wu et al. (2011) point out the importance of “intermediary” users, highly active and well-followed individuals who function as amplifiers for information deriving from a wide variety of sources. Thus, users with high betweenness can play an important role in information-sharing. The importance of relatively weak ties is as powerful as ever. In general, the structure should resemble that shown in Table 4.6, drawn from the literature.

At different levels of analysis, different structures are meaningful. The qualities that characterize connection broadly are different from the characteristics of more local structures. One example of this is the distribution of followers. In the network overall, the distributions of followers are often extreme as celebrities and other major public figures boast millions of followers; for example, at the time of writing, each member of the boy-band One Direction had at least 10 million followers, and President Barack Obama’s official Twitter follower count was over 41 million. Kwak et al. (2010) note that for both the number of users an individual is following (“friends”) and the number of users by which that individual is followed (“followers”) have median values of less than 100. The extreme values of these distributions, however, are indeed extreme - a few users have had several hundred thousand followers each, despite the fact that these most-followed individuals frequently following few other users themselves. In a particularly extreme case of this imbalance, the

official account of the performer Beyonce lists over 13 million followers to her 8 friends. Both Wu et al. (2011) and Kwak et al. (2010) suggest that this indicates that the overall structure of Twitter is one of one-way mass communications, but note that interpersonal communications tend to be much more reciprocated. Thus, when modeling Twitter it is important to incorporate highly-connected persons, but also to consider how the average individual experiences the platform.

That average individual experience is the subject of some debate: different researchers regard the Twitter network in different ways when it comes to the question of reciprocity and homophily. While Cha et al. (2010) claim that reciprocity is low, they ignore users with private settings. Weng et al. (2010) support the idea of high reciprocity between users, and Stefanidis et al. (2013) highlight the influence of homophily on network structure when it comes to the experience of the vast majority of users, who overwhelmingly associate and interact with other similar users. When it comes to measuring, Kwak et al. (2010) argue that information sharing on Twitter is frequently non-reciprocal and therefore more of an information-sharing network than a social network: certain “influencers” of opinion and taste are significantly able to shape information flows. They argue that Twitter could thus be viewed as a broadcast medium rather than a forum for social interaction. However, this may depend more upon the kind of information being shared, a subject which will be discussed further in the following sections.

Degree of User Activity

Users demonstrate a range of activity levels - while some users create an account, post once, and lose interest, there are also bots which constantly push advertisements, gibberish, or both out into the system. Huberman et al. (2009) argue that the number of users an individual is following is a more accurate predictor of activity than is the individual’s number of followers. However, Krishnamurthy et al. (2006) find that users with many followers update more frequently than those who follow many users - indeed, Kwak et al. (2012) report that a number of users they interviewed reported being frustrated by the

volume of tweets generated by celebrities, even unto the point of unfollowing them. Kwak et al. (2012) further note that the type of information that was being broadcast was the source of the frustration, as many of the users who unfollowed celebrities remained fans of the celebrities themselves - the issue was the content, not the relationship. Wu et al. (2011) note that different categories of users emphasize different types of content, and that those types of content exhibit lifespans of interest that vary substantially. The work of Hughes and Palen (2009) shows the varying number of tweets generated on different days of a series of events, including two hurricanes, demonstrating the characteristic spikes in tweet generation on the days that the hurricanes made landfall; while it is not clear that users actually become more active during crisis situations, they certainly focus on emerging crisis topics. Sutton et al. (2008) record that individuals in crisis situations purposefully utilised social media because other sources of information were insufficient in various ways, and Vieweg et al. (2010) note the increasing utilization of social media in crisis situations. Finally, leaving aside spambots, media outlets are the most active users of Twitter by a considerable margin (Wu et al., 2011).

Spread of Information

Having established general patterns for the connections that exist between individuals and the frequency with which information is transmitted through these connections, the characteristic spread of information through the a model of the platform should resemble that of the real-world. Information is introduced into the system from a variety of sources - Wu et al. (2011) note that ordinary users are receiving information from a variety of places ranging from mass media to their own personal experiences. Of the tweets that ordinary users receive, only 15% come from mass media accounts, despite their aforementioned activity levels. A user with a large follower count and past success in triggering retweet cascades is, on average, more likely to trigger a large cascade in the future, but the predictive qualities of these features are poor (Bakshy et al., 2011). In general, the lifespan of information is short: useful, interesting, and surprising content is more likely to be viral (Berger &

Milkman, 2012b) and tweets with URLs such as originate from media sources are widely but briefly retweeted (Wu et al., 2011). Galuba et al. (2010) report that social graphs tend to produce shallow and wide information cascades in general, a trend seen across sites including Flickr, blogging platforms like Blogger or WordPress, Digg, and Youtube. They note that the depth of the cascade is exponentially distributed and that the diffusion delay between URL tweets in a cascade is log-normally distributed, having a median of 50 minutes.


It is important to note, however, that in crisis situations there is a low rate of reply in what might otherwise be conversations. Hughes and Palen (2009) suggest that because users are pushing information out to a mass audience rather than engaging in interpersonal communication, the spread of information generates less chatter and more percolation. They further note that tweets which pertain to specific crisis events include more URLs than a sample of general Twitter traffic, suggesting that people utilize Twitter as an information-gathering tool in emergency events. Yardi and boyd (2010) find that local networks become denser when addressing local events, and that this clustering is particularly pronounced in the face of newsworthy situations - a local conversation about a missing child, for example, is likely to be discussed within the community but rarely outside of it. Starbird and Palen (2010) find that those Twitter users with highly retweeted tweets were almost always mainstream (and especially local) media members, service organisations, or accounts created for the purpose of covering the event, and that approximately 90% of tweets involved in these conversations were retweets or modified tweets. Given Kwak et al.'s (2010) finding that Twitter often functions as a news breaking mechanism and Starbird et al.'s (2010) suggestion that individuals use Twitter to recommend important information to others, this emergence and rapid propagation of information makes Twitter a good source for local information.

Summary

By ensuring that a model of the Twitter environment has both well-connected and highly active “hub” agents, the model developed in Chapter 6 and utilized in Chapter 7 should conform to the overall trends of the platform; by ensuring that individual users are embedded within social clusters, they will demonstrate reasonable levels of reciprocity and homophily. Ordinary users should focus their attention on crisis events during the period of the crisis, with their attention peaking on the most eventful days. The spread of information in these contexts should be less conversational and more of a top-down distribution of information, reflecting information being pushed out to the network. If these phenomena are present, the modeled structure resembles the real-world platform in significant and meaningful ways. The networks generated by the model are presented in Chapter 6, and the results of their operationalization are explored in Chapter 8.

4.5 A Worked Example

In order to give a sense of how information can be collected and analyzed to gain insight into a social phenomenon, this section presents a worked example of social media data analysis using a collection of tweets gathered during the course of the Waldo Canyon Wildfire in Colorado Springs, Colorado. Chapter 7 introduces the context of the study in full, but to briefly summarize the events, the city of Colorado Springs experienced evacuation efforts as the Waldo Canyon wildfire threatened and eventually burned a number of homes in the city. Evacuation orders were in effect from June 23 until July 5, with about 26,500 residents evacuated from the city on June 26 (City of Colorado Springs, 2013). The data presented here reflects some of the social media conversation surrounding the wildfire and the evacuation effort it prompted. The full dataset consists of tweets generated between June 10 and September 21, 2012, where the text of the tweet contained the word “fire”. All in all, the dataset contains 188,784 tweets. Of these, 9,568 tweets are geotagged with precise coordinates. Further, the dataset references 28 unique geotagged images, many of which



A



B

Figure 4.1: Comparison of rate of generation of all (red) versus specifically geotagged (blue) tweets, in linear (A) and log (B) scales

were widely retweeted and shared. In the remainder of this section, the dataset will be explored to show spatiality, phenomenon emergence, and social structure. All descriptions of the events which took place during the wildfire are based on the City of Colorado Springs After Action Report (2013).

4.5.1 Geotagging

Section 4.4.2 described some of the ways in which information can be extracted from social media, a line of research that will be expanded upon here. By exploring the set of data that was geotagged with specific coordinates, it is possible to gain a sense of the spatiality of the data being collected. Only about 5% of the data was tagged with these coordinates: Figure 4.1 shows the number of tweets generated per day over the course of the peak of the wildfire, both in linear and log scale. While the linear plot shows the relative proportion of non-geotagged versus geotagged tweets, the log scale comparison indicates that the order of magnitude changes in the generation of both types of tweets are relatively consistent. This suggests that geotagged tweets are subject to similar pressures, although during the peak of the evacuation effort geotagged tweet generation exceeds pre-crisis levels by more than one full order of magnitude, while non-geotagged tweet generation increases by about an order of magnitude. Given the number of individuals tweeting from their mobile devices rather than home or work computers during the evacuation period, a slight bump in the number of geotagged tweets is to be expected, and parallels the example of the high percentage of geotagged tweets following the Fukushima disaster in Japan (Stefanidis et al., 2013).

From these geotagged tweets, it is possible to gain a sense of whence individuals are tweeting, and to construct a map of commentators. Figure 4.2 shows a mapping of tweets to the greater Colorado Springs area, while Figure 4.3 gives a closer view of the city of Colorado Springs itself. The wildfire at its greatest extent is included to get a sense of where individuals are relative to the threat. Thus, during the peak of the wildfire, it is possible to gain a sense of when individuals were where, and to factor this into further analysis. This is an example of AGI in practice, and the emergent picture is informative.

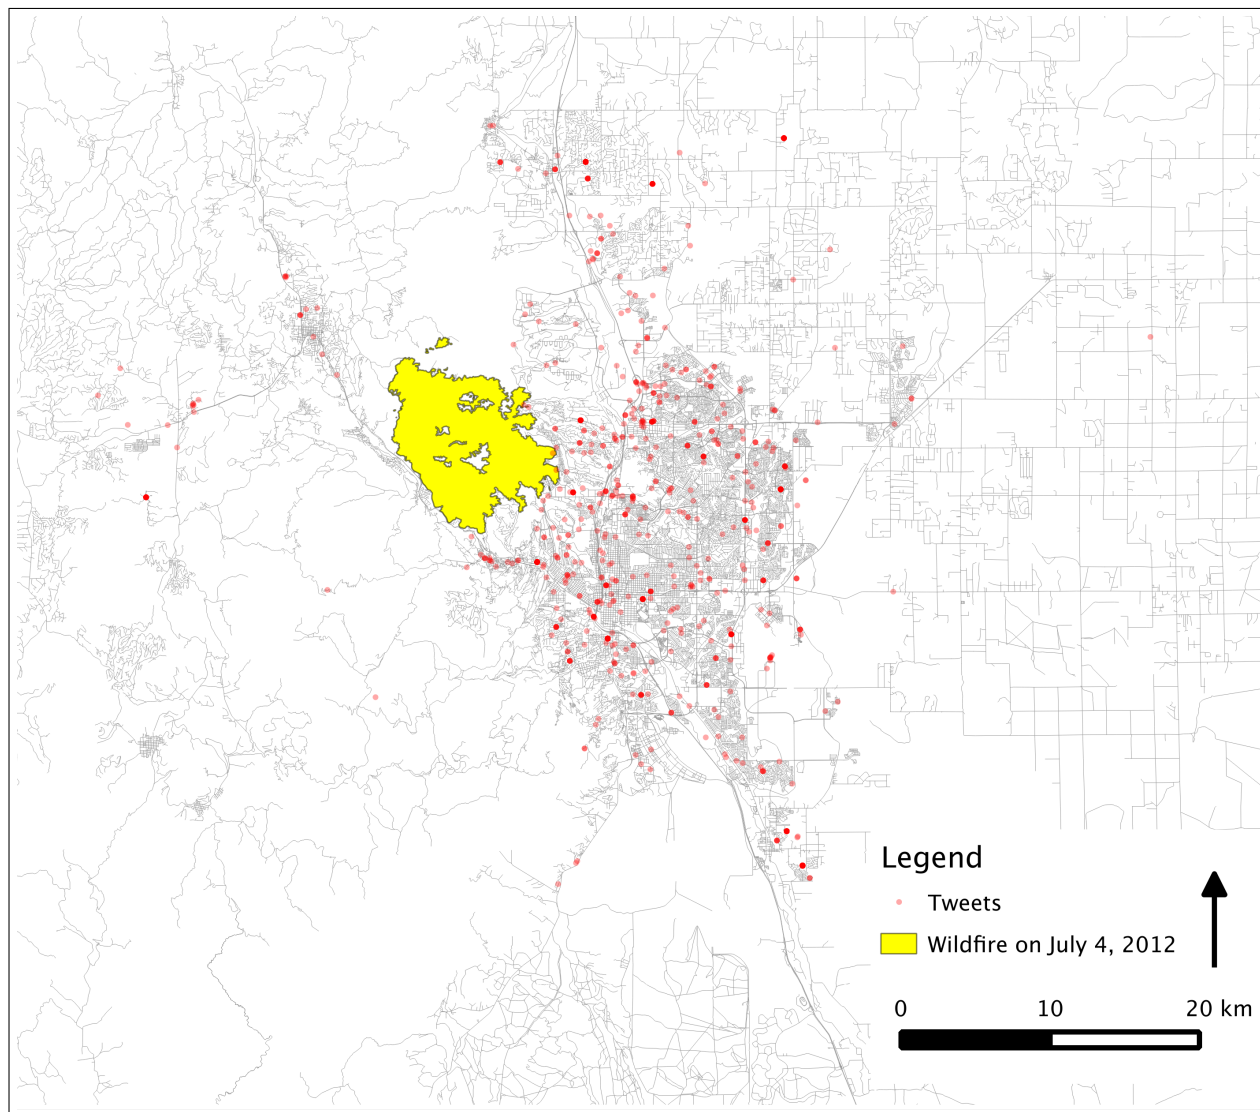


Figure 4.2: The set of geotagged tweets generated during the peak of the Waldo Canyon wildfire (from June 23 until July 5) mapped onto the Colorado Springs wider area road network to give a sense of Twitter user location during the peak of the wildfire. Individual tweets are designated by opaque red dots, so that multiple tweets from the same location are brighter red.

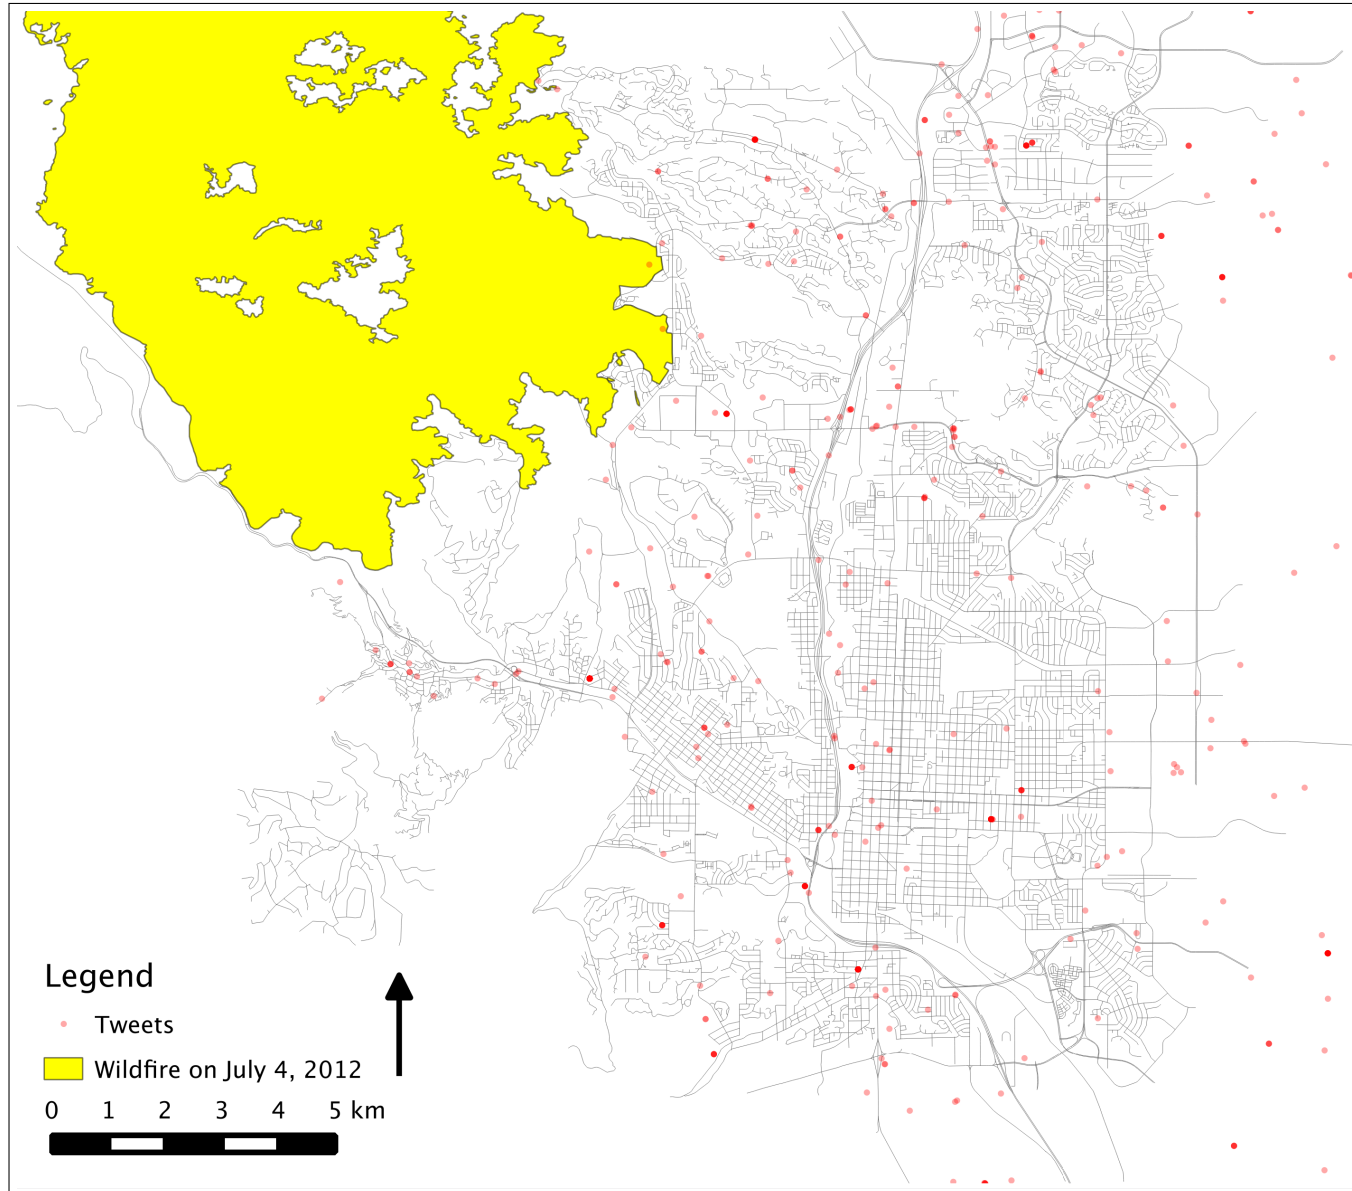



Figure 4.3: A closeup of the set of tweets generated during the peak of the Waldo Canyon wildfire mapped onto the Colorado Springs study area

4.5.2 Hashtag Tracking


As discussed earlier in Section 4.4, hashtags are used to indicate that tweet belongs to a particular conversation. Such indicators are particularly valuable for addressing some of the challenges described in Section 4.4.5 and Section 4.4.3. Toward the end of understanding how a term comes into use and how users converge on a particular hashtag to describe an event, Figure 4.4 tracks the usage of two hashtags over the course of the peak of the wildfire. The figure presents two candidate hashtags, #waldocanyon and #waldocanyonwildfire, comparing their adoption and usage over the course of the crisis. Figure 4.4A shows the usage of both terms on a linear scale, with the hashtags coming into being on the day the fire broke out. The hashtag #waldocanyonwildfire is obviously rapidly taken up by the community at large, with the less popular #waldocanyon hashtag trailing in popularity, although still being utilized throughout the peak. Figure 4.4B shows the changes in order of magnitude, indicating that the #waldocanyon hashtag actually initially peaked in its popularity on June 23, the day of the first evacuations. Despite the early rapid adoption of both, #waldocanyonwildfire emerges as the community standard for tagging discussions about the wildfire.

Also of interest is the rise and fall of the popularity of discussions in the #waldocanyonwildfire thread. Conversations peaked on June 24, the day the wildfire first moved into a residential area, prompting officials to close down the highway nearest to the fire. While the fire expanded by about 1,100 acres on June 25, there were no further evacuations until June 26, the day of the biggest spike in hashtag usage. June 26 saw the evacuation of approximately 26,500 city residents and the expansion of the fire by about 11,000 acres, with five parts of the city being issued mandatory evacuation notices. The wildfire continued to grow over the next few days, but no further evacuations were ordered and by the time the fire was declared contained on July 10, there was only a small increase in the number of tweets. #waldocanyonwildfire saw small peaks on July 3 and 5 in trends most visible in Figure 4.4B, days when mandatory evacuation orders were lifted. It is particularly interesting to compare the change in the number of conversations classified with these hashtags

with the findings presented in Section 5.6 regarding the tone and sentiment associated with different portions of the wildfire sample period.



A



B

Figure 4.4: Comparison of use of the “#waldocanyonwildfire” hashtag in linear (A) and log (B) scales. The lines indicate the rate of hashtag usage in all (red) or exclusively in geotagged (blue) tweets

4.5.3 Social Communities

In addition to social and spatial aspects of the social media dataset, it is possible to get a sense of the kind of conversations that are happening by exploring a network of active interaction: the retweet network. The virtues of utilizing retweet networks rather than structural networks was explained in Section 4.4.4. Tweets which were retweeted from other users were processed into a directed network, where each node represented a unique Twitter user and a link from Node A to Node B indicated that A had retweeted B's tweet. The link's weight was determined by the number of times A had retweeted B's tweets, so that for each of B's tweets A propagated, the link weight was incremented by one. From this simple graph of captured influence, it is possible to extract a number of clustered communities. Using a simple modularity-based cluster detection algorithm in the program Gephi, the subgroups shown in Figure 4.5 were identified and visualized. Not only is it obvious that different conversations are happening, it is clear that different structures are shaping the flow of information, giving real insight into the question of hierarchy, hubs, and authority introduced earlier in this chapter.

Of particular interest is the users participating in each of the communities. Figure 4.5A is populated primarily by Colorado social and news organizations - the highly connected nodes at the center of the graph include the users *csgazette*, *epcsheriff*, *krdonenewradio*, *springsalliance*. This cluster is the largest of all the generated clusters, with 2,136 unique users and 6,292 edges among them. There is clearly extensive conversation happening between the major players, suggesting interaction and exchange rather than information broadcasting in the traditional sense, as well as a large audience of information consumers. Figure 4.5B and Figure 4.5C show the retweet networks of some of the largest communities. The structure of these networks indicate that conversations are much more centralized, focusing around the twitter handles of the Denver Broncos and the Denver 9 News Channel respectively. Figure 4.5B contains 1,528 nodes with 1,908 edges, while Figure 4.5C has only 1,122 nodes and 1,530 edges. These communities are built around the hubs of the respective organizations, and consequentially show much less of the back-and-forth sharing apparent in

Figure 4.5A. Finally, Figure 4.5D centers on the communication between Colorado Springs-local official organizations, many of which are dedicated to emergency and news. Users include `coemergency`, `larimersheriff`, `rmaccfireinfo`, `larimercounty`, and a number of local news organizations. In all, the graph contains 899 nodes and 1,362 edges. Particularly interesting is the appearance of intermediate-level information propagators, individuals who convey information from these official sources to others, making it accessible to those who are not otherwise following official emergency channels. A wide range of types of communication is reflected in these different communities; while it is beyond to scope of this work to explore further, it would be fascinating to consider the kinds of information that are propagated through each of the graphs.

4.6 Summary

Human social networks are characterized by qualities dependent on the context in which they are formed and the ways in which they are reinforced and otherwise reshaped. The way that researchers conceptualize and study social networks varies depending on the purpose, context, and resources available for their research, ranging from measures of intimacy to frequency of contact to explicit, hard-coded social media network “friendships”. Section 4.2 shows that efforts toward measuring the relationships among humans as a function of these relationships have a long history, and the methods of research have changed with the introduction of technologies which allow humans new ways to interact. Even more importantly, these developments in technology and its resultant interconnectedness are not limited to the United States or a specific social class - these kinds of connections are appearing around the world. Twitter usage alone has a significant presence within other countries (see Kwak et al., 2012; Sakaki et al., 2010) and around the world more generally (Java et al., 2007). Efforts to understand social media through new technologies are already working with mobile phone information to understand communication, coordination, and influence (see Heinzelman & Waters, 2010; Lu et al., 2012; Starbird & Palen, 2012). As a function of the increasing permeation of social media described in Chapter 3, studying human social

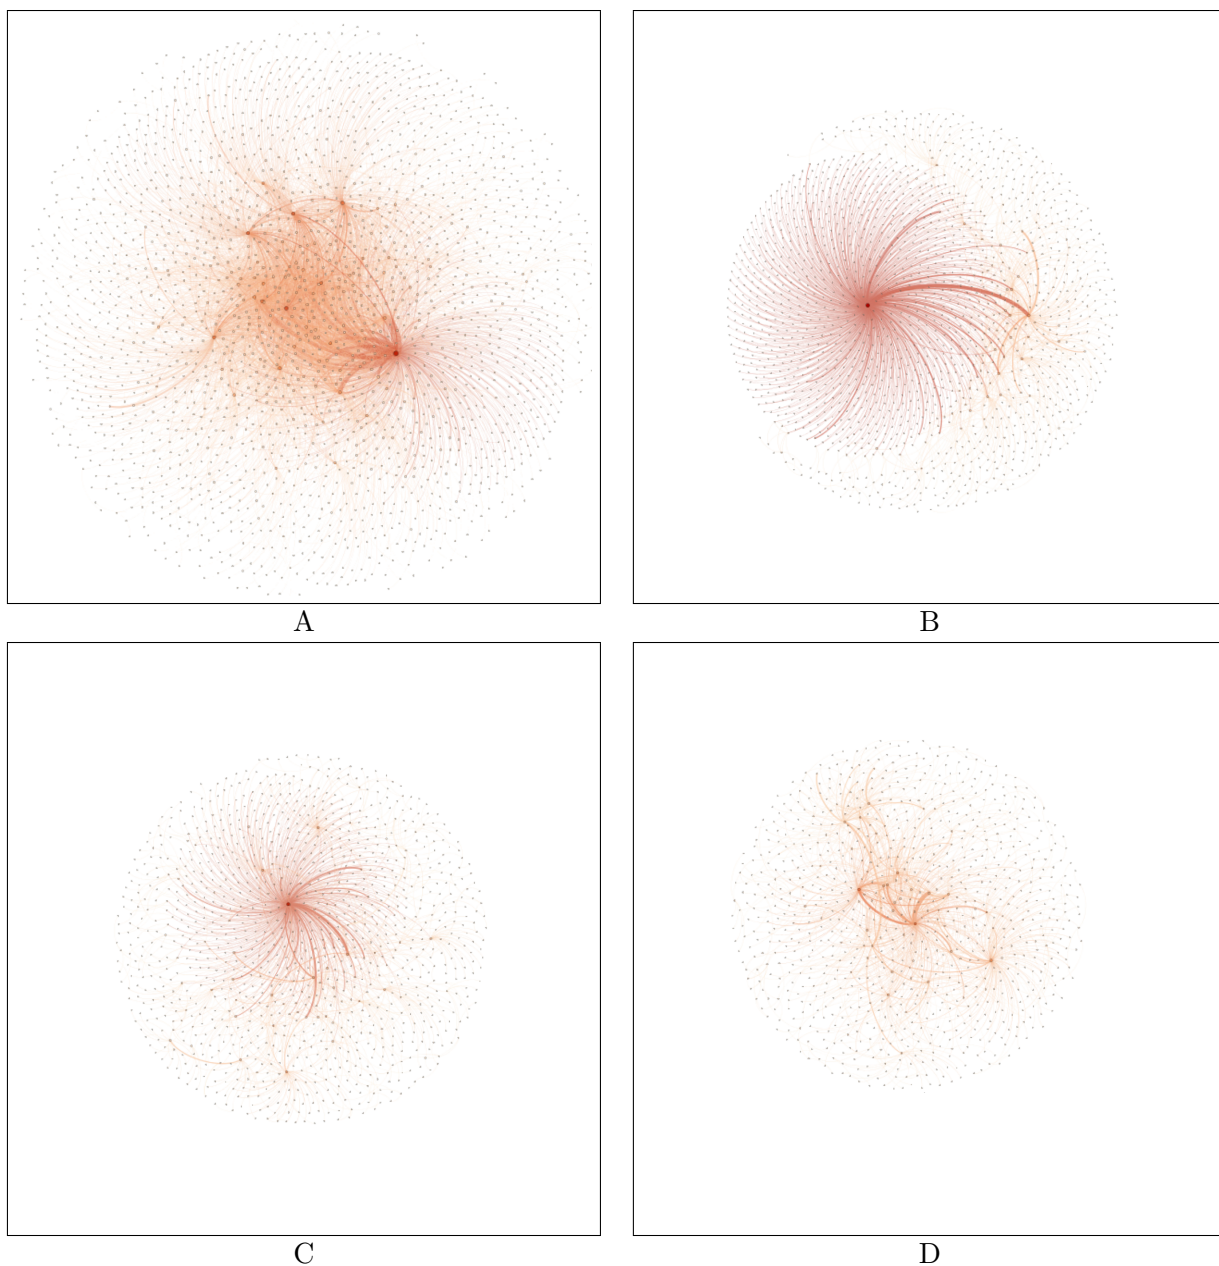


Figure 4.5: The four largest communities detected by modularity in the dataset. The color and size of the nodes reflect their degrees, with larger and darker nodes having relatively higher degrees. (A) consists largely of Colorado social and news organizations; (B) centers around the username of the Denver Broncos and (C) around the Denver 9 News Channel; (D) contains a large number of Colorado Springs official organizations, including emergency personnel and local media sources

networks will only become easier and more rigorous as time goes on, and the utility of such information in crisis situations will only grow. Based on the review of human social networks and especially social media networks presented in this chapter, Chapter 6 introduces a method for generating a synthetic network for interaction.

Part II

Methodology

Chapter 5: Sentiment Detection

Given the importance of emotions to decision-making as described in Chapter 2, there is obvious value in being able to ascertain how individuals and populations are feeling and, therefore, how they may behave. One way to gain a sense of the mood of an impacted population is to attempt to explore the media they produce, translating their output into measurements of sentiment. To that end, sentiment analysis is the codified study of people's opinions, emotions, or attitudes toward a given target. It is a discipline in its own right but also frequently utilized by other fields and industry, where it is sometimes called opinion mining, opinion extraction, sentiment mining, subjectivity analysis, affect analysis, emotion analysis, or review mining (Liu, 2012). Almost as fast-growing as the internet itself, sentiment analysis techniques have come increasingly into vogue and been applied to a wide range of applications in recent years. In the field of crisis informatics, sentiment analysis is particularly interesting because it allows researchers to gain an understanding of elements of the environment that would otherwise be invisible to them. At the individual level, sentiment detection can indicate how stressed or optimistic individuals feel; at a more aggregated level, it can track how the mood of the population is changing. As described in Chapter 3, it can turn individual social media feeds into sensors for an important yet elusive element of the population response to crisis.

The chapter that follows presents a brief history of the development of the field of sentiment detection in Section 5.1, followed by an overview of the ways in which sentiment detection is used in Section 5.2. Section 5.3 presents a review of the most prominent techniques currently being utilized, and Section 5.4 describes some challenges practitioners face. The sentiment detection approach utilized in this work is introduced in Section 5.5, and its effectiveness on a sample of real tweets from the Twitter data set presented in Chapter 4 is demonstrated.

5.1 History

Computational efforts toward extracting meaning, emotion, and opinion from texts have been under the academic microscope for some time: Blood and Phillips (1995) measured the sentiment conveyed by newspaper headlines, while researchers have explored automatically identifying affect (e.g. Huettner & Subasic, 2000) and subjectivity (e.g. Wiebe, 1990; Wiebe et al., 1999). Kantrowitz (2000) filed a patent for techniques dedicated to “analyzing affect and emotion in text” in 2000, but Pang and Lee (2008) suggest that the field essentially opened up in 2001 in light of new advances in machine learning methods and the ready availability of large amounts of data to analyze. Liu (2012) reports that the term sentiment analysis first appeared in the work of Nasukawa and Yi (2003), and the term opinion mining in Dave et al. (2003). Earlier work on sentiment and opinions (Das & Chen, 2001; Pang et al., 2002; Turney, 2002) laid the foundation for these new approaches, Liu (2012) asserts.

Recently, sentiment detection research has spread from the exclusive domain of academia into the often proprietary world of industry. As social media has exploded and the number of online conversations about products and brands has grown too large for even the most dedicated public relations department to follow manually, sentiment detection has become a valuable tool and companies offering sentiment detection have blossomed. From 40 to 60 American standalone companies (Liu, 2012) such as Twitratr, Social Mention, Tweetfeel, Twendz, and Twitter Sentiment provide sentiment detection services on demand (Kouloumpis et al., 2011; Jiang et al., 2011), while bigger corporations including Microsoft, Google, SAS, and Hewlett-Packard have in-house sentiment tracking departments (Liu, 2012). In its current form, the field of sentiment research spans the academia/industry divide and is applied to many problems.

5.2 Uses of Sentiment Detection

On a theoretical level, the utility of sentiment detection is that is capable of analyzing an enormous number of subjective assessments consistently and rapidly, allowing researchers

to consider the aggregate qualities of the population rather than a smaller subset (Hu & Liu, 2004). In the case of Twitter or SMS messages in particular, the researcher often has temporally and sometimes even spatially tagged information about a situation. The large quantity of data involved allows researchers to more confidently make assertions about population-level unrest, aggression, need, approval, or various other emotions of interest to responders or marketers. Given the obvious utility of this enormously powerful tool, it is no surprise that it has been widely adopted by a range of fields.

As sentiment detection has become more and more accessible, its use has spread to an ever widening range of domains. News articles, blogs, and reviews in general have all been subjected to targeted analysis (Glance et al., 2005; Koppel & Shtrimberg, 2006; Lavrenko et al., 2000; Pang et al., 2002; Pang & Lee, 2004; Wiebe & Riloff, 2005; Wilson et al., 2005). Industry has used it for assessing consumer products and services, but the fields of healthcare, financial searches, and political elections have also found the associated methodologies to be useful (Liu, 2012).

Much of the research that has been done has focused on detecting sentiment relative to specific kinds of targets such as products, reviews, or politics; these efforts often further attempt to compare these measurements to some associated trend. Predicting sales (Liu et al., 2007) or ranking products or merchants (McGlohon et al., 2010) are some of the most common applications of sentiment detection. Multiple studies try to glean financial information from investor message boards (Antweiler & Frank 2004; Das & Chen 2007). O'Connor et al. (2010) try to generate opinion polls, while Tumasjan et al. (2010) used the techniques to try to predict election results. Fink et al. (2012) attempt to map sentiment, location, and ethnicity in the service of better understanding the 2011 Nigerian presidential election.

Most applicable to this particular work is the use of sentiment detection in crisis or disaster situations. The open-source crisis mapping platform Ushahidi (initially discussed in Chapter 3) was utilized during the 2010 Haitian earthquake, and sentiment detection was applied to the vast body of SMS messages collected through the program (Heinzelman

& Waters, 2010). Using this kind of automation, humanitarian workers attempted to track stress levels in the aggregate population and thereby gauge the probability of violent outbreaks. Nagy and Stamberger (2012) apply sentiment detection to the Twitter discussion surrounding a gas explosion, theoretically to help manage the official response in light of public reaction. Iftene and Ginsca (2012) explicitly attempt to incorporate sentiment detection into crisis detection efforts, an effort similar to the explicitly disaster-oriented semantic detection work of Tung (2012). The timeliness and broad information-gathering capabilities associated with automated semantic detection are frequently cited by the researchers as benefits of this type of approach.

The works presented above are only a handful of examples taken from the vast body of research dedicated to the theory and application of sentiment detection, selected from the fraction of publicly available studies. Given the applicability and utility of the method, how can it be profitably carried out?

5.3 Analysis Techniques

The analytical techniques applied to the study and extraction of sentiment vary depending on the purposes and resources of the researcher. Section 5.3.1 reviews the types of text different sentiment detection efforts analyze, while Section 5.3.2 reviews the ways sentiment can be measured and quantified. Next, Section 5.3.3 reviews the methodologies researchers have employed to process and classify the data, while Section 5.3.4 and Section 5.3.5 present some of the linguistic features and lexicons utilized in the pursuit of sentiment detection.

5.3.1 Levels of Analysis

When it comes to sentiment analysis, researchers often target the analysis of sentiment at one of three levels: document, sentence, or entity/aspect. The document level attempts to calculate the sentiment associated with the entire document, summarizing the general tone of the piece (as in Pang et al., 2002; Turney, 2002). Detection at the sentence level attempts to capture the sentiment expressed by a particular phrase, outside of the context

of a piece of text (Narayanan et al., 2009; Pang & Lee, 2004; Tsur et al., 2010; Wilson et al., 2004; Yu & Hatzivassiloglou, 2003). Drilling down further, entity/aspect analysis attempts to extract a set of sentiments about either an entity or potentially different aspects of a given entity (Hu & Liu, 2004; Jakob & Gurevych, 2010; Popescu & Etzioni, 2005; Wu et al., 2009). The choice to analyze at a given level depends on the researcher’s reason for study. A researcher looking to determine the precise attitudes of voters toward a candidate (e.g. Fink et al., 2012) or consumers toward a piece of hardware (e.g. Hu & Liu, 2004) would obviously focus on the entity level; a researcher interested in the overall quality of a restaurant or movie would be more interested in summarizing reviews at the document level (e.g. Pang et al., 2002). Because tweets are limited in length to 140 characters, this work employs a sentence-level classification scheme.

5.3.2 Measures of Sentiment

Regardless of the level of analysis, the “value” of the sentiment may be described categorically or numerically, making it a problem of either classification or regression (Liu, 2012). In terms of simple classification, many researchers focus on polarity, measuring the sentiment as positive, negative, or neutral. Others classify the sentiment in terms of a number of categories or moods, trying to capture predefined emotions like happiness, sadness, boredom, fear, and gratitude (Davidov & Tsur, 2010) or calmness, alertness, sureness, vitality, kindness, and happiness (Bollen et al., 2011). Still others try to extract the categories themselves from a body of texts: Mihalcea and Liu (2006) try to measure happiness by deriving word associations with the help of blogger-defined mood labels, and Balog et al. (2006) similarly utilize bloggers’ mood-labeled posts to inform their analysis. Frequently, the form of the results desired dictates the methodology by which the researchers explore the data.

It is also important to note that researchers distinguish between a kind of ambient, target-independent sentiment in a document and more targeted opinions toward specific

entities. In contexts where a number of different sentiments are being expressed simultaneously, specifying the opinion and its target can be important - to take a real example from the data set presented in Section 4.5, the tweet text “The #HighPark fire is the worst Ive seen in my 15yrs in Colorado. Lets commit to praying for the firefighters & 1st responders” is arguably expressing negative attitudes toward the fire but positive attitudes toward response personnel. Many researchers have tried to address this question, and it remains an open topic of research (see Ding & Liu, 2007; Nasukawa & Yi, 2003; Hu & Liu, 2004; Jiang et al., 2011).

5.3.3 Approaches to Sentiment Detection

Here, a distinction is drawn between pure machine learning approaches and natural language processing, and between supervised and unsupervised learning within the category of machine learning. While natural language processing focuses more on the structure of human language and trying to exploit known features of a given language, machine learning is essentially the process of constructing a classification engine from either tagged (supervised) or untagged (unsupervised) pieces of information.

Machine Learning

Wang et al. (2011) report that machine learning approaches tend to have higher recall than natural language-based methods due to the strength of classifiers when it comes to generalizing their results, but that they are only as strong as their training data. While appropriate at the document level, therefore, classifiers may be less effective on short texts like tweets.

Supervised Learning: Given the strength of classifiers at the document level and the increasing availability of high quality tagged data, the bulk of the work that has been done in sentiment detection has used supervised learning (Pang & Lee, 2008; Liu, 2012). This is especially true of sentiment detection that uses regression. Pang et al. (2002) presented one of the foundational works in the field, comparing the effectiveness of three different

methods (Naïve Bayes, Maximum Entropy, and Support Vector Machines [SVM]) given a variety of different features (unigrams, bigrams, adjectives, and part of speech [POS] tags, among others). They found that the SVM classifier used in conjunction with unigrams outperformed the other techniques. Given the success of their methodology, many other researchers have continued along this line of inquiry (Barbosa & Feng, 2010; Wang et al., 2011; Jiang et al., 2011).

Unsupervised Learning: Unsupervised learning has grown in popularity in proportion to the amount of raw material from which relevant lexicons can be derived, and has benefited tremendously from modern innovations in tagging, parsing, and other subsystems which support language processing (Pang & Lee, 2008). One popular approach to unsupervised learning is part of speech (POS) tagging. Turney (2002) uses specific fixed part-of-speech syntactic patterns which are frequently used to express opinions to find sentimental phrases, then calculates a specific sentiment orientation score for each phrase based on the mutual information values between the phrase and predefined positive and negative seed words. Many other unsupervised approaches have built off of this foundation (see Pang & Lee 2008). Pang and Lee (2008) suggest that a strength of POS tagging is its ability to help with word sense disambiguation. Bootstrapping on top of the output of an unsupervised classifier is another technique, frequently associated with finding subjective phrases in particular (e.g., Riloff & Wiebe, 2003). One final popular approach is the lexicon-based method, in which a dictionary of words and phrases with given vectors of sentiment (usually the orientation and strength of the term) are combined with negation and intensification factors in order to generate a sentiment score. This can be done at the document level (Taboada et al., 2011) or at lower levels of analysis (Ding et al., 2008; Hu & Liu, 2004; Kim & Hovy, 2004).

Natural Language Processing

Approaches to natural language processing (NLP) were extremely popular before the 2000s, if for no reason other than the tremendous cost associated with gathering suitable corpuses before the flourishing of the internet and its accompanying deluge of data (Pang & Lee,

2008). In its current usage, NLP is frequently utilized to try to glean more precise and targeted meaning from data. As mentioned before, target-specific sentiment detection frequently utilizes NLP to try to tease out the specific relationships between targets and the writer’s opinion toward them (Jiang et al., 2011). Specifically, Nasukawa and Yi (2003) utilize a syntactic parser and a sentiment lexicon, while Ding and Liu (2007) try to apply linguistic rules to their analysis. Jiang et al. (2011) suggest that the major drawbacks of an NLP are the lack of coverage of the applied linguistic rules and the need for an expert linguist to construct the detection algorithm and expand it as necessary. Thus, while powerful and certainly appropriate for certain contexts, NLP can be an especially difficult technique to implement.

5.3.4 Effective Features for Sentiment Classification

Regardless of the specific methodology, level of analysis, or measure of sentiment, the success of any sentiment detection effort lies in the selection of a set of effective features (Liu, 2012). A few of the most popular - and useful - features include the following:

Terms

The presence of individual words (unigrams) and combinations of words (n-grams) (Pang & Lee, 2004; Barbosa & Feng, 2010).

Parts of speech

For example, Hatzivassiloglou and Wiebe (2000) report that adjectives are indicators of opinions, and Barbosa and Feng (2010) look for the presence or absence of verbs as a sign of subjectivity.

Sentiment shifters

Words or constructions that modify the sentiment of another word, including negations e.g., the “don’t” in “I don’t like”), intensifiers (words such as “more”), or diminishers (words

such as “less”) (Pang et al., 2002; Go et al., 2009; Davidov & Tsur, 2010; Iftene & Ginsca, 2012).

Microblogging features

Abbreviations, context-specific intensifiers (the use of all caps in “SO GREAT”, or repeated letters in “amaaaazing”), and emoticons (Kouloumpis et al., 2011; Nagy & Stamberger, 2012; Barbosa & Feng, 2010).

Punctuation

Davidov and Tsur (2010) consider the presence and number of “?” and “!” marks within tweets.

It bears repeating that the specific target and media being studied hugely influence the types of features that will be useful. Kouloumpis et al. (2011) suggest, for example, that part of speech features may not be useful for sentiment analysis in a microblogging context due to the typical linguistic structures employed. Similarly, any study of texts from before the 1980s would find the inclusion of emoticons unhelpful. Finally, the process of “stemming” words in order to simplify them to their stem or root form has been utilized since the 1970s in order to convert words to more common forms (Dawson, 1974). In cases where the root form is associated with sentiment, this can in theory vastly increase the power of sentiment detection efforts by deriving sentiment from infrequently utilized words: if a sentiment lexicon contains only the word “evacuate” and fails to assign sentiment to “evacuates”, “evacuating”, “evacuee”, a great deal of sentiment information can be lost. A detailed review of stemming algorithms can be found in (Hull, 1996).

5.3.5 Selecting a Lexicon

Frequently, research utilize specific lexicons of sentiment words. Some of the most popular general-purpose lexicons include:

- AFINN (Nielsen, 2011) - used by Nagy & Stamberger (2012)
- General Inquirer (Stone, 1968) - used by Wang et al. (2011)
- MPQA subjectivity lexicon (Wilson et al., 2009) - used by Kouloumpis et al., (2011)
- POMS lexicon - expanded upon and used by Bollen et al. (2011)
- Sentiment Lexicon - developed and used by Hu & Liu (2004)

Dictionaries which focus on slang or abbreviation include the Internet Lingo Dictionary (Kouloumpis et al., 2011) and NoSlang (<http://www.noslang.com/dictionary/full/>). Unfortunately, it can be difficult to take into account for either the accidental or the purposeful misspelling of terms, so that many terms will not appear in any of these records. These kinds of difficulties are particularly aggravated in certain contexts; the casual and linguistically avante-garde nature of Twitter makes for something of a moving target with regard to lexicons. That being said, misspellings are by no means the only challenge researchers face in this context.

5.4 Challenges

Sentiment detection is subject to a number of challenges which bear explicitly noting here. Firstly, it frequently draws from a range of other extremely difficult subproblems, including NLP, machine learning, data mining, and information retrieval. The level of NLP involved in sentiment detection is particularly challenging, Liu (2012) points out, because nearly every aspect of NLP is involved: coreference resolution, negation handling, word sense disambiguation, and other contextual cues, not to mention the challenges of cleaning the typically noisy data. Handling context correctly is a particularly hard problem for researchers. For example, it can be both difficult and important to correctly match sentiment and the entity toward which it is directed - a user who says that a given product “sucks” means very different things depending on whether she is talking about a vacuum or a car. Sentiment can be expressed using no sentiment-laden words at all, especially cases

of implied expectations (“this razor blade lasted for two years” and “this car lasted for two years” do not imply the same level of satisfaction). Liu (2012) asserts that in cases where researchers only seek to determine whether sentiment is positive or negative the task is relatively constrained and therefore easier, but on the whole NLP remains troublesome.

Not only are the underpinnings of the methodologies difficult in and of themselves, practitioners must make a large number of choices regarding what technologies, lexicons, database management software to use, what targets to study, and what level at which to conduct analysis. To take lexicon selection as an example, while many sentiment-tagged lexicons exist, many researchers build upon preexisting lexicons in order to improve their success rates (Barbosa & Feng, 2010; Nagy & Stamberger, 2012), construct them entirely from scratch (Iftene & Ginsca, 2012), or manipulate them in various other ways to better suit their specific purposes (Wang et al., 2011; O’Connor et al., 2010; Bollen et al., 2011). The choice of whether to operate on a unigram level or to incorporate lexicons of phrases (“not worth the paper it’s printed on” and so forth) adds to the burden of choice on the researcher. Adding to the confusion, there is much discussion as to whether different forms of media warrant separate approaches to analysis. For example, Barbosa and Feng (2010) suggest that the successful unigram-heavy approach of Pang and Lee (2004) might not be appropriate for tweets due to the medium’s inherently short length. O’Connor et al. (2010) cite a concern that text analysis techniques designed to optimize the classification accuracy associated with single documents might skew population-level proportions and therefore be inappropriate. Further, all of the techniques described here require the texts they analyze to be minimally cleaned and standardized (Liu, 2012). Given the wildly varying conventions in usage of spelling, grammar, and meaning that characterize text taken from, for example, Twitter versus New York Times articles, a variety of preprocessing methods may be necessary to produce results that can be compared across different text-producing communities.

Sentiment detection is a powerful tool when used correctly, but there is a great deal of latitude for researchers to make bad choices and wrong decisions in terms of how they clean,

analyze, and present their data. It is important to be mindful of these considerations when attempting to design a sentiment detection algorithm.

5.5 Methodology Employed in this Work

The methodology employed here is very simple and is meant to be applicable to a wide range of target topics with very low overhead costs. It is applied to the data set initially introduced in Section 4.5. As mentioned in Section 5.3.1, the analysis proceeds at the level of the sentence, given the 140-character restricted length of tweets. Given the desirable emphasis on rapid classification and the extensibility of the methodology to non-English languages, the goal was to construct a framework into which one could potentially substitute a rapidly-constructed, non-English valence list. In general, the work follows the framework laid down by Dodds et al. (2011). Of the features mentioned in Section 5.3.4, terms and microblogging features such as emoticons are utilized to drive the calculation, while experiments with sentiment shifters ultimately proved not to be particularly fruitful. Section 5.3.2 addressed some of the questions associated with trying to express sentiment as a measurement: in this work, the sentiment measures produced ranged from 0 to 10, with 0 being the cutoff for the most negative and 10 the cutoff for the most positive values.

5.5.1 The Data

In an effort to capture a large number of emotionally-charged tweets, a set of sample tweets is selected from the population of unique tweets posted on June 26, 2012, a day on which over 25000 citizens were evacuated from Colorado Springs. From this corpus of unique tweets, 300 random tweets are selected with which to develop the methodology and 300 other, separate random tweets were selected to test the effectiveness of the developed method. These two sets of tweets are referred to as the training and testing set. The training set was utilized as part of the effort to develop the method, and was frequently observed during the process of the development of the method in order to improve classification efforts. The testing set was kept in reserve until it was used to quantify the success of the method on a

completely unknown set of tweets. For the testing set, the author went through and marked each tweet as positive, negative, or neutral. These classifications are treated as the gold standard against which the effectiveness of the automatic classifier is judged. In all, the testing set included 31 positive texts, 163 neutral texts, and 106 negative texts; the success of the method in classifying these tweets is presented in Section 5.5.3.

Lexicon

Given the importance of the lexicon to the success of the method discussed in Section 5.3.5, this work presents a comparison of the effectiveness of three different sentiment-weighted lexicons, including a modified version of one of the sets. Specifically, the lexicons include AFINN (http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=6010) and SentiWordNet (<http://sentiwordnet.isti.cnr.it/>) wordlists. The AFINN word list ranks all of its terms with integer-value valences of positivity or negativity ranging from -5 (most negative) to +5 (most positive). SentiWordNet, on the other hand, measures the positivity, negativity, and neutrality of a word. Because SentiWordNet provides multiple definitions for the same word depending on part of speech and POS tagging is not implemented in this work, the sentiment detectors utilizes the “primary” definition of a word.

In addition to the existing terminology, the lexicon has been expanded with several sentiment-heavy terms that appear frequently in the text to the AFINN results. These words include terms such as “ugh” and “yikes”, or online shorthand such as the positive “lol” (“laughing out loud” or less commonly “lots of love”) and the negative “smh” (“shaking my head”, used to indicate unhappiness or frustration). Additionally, because the search term used to collect tweets was “fire”, itself coded as being a -2 valence word in the AFINN lexicon, in the modified AFINN dictionary the term “fire” has been excluded from the sentiment calculation. The effectiveness of this modified dictionary in increasing the accuracy of the classification process is addressed in the results section.

In addition to the wordlists, the sentiment detector also utilizes a list of emoticons partially derived from Tung (2012), building on this basis with information from the same

source as Tung originally used (en.wikipedia.org/wiki/List_of_emoticons). This list has been expanded upon as well (making the assumption that “:-)” and “:.)” are semantically equivalent, for example).

5.5.2 The Algorithm

The sentiment calculation presented here is equipped with a number of options, all of which are explored and compared in order to maximize the information content extracted. Thus, the metric is capable of analyzing the information with different types of normalization, negation, and stemming applied to the calculation. After the phrase is stripped of punctuation and URLs, the aforementioned set of stopwords is removed and all of the terms are converted to lower case. Punctuation which relates to emoticons is preserved, and groups of emoticons are treated as valence-bearing terms in their own right. If word stemming has been activated, it is at this point that all terms are stemmed.

For each of the N terms which remain of the phrase, each word w_i with valence $v(w_i)$ is manipulated in accordance with the given normalization factor (p_i , discussed below) and, depending on whether the $i - 1$ term is a negation, negation (represented in the following as n_{i-1} , the negation-status of the previous term). The tested metrics are structured as follows, although the various parameters may be set to default values in order to “deactivate” any of the types of functionality:

$$sentiment(phrase) = \sum_{i=1}^N n_{i-1} v(w_i) p_i \quad (5.1)$$

In situations where no valence-bearing terms exist, the phrase is assumed to have perfectly neutral sentiment.

Normalization

The work done here tests three different types of normalization of terms: normalization by the number of valence-bearing terms in the phrase, normalization by the frequency of

the use of the term in the overall corpus, and the lack of any normalization at all. For the frequency-based normalization, the frequencies were derived from the corpus itself. Of the millions of unique words drawn from the corpus of tweets, only the 4000 most common terms are assigned a weight; a default “minimum” frequency is assigned to all others. This cutoff corresponds approximately with frequencies lower than .001% in the corpus-derived body of terms. The top 4000 terms were selected because frequency tailed off substantially after that point. For frequency-based normalization, the term p_i is derived as follows:

$$p_i = \frac{f_i}{\sum_{j=1}^N f_j} \quad (5.2)$$

In the case of phrase-based normalization, p_i is simply set to be $\frac{1}{N}$, and in the case of no normalization p_i is 1.

Negation

In addition to comparing the effectiveness of the different valence-bearing lexicons, the impact of a simple negation function was investigated. If any of the words “not”, “cannot”, “never”, “can’t”, “won’t”, “ain’t”, “don’t”, or “didn’t” appear directly before another valence-bearing term, that term has its valence flipped before being incorporated into the calculation. That is, if a word has valence v and is directly preceded by “won’t”, the term is temporarily assigned valence $10 - v$ for use within the metric.

Stemmer

A simple implementation of a word stemmer is taken from the SnowballStemmer project (<http://snowball.tartarus.org/>), which is a rapid and Java-based option. The stemming process happens before the valence is calculated, and potentially allows for less common conjugations and forms of certain words to be identified as valence-bearing.

5.5.3 Results

The results presented here indicate the percentage of all texts assigned to the correct category - positive, neutral, or negative. In Tables 5.1-5.3, the successful classifications are all grouped together, while the more successful combinations of parameters are explored in greater detail later. Tweets are regarded as successfully classified if they are less than, precisely equal to, and greater than 5 for negative, neutral, and positive classifications, respectively.

Table 5.1: AFINN results comparison: total success scores

	Stemmer		No Stemmer	
	Negation	No Negation	Negation	No Negation
Raw	0.590	0.597	0.620	0.623
Frequency	0.567	0.570	0.610	0.613
Count	0.313	0.320	0.347	0.353

Table 5.2: SentiWordNet results comparison: total success scores

	Stemmer		No Stemmer	
	Negation	No Negation	Negation	No Negation
Raw	0.463	0.460	0.497	0.503
Frequency	0.500	0.493	0.517	0.517
Count	0.540	0.537	0.573	0.580

From these results, it seems the overall AFINN does the best job, and that it performs best of all with no stemming or negation involved. When SentiWordNet, negation is useful in conjunction with the stemmer, but overall the best results come from the count-normalized, unstemmed, non-negation option. The ANEW lexicon produces similar results: the unstemmed options are preferable, although negation does not impact performance when

Table 5.3: ANEW results comparison: total success scores

	Stemmer		No Stemmer	
	Negation	No Negation	Negation	No Negation
Raw	0.483	0.483	0.513	0.510
Freq	0.503	0.503	0.550	0.547
Count	0.487	0.487	0.533	0.530

used with stemming and actually improves the performance when no stemmer is utilized. Thus, overall the best option seems to be the AFINN lexicon.

Table 5.4: AFINN results comparison: breakdown by classification scores

		Stemmer		No Stemmer	
		Negation	No Negation	Negation	No Negation
Raw	Total	0.590	0.597	0.620	0.623
	Positive	0.516	0.516	0.677	0.677
	Neutral	0.748	0.755	0.706	0.706
	Negative	0.368	0.377	0.472	0.481
Frequency	Total	0.567	0.570	0.610	0.613
	Positive	0.484	0.484	0.613	0.613
	Neutral	0.748	0.755	0.724	0.730
	Negative	0.311	0.311	0.434	0.434
Count	Total	0.313	0.320	0.347	0.353
	Positive	0.419	0.419	0.613	0.613
	Neutral	0.018	0.025	0.037	0.043
	Negative	0.736	0.745	0.745	0.755

Diving further into the performance of AFINN in Table 5.4, the metric shows very different success rates with regard to the different classifications. For example, the stemmed, raw, negation-enabled neutral detection rate is 74.8%, while it catches only 51.6% of the positive and 36.8% of the negative phrases. This contrasts with the unstemmed, count-driven, negation-disabled detection rate, which is 61.3% for positives and 75.5% for negatives, but only 4.3% for neutral phrases. The 62.3% overall success rate of the unstemmed, raw, negation-disabled AFINN option compares favorably with the other options, making

it the best choice of the options. From this point, all references to the sentiment detection algorithm will refer to the unstemmed, raw, negation-disabled AFINN option.

5.6 Sentiment Analysis of Real Data

To give a sense of the range of items that were similarly classified, Table 5.5 presents a comparison of items which have been classified by the sentiment detection algorithm. Each of these cases was taken from the real dataset, and were selected to show the range of items which were classified in one way or another. As expected, the algorithm has some troubles: the phrase “fraught with drama!” is arguably not a natural phrase; nor is a report on a healthy puppy which includes a smiley emoticon. Positively-classified examples include the dubiously classified “hope...we don’t die horribly in fire” as well as the more obvious and correct “thank you!!”. Negative classifications were more strongly negative, with “firefighter scam” being perhaps the most neutral example. Thus, there is certainly room for improvement in the algorithm, especially with regard to the much sought-after sarcasm detection (see Pang & Lee, 2008; Mejova, 2009; Liu, 2012). However, overall the negative results were qualitatively extremely reasonable, as were the (nonsarcastic) positive classifications.

Given this sampling of actual results, the metric is extended to the set of all data collected over the period of time. Figures 5.1- 5.4 present the raw and log-scaled number of tweets of each valence generated over the entire period of time sampled and the peak of the Waldo Canyon fire specifically. Specifically, Figure 5.1 shows the sentiment values associated with all tweets over the entire period that was sampled. It is clear that, over the course of longer periods of time, relatively neutral phrases are the norm - there is certainly an increase in more extremely charged phrases during spikes of activity, but the texts are overwhelmingly neutral. This trend is even more obvious in the log-scaled sample data shown in Figure 5.2, which presents the changes on a log scale to emphasize the change in orders of magnitude of the difference valence measures relative to one another. The extremes of sentiment are relatively infrequently observed: the more neutral the sentiment




Figure 5.1: Tweet valences over the entire sample period




Figure 5.2: Tweet valences over the entire sample period, log-scaled to give a sense in the change of rank

Table 5.5: Comparison of valence classifications of real texts

Negative	
0	And it gets more bizarre: Fake #highParkFire fire-fighter pulled a similar scam on #LowerNorthForkFire: http://t.co/MV64BdKV
2	@AKbirder I have my stuff ready. Very uneasy with the current situation w/fire to the north. Most pics are on hard drives. Ready to go.
Neutral	
5	puppy is doin much better =)
5	Whew. Dropping off the Youngest at camp was fraught with drama! Highway closed behind me, due to forest fire on Hwy 50. #Colorado
Positive	
7	Fire Fighters Rock - Thank You !! http://t.co/Kiqr5EUP
10	Dear Universe, Hi, hope everything works out & we don't die horribly in fire.. Thanks.. Sincerely, ME

measure, the more frequently it is observed, without any noticeable exception. Obviously there is a bias toward this in the data - both positive and negative texts are more likely to be misclassified than neutral texts - but the relationships among the different extremes seems to stay quite constant over time. The relative rank of valences over the entire sample period is noisy but consistent.

Turning to the most intensive period of discussion of the wildfire event, Figure 5.3 shows the comparative rates of tweets generated over the span of the Waldo Canyon Wildfire's activity period, from June 21, 2012, until July 12, 2012, two days after the official announcement that the wildfire had been contained. The valences over the peak period show a few interesting trends, including the relatively lower rates of strongly negative tweets in the immediate aftermath of the evacuations. The relationships among the valences stayed relatively constant over the entire sample period, but the peak period does show interesting trends in the extremes of the valence measures. Considering the log-scaled peak period values highlighted in 5.4, the uncharacteristic spike in 0 valence tweets relative to the other extreme valences on June 29 corresponds with a piece of information that was released and went viral on that day - media outlets announced the discovery of a body in an evacuated

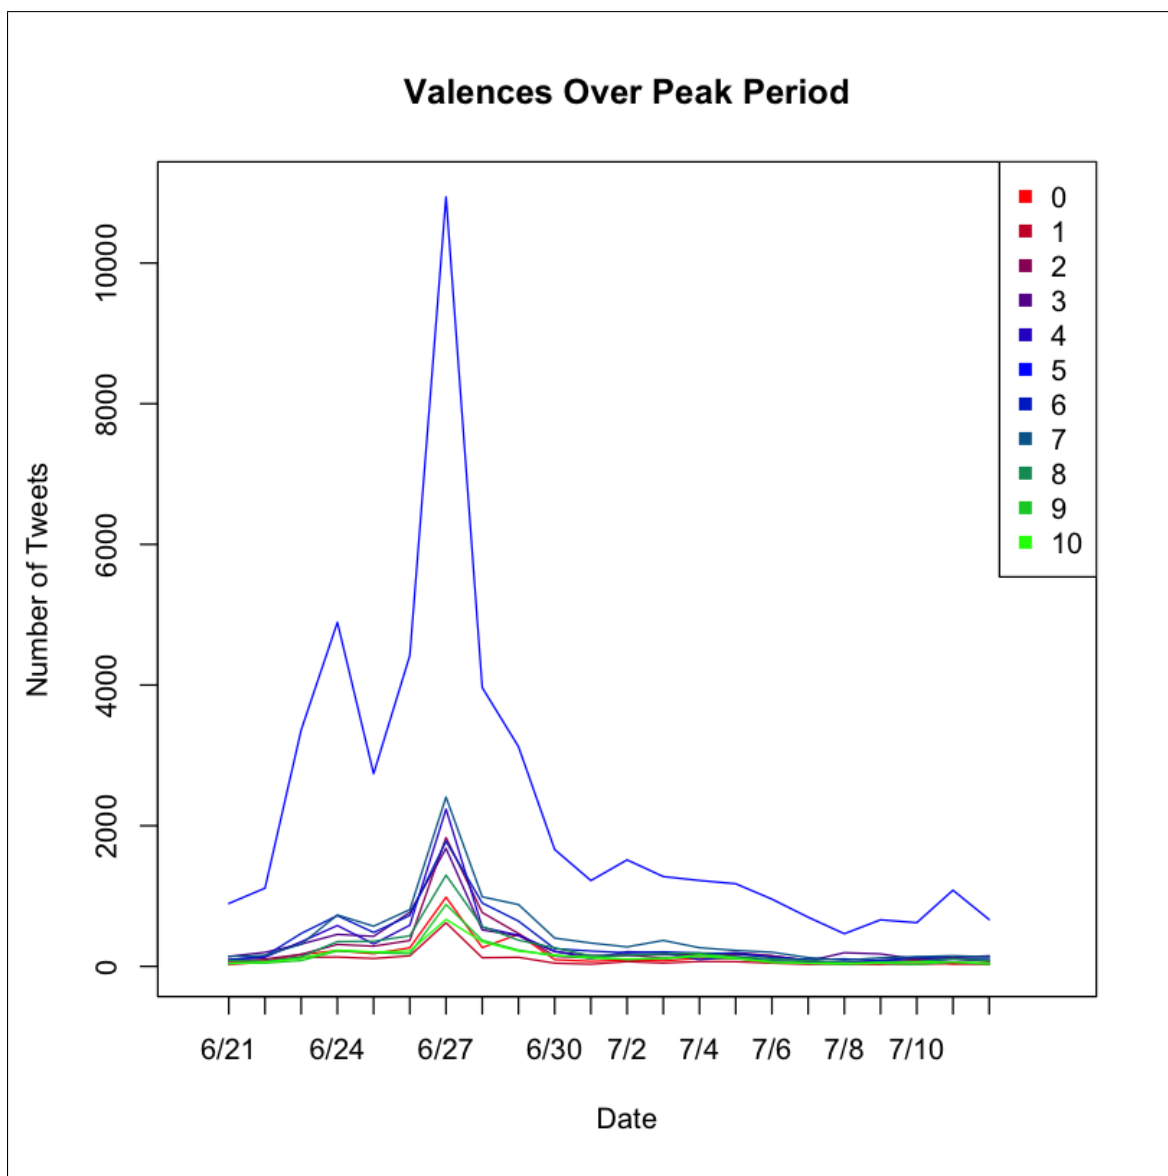


Figure 5.3: Tweet valences over the peak period of the Waldo Canyon wildfire evacuations




Figure 5.4: Tweet valences over the peak period of the Waldo Canyon wildfire evacuations, log-scaled to give a sense in the change of rank

area (“RT @9NEWS: We are sad to report a death in the Waldo Canyon Fire. A body was discovered in a home in the 2900 block of... [url]” or “RT @kktv11news: 2910 Rossmere St. Home destroyed by fire. 1 person found dead inside. #WaldoCanyonFire”). In the immediate aftermath of the evacuations, however, the extremely positive valences outpace the extremely negative valences - as the fire was contained, there emerged highly retweeted phrases such as the July 1 “RT @DenverChannel: Firefighters say they made good progress at the #WaldoCanyonFire Sunday. Fire is 45% contained.” As progress was made in containing the fire, the positive tweets continued, but dropped off as houses were released for reentry conversation fell away from the Waldo Canyon wildfire.

Overall, a simple sentiment analysis of the tweets generated during the period of time under study suggests that while the volume of Twitter communication increases in a time of crisis, the relationships between the difference valence classes stay relatively constant but can see mild reorderings as certain pieces of information are introduced to the system. Much of the on-topic discussion of the wildfires was relatively positive, thanking firefighters for their efforts and encouraging others to contribute to wildfire-specific charities, perhaps because the conversation included individuals who were not themselves in the affected area. However, the overwhelming majority of the conversation was relatively neutral in tone, suggesting that the increase in chatter is the most characteristic signature of the crisis event.

5.7 Summary

Based on the body of existing sentiment detection efforts, this chapter has presented a simple method of sentiment detection which is capable of rapidly parsing a large set of uncleaned data and producing broad but satisfactory classifications of their emotional content. From these coarse measurements of emotional valence and the exploration of how it varies at the population level, it is possible to explore how population levels of sentiment vary over time. The dramatic increase in communication, in both absolute and order of magnitude terms, is an important dynamic to be aware of in constructing a model of communication in crisis

situations. Despite the dramatic increase in information sharing, the sampled data suggest that much of the texts produced even at the height of the crisis are relatively neutral in tone, perhaps reflecting Twitter’s role as an information-sharing network rather than a forum for strictly intimate interpersonal communication, as discussed in Chapter 4. The implications of this study inform the design of the communication patterns in Chapter 7, and serve as a point of comparison for the results presented in Chapter 8. By building upon the work done here, it is possible to capture lower-level sentiment and to incorporate that into individual behaviors in the model presented in Chapter 7.

Chapter 6: Population Synthesis

One of the challenges agent-based modeling faces is obtaining a realistic population upon which to base the simulation. Acquiring good data is frequently a challenge in research efforts, and many techniques have been devised in order to address this need. Agent-based modeling in particular is a data hungry methodology, as the model can incorporate theoretically boundless heterogeneity, and consequently can be designed to take extremely rich data. As discussed in Chapter 3, in some cases that data may simply not exist, in which case it is necessary to generate a synthetic population with realistic aggregate characteristics. Even if population data does exist, it is impossible to know certain important aspects of the population - it is unlikely that a neighborhood demographic survey will capture which individuals will turn to one another for information or assistance, nor will it reveal who talks to whom on Facebook. However, these structures and relationships shape the way information is exchanged and communities respond. Considering the emerging importance of social media to the spread of information in crisis situations in particular, as highlighted in Chapter 3, capturing these dynamics is of crucial importance to a model. This chapter will present the process by which this work creates the set of heterogeneous individuals which populate the simulation.

In Section 6.1, a discussion of population synthesis as a field provides background for the initial stage of the generation of individuals, contextualizing the current work relative to the state of the art. Following this section are descriptions of each of the steps involved in creating the relevant population, complete with samples of the generated data and progressive validation. Specifically, Section 6.2 gives an overview of the data sources utilized in the simulation. Section 6.3 reviews the process whereby houses are generated before, in Section 6.4, a set of individuals is created and processed into households which are assigned

to the houses. Section 6.5 outlines the creation of workplaces and how individuals are assigned to jobs. Individuals are connect to one another, both by intimate social networks in Section 6.6 and by information-sharing networks in Section 6.7. These networks allow them to communicate with one another during the course of the simulation, and guide the flow of information at a distance. As a final step, in Section 6.8 the population of individuals is merged into behavioral units which form the agents described in Chapter 7. The justification and presentation of results at each step delineates the work done and the assumptions made, so that the entire process of creating agents is presented in this chapter. In all, the process generates a sample of a little more than 170,000 individuals assigned to 78,000 households. The households are distributed among approximately 91,000 possible houses, and employed individuals work at one of 130,000 possible job sites. The area for which this population is generated is approximately 20km by 20km. With this sense of scale, the following methodology is presented. The code used to carry out these processes is included in the repository of code associated with this project and is available at www.css.gmu.edu/swise/thesis.

6.1 Introduction

Planners and researchers sometimes find themselves faced with a problem for which the heterogeneous and individual-level characteristics of people matter a great deal, yet without information about who in particular has those characteristics and in what combinations. Even in situations where data about the characteristics of an entire population is available, questions of privacy usually prevent researchers from accessing it (Barthelemy & Toint, 2012). Trying to estimate questions regarding health, transportation, or need for government services requires a certain amount of personalized information - for example, two similarly-sized populations with equal numbers of children will demand very different kinds of government support if, in one, almost all of the children are concentrated in households in the lowest income bracket, while in the other children are relatively uniformly distributed over all income brackets. However, information about the makeup of a population is usually

collected only infrequently or in staggered survey efforts, (Müller & Axhausen, 2010) and in highly changeable and unsettled situations such as refugee camps or developing countries perhaps not at all. Thus, it becomes necessary to try to generate a kind of synthetic population on which to test ideas, policies, and structures. As the tools for generating synthetic populations have become more accessible, and as Census and survey information has made the raw information that supports such research have become both more extensive and more accessible, synthetic population generation has blossomed as a field.

In 1957, Orcutt proposed the development of simulation models for the creation of synthetic populations for use in policy research, noting that “current models of our socio-economic system only predict aggregates and fail to predict distributions of individuals, households, or firms in single or multi-variate classifications.” His paper is regarded as the genesis of the microsimulation method, which consists of taking a representative sample, creating a hypothetical sample with the attributes in question, and classifying their attributes under the new set of conditions (Gilbert & Troitzch, 2005). In general, microsimulation models can be distinguished as either being static or dynamic, where static models simply reweigh the attributes depending on the process being modeled (see Tomintz et al., 2008) and dynamic models iterate a population through space and time, aging the sample and subjecting the population to a set of influences which may probabilistically cause them to transition between attribute states (see Birkin & Wu, 2012). Dynamic microsimulations have obvious applications to the question of synthetic population generation, and are frequently used for this purpose (see Ballas et al., 2005; Harland et al., 2012). However, a pure microsimulation approach is not pursued in this work because, by its construction, microsimulation does not allow for agent to agent or agent to environment interactions (Gilbert & Troitzsch, 2005).

As it is currently practiced, synthetic population generation is applicable to a wide range of fields, including estimating taxes, benefit payouts, pensions, and questions of health and transportation (see Birkin & Wu, 2012). Synthetic populations of other kinds of units exist as well: Müller and Axhausen (2010) highlight the similar generation and grouping of

employees into firms, vehicles associated with households, and tenants with their buildings. It is important to note also that population generation can attempt to work at different levels, generating not only populations of individuals but also households according to given constraints (Arentze et al., 2007). Techniques may emphasize or completely ignore spatiality as a component of their considerations, with some researchers trying to account for heterogeneity between simulated regions (Arentze et al., 2007; Beckman et al., 1996; Müller & Axhausen, 2010). And of course there are many different approaches to the generation process itself.

To broadly outline the major aspects of synthetic population generation along the methods of Müller & Axhausen (2010), there are two steps to synthetic population generation, namely fitting the population (that is, establishing the relationships among the aggregated constraints) and then allocating it (generating individual units based on the fitted population). Approaches to synthetic population generation tend to be broken down into two categories: synthetic reconstruction and combinatorial optimization (Barthelemy & Toint, 2012). However, both of these approaches require that the researcher have on hand a sample of the population as well as aggregate statistics about the population taken from a source other than the sample, resources which are not necessarily available. In light of these dependencies, some other techniques have been developed. Below are some of the major methodologies and a few examples of their usage.

6.1.1 Synthetic Reconstruction

Synthetic reconstruction deals with using random sampling to generate individuals from a set of conditional probabilities. This technique allows for researchers to utilize constraints derived from a wide range of sources, as the conditional probabilities can be calculated from diverse tables of data (Huang & Williamson, 2001). Simply put, the reconstruction consists of calculating the joint distribution of characteristics of interest and then randomly drawing from it to add individuals to the population. After the weight associated with each grouping of characteristics is calculated, the allocation process is carried out by drawing from the

distribution of groups, the drawing being handled by Monte Carlo random drawings, greedy deterministic methods, or various other selection schemes (Müller & Axhausen, 2010).

The calculation of joint distributions is frequently carried out via iterative proportional fitting (IPF). First introduced by Deming and Stephan (1940), IPF has been shown to minimize the relative entropy and preserve cross-produce ratios, meaning that the resulting table generated from the initial table is the most similar possible table that fits all of the constraints upon it (Müller & Axhausen, 2010). It is the most widely used approach to the satisfaction of constraints (Frazier & Alfons, 2012), although it has a number of drawbacks: Frazier and Alfons (2012) note the difficulty IPF has in matching hierarchical distributions. IPF also suffers from the “zero cell” problem, whereby a Census might pick up on an unusual combination of traits that a sample does not capture, creating problems when the marginal count is greater than zero although the conditional probabilities suggest that every cell in the row should in fact be zero. Finally, internally inconsistent target constraints can prevent the system from ever fully equalizing, and only in very simple situations can heuristics address this problem (Rich & Mulalic, 2012).

Numerous workarounds exist to make IPF applicable to various kinds of problems. Arentze et al. (2007) perform a two-step IPF on individuals and households in order to generate household units; this is necessary because a basic run of IPF can work at only one level at a time, (Müller & Axhausen, 2010). Rich and Mulalic (2012) address the problem of contradictory constraints, do a two-step cleaning operation to ensure that the IPF will eventually terminate. Spatiality can also be incorporated by pushing separate zones through the IPF and comparing the aggregated output with the population constraints (Beckman et al., 1996; Müller & Axhausen, 2010).

6.1.2 Combinatorial Optimization

While synthetic population generation aims to create a population from a set of constraints, combinatorial optimization essentially creates a population and then modifies it until it meets a set of constraints. Thus, the aim of combinatorial optimization is to rapidly generate

a population by creating different combinations of units taken from a set of survey data. The key to the methodology is to take, for example, an initial set of households from the survey data and to iteratively consider switching random households from the generated population with households from the survey. If the switch would bring the generated population more closely into line with the higher-level constraints, the switch is carried out, and the process continues until the generated population is within a pre-defined threshold of the required constraints (Huang & Williamson, 2001).

Combinatorial optimization also suffers from the “zero cell” problem (Frazier & Alfons, 2012), because the real population is at least as heterogeneous as the sample and a combination of characteristics must appear in the sample to appear in the generated population. It also has the challenge of trying to fit populations that vary in composition over space; for example, it has difficulty identifying smaller communities of religious or ethnic groups which are clustered in small areas of the city, drastically varying from the higher-level constraints. However, Huang and Williamson (2001) note that so long as a joint distribution table exists for the divergent areas, the problem can be addressed by generating populations for the most divergent areas first and maintaining those subpopulations as the rest of the population is generated. The approach has low memory requirements, but is time consuming as the population converges to an acceptable composition (Farooq et al., 2013).

6.1.3 Other Methodologies

Methodologies other than pure synthetic reconstruction and combinatorial optimization exist. Frequently, these alternative approaches have been developed in light of the extensive data needs of these two main approaches, not to mention their failure to capture the underlying heterogeneity of the population. Some of these efforts merely tweak one or the other method, for example by relying more heavily on data to make up for the lack of aggregate information. In one such case, Mussavi Rizi et al. (2013) fuse remote sensing data with surveys collected via various sampling techniques, adopting a combinatorial optimization-type approach to generate households in villages according to the villages unassigned land

wealth. Barthelemy and Toint (2012) do away with the sample altogether by generating individuals, then calculating joint distributions of household types and assembling the generated individuals into households. The extensibility of Barthelemy and Toint's (2012) approach, at least, has been questioned; for example, the difficulties involved in designing the context-specific matching rules that assign individuals to households and the multi-level hierarchical fitting efforts may offset the benefits of forgoing a sample (Farooq et al., 2013). However, in situations where no sample exists, the methodology is a powerful tool. The flowering of population synthesis approaches in recent years suggests a growing trend, and the sample-less reconstruction efforts will become more and more important as they are applied to displaced populations, rapidly-changing areas, and otherwise marginalized and poorly understood groups.

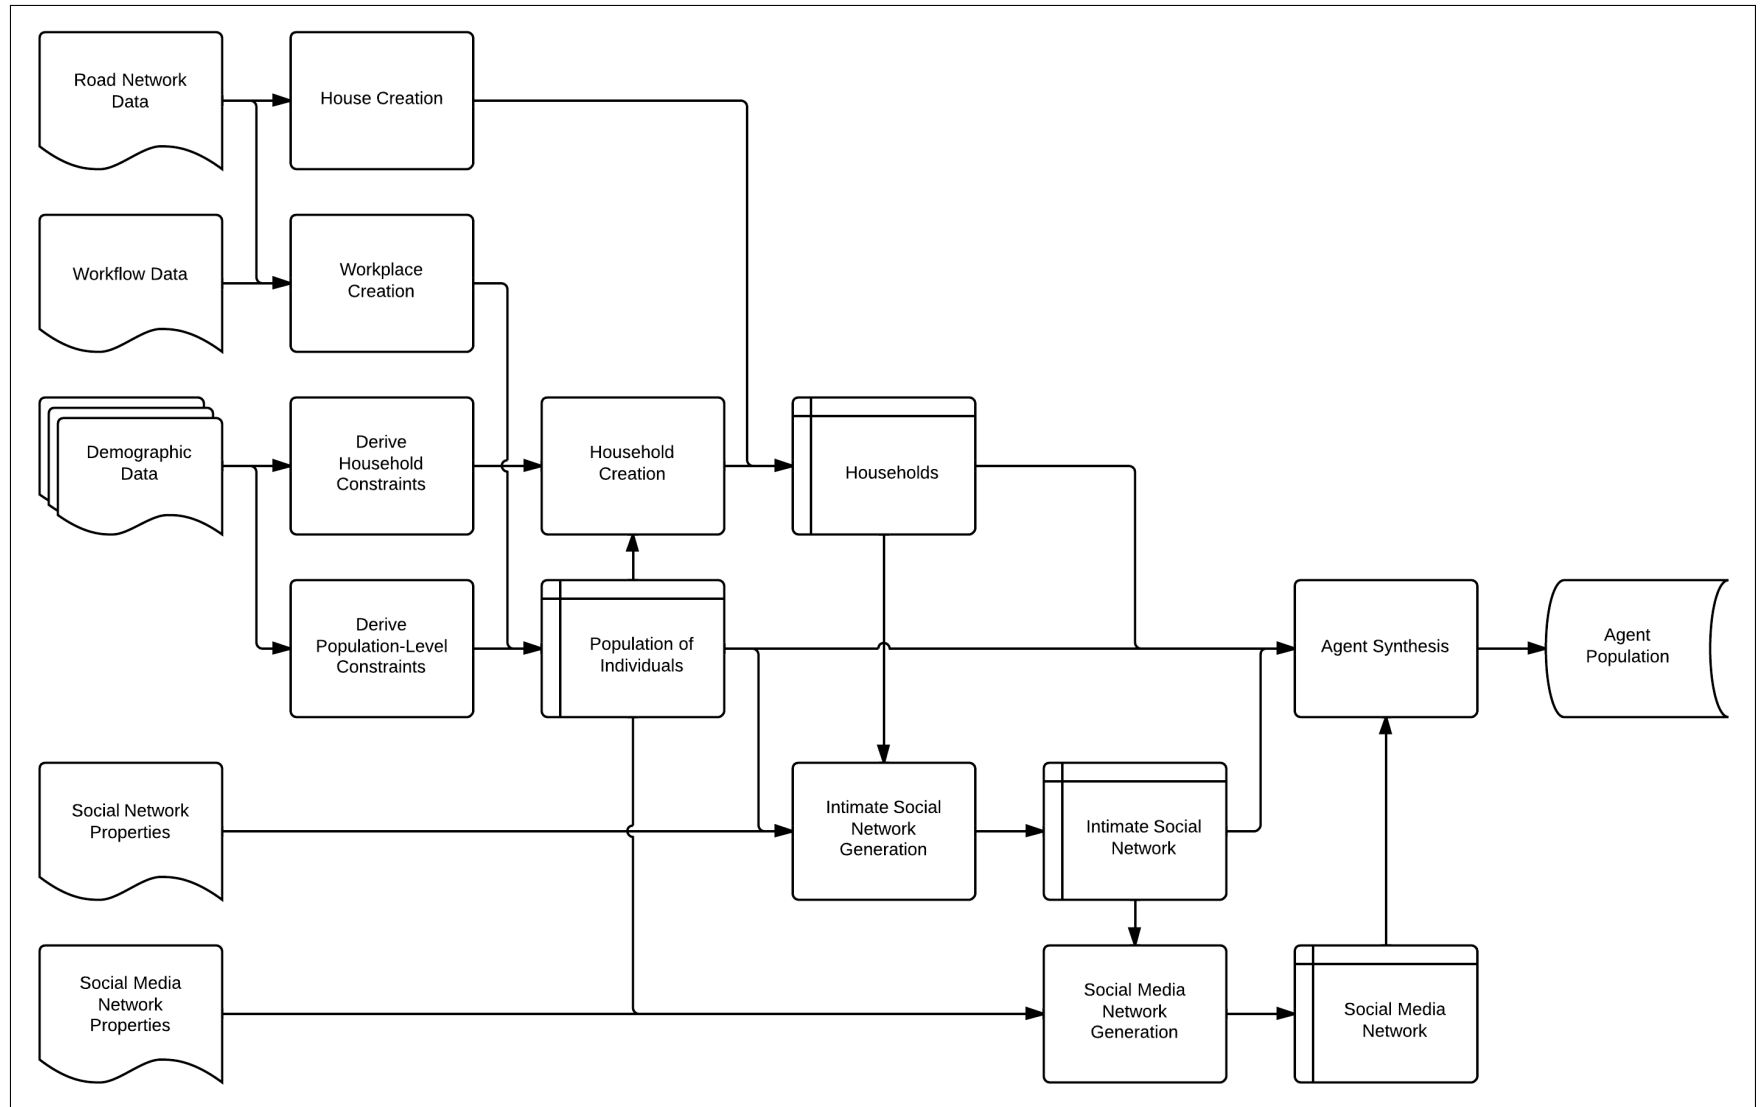


Figure 6.1: A flowchart of the population synthesis process

6.1.4 Methodology Utilized in this Work

This work follows the example given by Barthelemy and Toint (2012), both because it is more extensible to the kinds of data-sparse situations to which this framework may hopefully be applied and because of the quality of the results it produces. Given that this work is intended to be applicable in situations where no sample data is available, the Barthelemy and Toint (2012) method is the most powerful and applicable given the kinds of datasets that will be accessible. Thus, by applying the methodology to a population that is relatively well understood, it is possible to gain a sense of the quality of the methodology as well as to build toward a tool that is useable in less well documented circumstances.

Figure 6.1 shows the process by which the synthetic population is generated. The population synthesis process begins offline, with gathering and then manipulating the data into a usable form. The data is fed into the program, which derives population-level constraints. Based on the population constraints, a set of individuals is generated. The individuals are then assembled into households based on the statistical characteristics of households given by the input data. In this way, both the constraints on the higher-order structures of households and the individual characteristics including age and sex are accurately represented in the final population. Households are assigned to houses, and workers are assigned to workplaces. Intimate social networks are generated, connecting individuals to one another; drawing upon these relationships, a representative population of individuals are connected to one another via social media networks. While the social media networks inform their construction based on the already-constructed intimate social networks, they are distinct phenomena, and represent the intimate social ties between Agents (the intimate social network) and the information-sharing networks in which Agents are embedded (the social media network) respectively. However, the distinct networks represented by the simulation reflect different modes of communication, and the overlap between them reflects the fact that there are intentional and substantial redundancies between them. An overview of all the stages of this population generation process as operationalized for the test case is presented below. The sections are described in the order in which their associated processes

Table 6.1: Types of data required by the simulation

Data Type	Requirements
Demographic	size of population, proportion of individuals by age and sex, average household size, proportion of household types, presence of non-household group populations; variations among all these factors
Road Network	road geometries, tagged by usage
Employment Flows	calculation of commuter flows between areas
Social Media Network	information about social media usage in the area, broken down by demographic characteristics

are carried out.

6.2 Data Gathering

The population constraint data were derived from United States Census Bureau 2010 Census demographic information (2011). In addition to the population constraints, information about the road network, Census employment population flows, intimate social network structure, and social media usage were utilized. The specifics of these datasets will be discussed in the relevant sections, but the general requirements which must be provided, created, or assumed are summarized in Table 6.1. The generation of the various social networks is not drawn from explicit data records and instead proceeds from the first principles, the literature of which will specifically be referenced in those sections.

6.3 House Creation

The first step of the population generation effort involves creating the houses in which the households live. By combining the information from OpenStreetMap about residential streets versus highways and commercial areas with the expected number of households in an area, it is possible to generate a set of residences along streets which are specifically

labelled as residential spaces, creating a reasonable set of domiciles in absence of building footprints.

6.3.1 Derive Household Constraints

Generating houses requires the fusion of two different sources of data, namely the United States Census Bureau 2010 Census demographic data set (2011) and a set of road shapefiles drawn from OpenStreetMap on October 22, 2012. The specific Census record of the number of housing units per tract is DP0180001, and residential roads are those whose OpenStreetMap-designated “type” attribute is set to “residential”. If the road has no associated “type” attribute, it is counted as being potentially residential.

6.3.2 The Process

The Census data includes a record of the number of homes within a given Census tract, as well as the total population living in group housing, but the shapefiles distributed by the United States Census Bureau do not include the specific geographic locations of households within the tract. In an attempt to automatically generate reasonable housing distributions, the set of all roads labeled “residential” by OpenStreetMap are extracted, and roads are assigned to Census tracts if both of their endpoints are located within the tract. Going tract by tract, the number of homes in the tract from the Census data is determined and the implied density of housing is calculated through the relative cumulative length of residential roads and number of housing units associated with the specific tract. The process then generates houses at the given density along all of the residential roads. After they have been assembled, households are assigned to these generated houses. The set of generated houses is shown in Figure 6.2, with a closer view of the distribution of houses along roads in Figure 6.3.




Figure 6.2: The set of all generated house locations

6.4 Households

6.4.1 Derive Population-Level Constraints

Both the population-level and the household constraints were derived from United States Census Bureau 2010 Census demographic information (2011). Specifically, the demographic




Figure 6.3: A close-up of generated house locations

traits of interest in this context were age and sex, as well as the household structure in

which the individual is embedded. Thus, when synthesizing the population for any particular Census tract, the fields utilized for the generation of individuals were DP0010021-DP0010038 and DP0010040-DP0010057, which denote the number of males and females (respectively) within 5-year age bands. During the generation of households, the DP0130001-DP0130016 attributes were utilized to generate households of the appropriate types as well as DP0120014 to assign the appropriate proportion of the population to group housing. Other fields such as DP0140001 (the total number of households with individuals under 18 years) and DP015001 (the total number of households with individuals 65 years and over) were used to assess the quality of the household generation results. Based on these constraints, population synthesis proceeds.

6.4.2 The Population of Individuals

The generation of individuals draws on the methodologies introduced in Section 6.1, and particularly on the work of Barthelmy and Toint (2012). Individual members of the population were generated by constructing a distribution of the different documented combinations of age and sex and randomly drawing from this distribution as many times as there were individuals within the population. This generation of individuals occurs at the level of the Census tract, and is repeated for each Census tract within the study area, so that location information is preserved for the subpopulation of individuals as well.

6.4.3 Household Creation

Based on the set of individuals generated during the previous step, the existing individuals are assembled into households. First, a number of individuals are removed from the population in order to account for the population in group housing. Then, a joint distribution of household types is calculated based on the data drawn from the Census Bureau demographic data (2011). Then, for as many households as exist in the real-world, a household type is drawn from the distribution, selecting an appropriate householder and constructing

a set of individuals around him or her. This household construction process is the most notional and least data-driven aspect of the generation, and therefore requires a more detailed description.

There are two steps to the household generation process: the creation of base households and the assignment of unassigned individuals to family households. After the basic requirements of the household population constraints are met in the first step, individuals are added to family households until the population of individuals correctly reflects the underlying demographic patterns and all of the individuals have homes. These steps are elaborated upon below.

In the first step, a Census-derived number of “base” households are generated. Each time a new household is generated, it is assigned a “type”. The Census data specifies a number of “types” of households, specifically family households, husband-wife households, individual-led households, or nonfamily households. The Census does not identify households headed by same-sex married couples, classifying them instead as unmarried partners and therefore either nonfamily or family households, depending on whether other members of the household exist and are related to the couple. Within the specified categories, the Census distinguishes whether the households contain other family members, children under 18, or senior citizens. Thus, when a new household is generated, the type of household is drawn from the distribution of household types and the household itself is populated with the requisite members. These categories are operationalized so that the appropriate sex and age characteristics of the householder are met, their possible spouse is within an appropriate age range, and any children of the householder(s) are within an appropriate age of their parents and each other. The number of children is drawn from a configurable distribution based on the United States Census Bureau’s 2010 Households and Families Brief (2011), but can be no fewer than one. If a household is a non-family household, it consists of either a lone individual or unrelated cohabitants, and is assembled out of a random set of adults. The number of individual adults in a nonfamily living situation is drawn from a configurable distribution, but must be at least two lest it be classified as an individual

living alone. As individuals are added to households, they are removed from the unassigned population. Finally, family households are specifically added to the set of households which can have additional members added onto them in the second step.

In the second step, family households are assigned individuals until they have an adequate number of members. Given these base households, the remaining unassigned population is allocated randomly among the family households. Because family households can contain the parents, siblings, children, stepchildren, grandchildren, nieces or nephews, in-laws, and so forth of the householder, the age range of potential live-in family members is quite broad. After all of the individuals have been assigned to families, the group housing population is added back to the population, with their designated “household” being the group housing unit. In this way, all of the individuals are assigned to meaningful household structures.

6.4.4 Results

The results of the generation of synthetic individuals and households are presented here, with the “individual fit” representing a Pearson's χ^2 test of the fit between the population constraints and the ultimate populations of individuals, with 35 degrees of freedom (36 categories - 1). Thus, for populations where the chi-squared test returns a value of at least 50, the hypothesis that the generated population reflects the true population should be rejected with $\alpha = .05$. All but one of the generated populations meet this criteria; for $\alpha = .10$, four of the 49 Census tracts have populations which deviate from their true distributions (with $\chi^2 = 47$).

In addition to the population fit, Table 6.2 presents a number of statistics that reflect, for example, the percentage error of the true and synthesized proportions of households with children under 18 (House with Minors Err) and members over 65 (House with Seniors Err). It also includes the percentage error between the true and the synthesized average household size (Avg House Size Err) and family group size (Avg Fam Size Err). These numbers are shown in conjunction with the Census tracts overall population (Individuals)

and population in group housing (Group) to give context to the specific kinds of populations. In every case, negative numbers indicate that the synthetic population contained more units than the true population; likewise, positive terms indicate fewer units than expected. The Census tracts are presented in decreasing order of population size.

Table 6.2: A comparison of the generated and real populations and households

Census Tract	Individuals	Group	Households	Avg Fam Size Err	Avg House Size Err	House with Minors Err	House with Seniors Err	χ^2 Indiv Fit
39.06	5985	0	2404	-2.8	0.0	-23.8	-30.2	36.760
29	5892	10	2557	-3.5	-0.2	-21.6	-5.9	39.572
37.09	5822	149	2610	-2.3	0.0	-21.4	-19.6	48.798
67	5649	79	2829	-2.6	0.2	-4.1	-18.4	36.571
37.05	5522	11	2234	-2.9	-0.1	-22.7	-25.9	35.723
28	5476	119	2467	-2.1	0.0	-20.8	-15.3	29.056
77	5407	6	2557	-2.4	-0.2	-26.6	-22.3	27.692
37.02	5140	133	1934	-3.5	0.1	-20.2	-48.7	29.789
13.02	5059	9	2410	-3.4	0.0	-16.1	-14.0	29.817
47.01	4680	0	1789	-3.9	0.2	-23.9	-32.4	33.663
24	4536	53	2432	-1.7	0.3	-13.9	-15.6	27.255
30	4499	56	2368	-0.6	0.0	-13.4	-16.2	25.319
53	4382	96	1636	-1.0	0.1	-22.5	-53.9	26.206
37.06	4359	0	1682	-1.4	-0.1	-32.1	-23.3	34.071

Continued on next page...

Table 6.2: A comparison of the generated and real populations and households

Census Tract	Individuals	Group	Households	Avg Fam Size Err	Avg House Size Err	House with Minors Err	House with Seniors Err	χ^2 Indiv Fit
37.07	4242	0	1588	-3.7	0.0	-15.4	-34.5	33.862
52.01	4216	112	1737	-4.9	0.1	-21.9	-34.4	38.248
19	4173	153	2170	-0.3	-0.2	-10.6	-18.9	37.479
80	3995	0	1873	-2.7	-0.1	-26.2	-28.2	46.284
2.03	3855	0	1708	-4.4	0.1	-27.4	-14.2	39.161
25.02	3761	43	1974	-4.7	0.2	-14.9	-13.2	26.315
34	3638	65	1646	-5.1	0.0	-36.4	-17.5	26.811
78	3586	884	1629	-5.8	-0.1	-71.2	9.6	30.968
39.05	3548	0	1548	-3.3	-0.1	-26.8	-16.1	30.809
3.02	3468	321	1746	-2.4	0.2	-20.4	-0.4	38.807
21.01	3453	219	1797	-5.8	-0.1	-3.4	-36.4	34.331
14	3448	10	1750	0.1	0.0	-6.0	-9.3	37.798
25.01	3345	0	1404	-5.6	-0.1	-16.9	-23.5	21.443
7	3182	0	1417	-3.7	0.2	-26.4	-11.0	27.233

Continued on next page...

Table 6.2: A comparison of the generated and real populations and households

Census Tract	Individuals	Group	Households	Avg Fam Size Err	Avg House Size Err	House with Minors Err	House with Seniors Err	χ^2 Indiv Fit
11.04	3156	0	1460	-5.3	-0.1	-14.3	-29.1	18.811
6	3013	0	1263	-4.2	0.2	-31.1	-19.5	34.473
3.01	2929	68	1383	-6.0	0.1	-19.0	-27.9	35.132
37.08	2914	0	1284	-0.8	0.0	-24.5	-24.3	39.232
27	2872	125	1467	-3.2	0.1	-20.3	-12.7	49.587
22	2681	56	1352	-3.6	-0.2	-16.8	-12.9	30.702
8	2595	18	1323	-3.8	-0.1	-3.2	-16.8	32.417
66	2527	12	1166	-2.6	0.1	-20.2	-12.7	37.291
9	2320	17	1139	-1.1	0.2	-16.3	0.5	31.899
4	2283	2	1073	-2.9	0.1	-21.5	-8.6	50.575
10	2235	68	1060	-10.2	0.1	-20.8	-35.9	28.086
15	2224	144	1050	-3.2	0.1	-14.2	-29.5	43.940
79	2215	1	1128	-17.8	-0.2	-26.5	-20.0	40.313
13.01	2197	79	1113	0.8	-0.2	-9.7	-26.1	36.494

Continued on next page...

Table 6.2: A comparison of the generated and real populations and households

Census Tract	Individuals	Group	Households	Avg Fam Size Err	Avg House Size Err	House with Minors Err	House with Seniors Err	χ^2 Indiv Fit
18	1985	0	1031	2.0	0.2	-18.7	-2.6	47.080
5	1973	0	1012	-2.5	0.0	-12.5	-19.4	34.896
16	1655	2057	1025	-5.7	-0.3	-148.2	47.5	42.785
17	1564	0	826	-3.2	-0.2	-17.2	-10.3	20.961
23	1338	208	854	-0.1	0.2	-6.8	20.3	43.254
11.01	1286	0	705	-1.7	-0.2	3.1	-9.9	28.579
38.02	2	4484	1	0.0	0.0	0.0	0.0	1.032

The synthetic average family size tends to be very slightly larger than the true average family sizes, and there tend to be too many households with minors or seniors. However, the average household size is quite accurate. The greatest error rates seem to occur because a tract is primarily composed of a university, military base, retirement center, or prison. These institutions can significantly skew the populations and household types: the Census Block which contains the Air Force Academy classifies all 4484 young, predominantly male students are part of a single household. However, because the Academy occupies its own Census Block and is so extreme in its structure, it actually has the lowest individual fit error - 1.032. However, even in less extreme cases populations concentrated by demographic characteristics pose a challenge for the algorithm. The generated population of Census Tract 16 contains significantly fewer households with senior citizens and significantly more households with minors than expected: much of its sizable group population lives on the campus of the college, meaning that many of the young students who might otherwise live in households are clustered in school housing. In general, the senior and minor populations seem to be clustered into many fewer households in reality than they are in the generated population, which could perhaps be addressed by explicitly clustering young adults without children into “roommate” structures.

Table 6.3: A sample of the generated households showing individual members, where each line represents a household and each row member represents an individual as Age(Sex)

50 M	45 F	5 F	10 M	
55 M	65 F			
60 M	45 F	15 M	15 F	40 F
50 M	55 F	5 M		
55 M	45 F	5 F	5 F	
40 M	20 F	5 F	5 M	
60 M	50 F	10 F	70 M	
70 F				
40 M	35 F	10 M	10 M	
50 M	50 F	5 F	5 F	

Regardless, the population synthesis algorithm produces reasonable distributions of individuals in terms of age and sex and assembles them into households with recognizable structures. For example, Table 6.3 shows the first ten families generated for Census Tract 79, reflecting the composition relative to individual members’ ages and sexes. Ages are reported as bands, so that an individual marked “50” is in the range of 45-50 years old. These represent a number of households with reasonable distributions of family members: perhaps a man in his late forties married to a woman in her early forties with two children below 10 years old; an older couple with no children in the home; an older husband and younger wife with two daughters and the wife’s sister; an older couple caring for their young grandson; and so forth. While it is difficult to determine whether cross-correlations may exist within the data, the population generated tracks with the available data at the finest grain of detail available, suggesting that it is useable in this context.

6.5 Workplace Creation

For individuals with jobs, the location of their employment significantly impacts the way they travel through their environment and the times at which they do so (Lau, 2009). Further complicating this problem is the fact that Census data typically records residences without recording workplaces, which gives researchers a good sense of nighttime location without a rigorous understanding of daytime location (McPherson & Brown, 2004). To address this need, workplaces are generated and Agents are assigned to them in accordance with their home locations.

6.5.1 Data

As with the problem of generating homes, there is limited information about the specific locations of workplaces and the number of workers associated with them. The Census Bureau provides information about travel to work time and the relative flow of workers between counties, but this information must be disaggregated to be useful for workplace assignment. The OpenStreetMap road network is used to locate notional workplaces, using

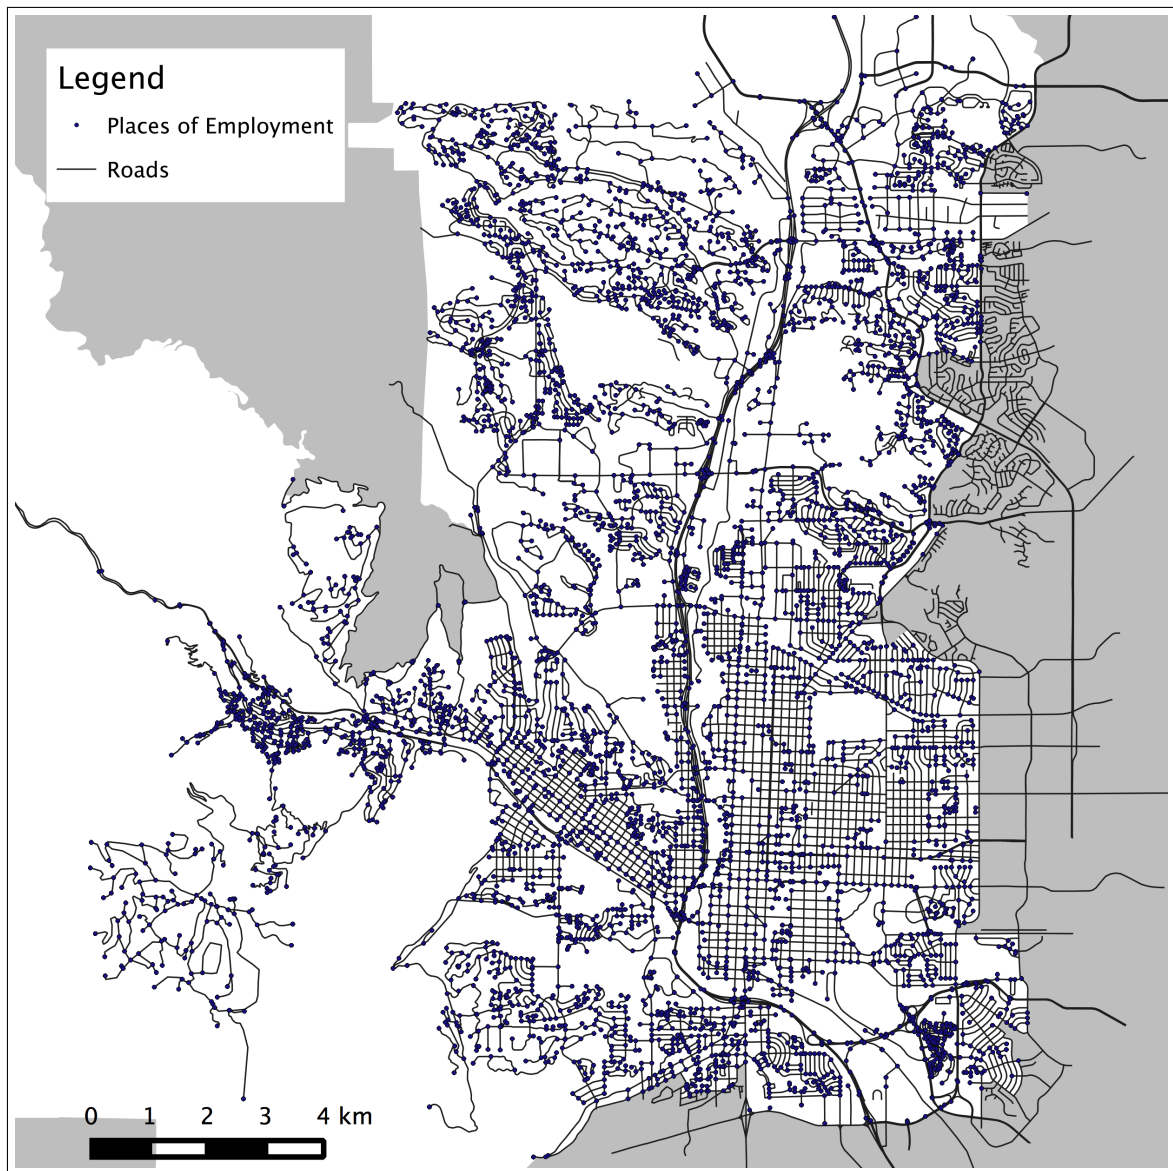


Figure 6.4: Generated workplace locations

nodes in the road network to represent work destination locations.

6.5.2 The Process

Because counties consist of multiple Census tracts joined together, the process begins by aggregating the tract-delimited populations up to the county level. The ratio of workers

traveling from one county to another is calculated, and the percent of the population in that county that is employed is determined. This is a necessary step because many of the counties contain Census tracts that are not populated as a part of the simulation; therefore, assigning all of the jobs would overrepresent the employed population represented in the simulation. Thus, the process assumes that the rate of employment is constant across Census tracts within a given county.

Given this appropriately-scaled set of jobs per county, the counties are processed one by one. For a given county, a set of workers is assembled out of the population of county citizens who are at least 16 years old and not already employed. Then, for each job held by a resident of the county, the county where the resident works is selected from the distribution of commuter flows. A work point is randomly generated from the set of nodes located within the county where the employee works, and a random worker is selected to fill the job. After all of the jobs have been assigned, the remaining individuals are assumed to work informally within the home. The set of generated workplaces is shown in Figure 6.4. In future work, it might be possible to use OpenStreetMap information about points of interest to attempt to generate more precise workplace locations.

6.6 Intimate Social Network Generation

Given the importance of the network of communication between family members and other intimate persons to the spread of information highlighted in Chapter 4, it is necessary to create synthetic intimate social networks. The goal of the intimate social network generation process is to generate a realistic, empirically-based set of undirected ties between individuals. It is important to clarify here that the simulation intends to represent only intimate social ties, or ties that an individual would activate to call for information or to proactively spread information out of concern for the other member of the relationship. These connections and their significance were discussed in Section 4.2. Because this kind of relationship is less concrete and apparent than, say, the network of exchanged phone calls or even a social interaction graph, it is important to ground this in the existing literature

of social interaction.

While the represented intimate social networks do shape communication at a distance in the model (emulating the use of cell phones, email, SMS messages, and so forth to communicate), individuals are still able to communicate with non-intimates and even complete strangers. The intimate social network does not constrain, but rather guide the flow of information.

6.6.1 Data

To reiterate some of the most important the findings discussed in Chapter 4, Wellman and Wortley (1990) suggest that a social network should be low density with smaller-scale clustering behaviors. Kilduff and Tsai (2003) and Watts (2003) concur that a moderate clustering coefficient is reasonable. Watts (2003) and Szabó et al. (2003) agreed in their findings that the average network distance should be low; Bollobás (2003) and Durrett (2007) suggest that the average network distance should be approximately equal to $\frac{\log(n)}{\log(d)}$ (where n represents the number of people involved and d is the average degree of all individuals). This data gives us a sense of how the network overall should look in order to resemble a real-world social network.

As for the personal egocentric networks, Albert and Barabási (2002) cite an approximate power-law distribution of node degree in a network of long-distance phone calls, giving a sense of the range in number of social contacts an individual might have. They put the power law exponent of such a system at approximately 2.1. Wellman and Wortley (1990) determined that an egocentric network might contain about 137 “socially close” contacts. McCarty et al. (2001) found that individuals reported knowing an average of 3.5 immediate family members, 24 other birth family members, 12.3 members of their spouse/significant other, 35.6 coworkers, 4.3 best friends or confidantes, 12.8 neighbors, and 22.6 individuals classified as “just friends”, among various other categories. These groups sum to 115.1 members, and provide a rough sense of the breakdown of social contacts

into various categories. Of these, only a few are likely to be contacted directly in the event of an emergency, so this work suggests that the number of intimate contacts should contain the immediate family members, best friends, and a few other family members and perhaps neighbors. The other close contacts, especially neighbors and coworkers, will only communicate with the individual if they are around one another as a result of their daily habits - an important, but less intimate, set of interactions. One final important aspect of a social network is homophily, or the tendency of individuals to associate with like individuals (McPherson et al., 2001). This results in individuals seeking out the company of similar individuals, a tendency which the intimate social network generation process should reasonably attempt to capture.

6.6.2 Process

The intimate social network generation assembles the ego network of a single random individual at a time, so that over the course of the process a comprehensive, higher-level network emerges. First, the individual agent draws the number of connections it has - its social degree - based on a power law distribution. The individual then assembles a list of individuals to whom it may be linked, first considering any friends of friends and then selecting random individuals. The assembled list of potential contacts is sorted by the social distance between agents, where the social distance is a weighted score based on the similarities between sex, age, and home location of the two agents. More similar agents are added as friends first, until the desired number of friends has been added. Friends of friends receive a social distance “discount” to promote the formation of clusters of similar individuals. The parameters of the power law, the social degree, and the importance of social distance are drawn from the texts discussed in greater detail in Chapter 4. These relationships are not directed, so that if an individual has acquired friendships through the egocentric network formation of others, it builds upon this existing structure in constructing its own egocentric network.

A number of assumptions exist here. For example, the social degree of the individual is

Table 6.4: Parameters of the generated intimate social network

Parameter	Value	Target	Target Citation
Num Nodes	170282	-	-
Num Edges	1356657	-	-
Avg Degree	15.9	10	Eubank et al., 2004
Network Diameter	9	6	Eubank et al., 2004
Modularity	0.616	-	-
Avg Clustering Coeff	0.208	0.480	Eubank et al., 2004
Avg Path Length	6.9	6	Albert & Barabási, 2002

assumed not to vary with age, sex, and so forth. Further, by assigning a weighting based on a notional social distance, the process assumes that all individuals perceive social distance with the same specific weighting; a 10 year old boy and a 14 year old girl perceive the same social distance between themselves as do an 80 year old man and woman. All individuals are assumed to be socially engaged with others in the study area, which obfuscates the fact that the variation in social degree can account for differences not only in personality but also in familiarity and experience with the given study area. So, for example, an individual who has moved to Colorado Springs for work two days before the simulation begins may have zero intimate contacts in the study area, but in fact be hugely sociable and have 50 very close contacts with individuals outside the modeled network. Teasing out the variation at such a level is beyond the scope of this work, but bears mentioning.

The intimate social network generated as a result of this process is shown in Figure 6.5, and the structure itself was analyzed using a variety of statistics. The qualities of the network and the way the generated model compares to real-world models are explored in the next section.

6.6.3 Results

Table 6.4 gives the network-level properties of the generated intimate social network. From this network, a number of reasonable findings emerge. The average degree of 15.9 for intimate contacts is reasonable relative to the number of social ties presented in the previous




Figure 6.5: The large and very dense generated intimate social network. The color of a node indicates the community to which the node has been assigned, each community being represented by a different colour

sections and Table 4.5, and the degree distribution is shown in Figure 6.6 to be power-law distributed. Both the network diameter of 9 and the average path length align with Milgram's (1967) small world experiment. The clustering coefficient of 0.208 is indeed moderate, as suggested by Watts (2003). Given these parameters, it seems that the generated intimate social network tracks well with the literature regarding real-world social networks.

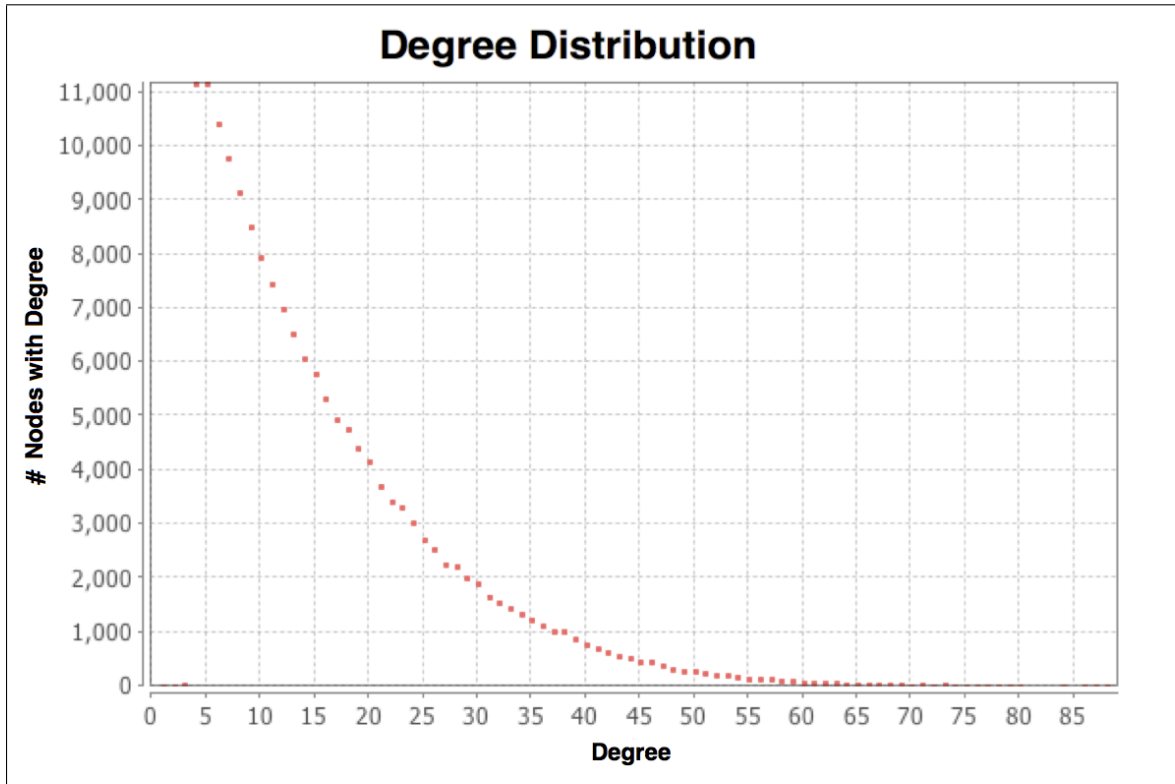


Figure 6.6: The degree distribution of the generated intimate social network.

6.7 Social Media Network Generation

As explored in Chapters 3 and 4, social media networks represent an important new medium of communication, one which is frequently utilized in crisis situations to distribute information (see Sutton et al., 2008; Hughes & Palen, 2009; Sarcevic et al., 2012). To that end, it is important to capture the way information spreads masspersonally, and therefore the way individuals are linked to one another through social media. This section draws upon the Twitter model structure described in Section 4.4.6, allowing individuals to share information based on their social media usage habits.

6.7.1 Data

The data used to support the social media network generation is drawn from two 2012 Pew Data surveys here, specifically the Demographics of Social Media Users Report (Duggan & Brenner, 2013) and the Digital Differences Report (Zickuhr & Smith, 2012). The former captures information about the specific demographic characteristics of users, while the latter includes information about internet accessibility broken down the by age, sex, race, educational level, and so forth. The Demographics of Social Media Users Report breaks down social media usage and internet usage, and additionally presents data that is not included in the model but which is included for context and to inform future directions of work. For example, while household income is not a major predictor of social media usage among internet users, it plays an enormous role in whether an individual is an internet user in the first place. Thus, incorporating these attributes into the simulation could have a significant impact on understanding how information flows through the population. Such attributes are beyond the scope of this thesis, but reflect the potential for future expansion.

6.7.2 Process

Firstly, the aggregate statistics about the relative distributions of both social media and internet users by age and sex are read into the simulation. These distributions, summarized in Table 6.5, are taken from the breakdown presented in Duggan and Brenner (2013), with

Table 6.5: The percentage of different demographic groups which utilize Twitter. (Source: Duggan & Brenner, 2013)

Twitter Usage - (% of Pop.)	
Sex	
Men (n=846)	17
Women (n=956)	15
Age	
18-29 (n=318)	27
30-49 (n=532)	16
50-64 (n=551)	10
65+ (n=368)	2
Household income	
Less than \$30,000/yr (n=409)	16
\$30,000-\$49,999 (n=330)	16
\$50,000-\$74,999 (n=330)	14
\$75,000+ (n=504)	17
Urbanity	
Urban (n=561)	20
Suburban (n=905)	14
Rural (n=336)	12

simple conditional probabilities depending on both age and sex attributes being used in a rough version of IPF to establish how many individuals should exist within each age/sex category. Once the counts match the described distribution, the ratios of participation within each category is calculated. Each individual agent is assigned to be a social media user or non-user based on the probability of their specific age/sex combination utilizing Twitter. This population of users is clustered using the same algorithm that was applied to intimate social networks, linking individuals to the friends of friends and sampling random individuals, adding the most “similar” agents to ones social media network first. The generated social media network is shown in Figure 6.7.

A number of assumptions are made in utilizing this data, which are important to acknowledge. For example, the distribution of users is assumed to follow the nation-wide distribution of users. That is, the process assumes that Colorado Springs demonstrates the

same levels of Twitter adoption by sex and age as does the United States overall, regardless of the local variations in Twitter usage by geographic region. This is assumed despite the fact that on the average urban, suburban, and rural populations demonstrate different degrees of Twitter adoption (Duggan & Brenner, 2013) because the data provided only described the aggregate levels of adoption, failing to break down the location-based data by age and sex. This would require making the assumption that that rural and urban 35 year olds utilize Twitter at corresponding rates, for example. Social media usage data is not available at the level of individual areas, so that it is difficult to identify the kind of variation at this level that might influence usage rates. Thus, local patterns of adoption have trends for which adequate data does not seem to exist; the process must rely on the higher-level statistics. This discrepancy in existing data sources makes this line of research a particularly interesting one, and a rich area for potential future work.

Further, the process assumes a uniform distribution of Twitter usage within age groups, so that when the data reports that 17% of 25-35 year olds utilize Twitter the process that Twitter is being used by 17% of 26 year olds, 17% of 27 year olds, 17% of 38 year olds, and so forth. There is also the assumption that sex and age are independent as determinants of Twitter usage. This ignores the possibility that there might exist a massive population of young women who comprise 80% of female Twitter users while 80% of the 65+ Twitter users are males.

Finally, the process includes internet access statistics into the assignment of individuals to be social media users. The Pew Data regarding social media usage applies only to those who already use the internet, so it is important not to treat all adults as potential users. However, internet usage tracks closely with education and income information, (Duggan & Brenner, 2013) which is not currently included as characteristics of the generated population. These trends can therefore not be represented at the lowest level of granularity, which might be a desirable adjustment in future work.




Figure 6.7: The generated social media network. The color of a node indicates the community to which the node has been assigned, each community being represented by a different colour

6.7.3 Results

The generated social media network shows a number of promising and reasonable properties. As summarized in Table 6.6, the network being generated does not attempt to create celebrities or other highly-connected individuals, the network reflects a degree distribution

Table 6.6: Parameters of the generated social media network

Parameter	Value	Target	Target Citation
Num Nodes	18458	-	-
Num Edges	157874	-	-
Avg Degree	17.1	18.86	Java et al., 2007
Network Diameter	7	6	Kwak et al., 2010
Modularity	0.487	-	-
Avg Clustering Coeff	0.066	0.106	Java et al., 2007
Avg Path Length	3.895	4.12	Kwak et al., 2010

for the way individuals follow one another that matches Starbird and Palen’s (2010) reports of the network structures apparent in local communities. Again, Table 4.6 serves as a point of comparison for the synthetic structure, and tracks well with the social media network properties introduced in Section 4.4.6. A local news organization is included as an input to the social media network during the course of the simulation, ensuring that these hubs of information are represented in the flow of information as both Java et al. (2007) and Hughes and Palen (2009) suggest they should be. This hub and authority structure, combined with the social media network generation, ensures that there are individuals whose presence online is extremely limited as well as heavier local users, as shown in Figure 6.8. The network diameter is low at 7, which is again in line with Milgram’s (1967) suggested target and especially appropriate for a social media network localized to a particular community. The clustering coefficient is low at 0.066, which reflects the hierarchical information propagation structures described in Chapter 4 and particularly good compared to the issues other algorithms have had with this parameter (see Stonedahl et al., 2010). As with the generated intimate social network, the given values track well with the literature and suggest that this network will allow and constrict information along its appropriate channels of communication.




Figure 6.8: The degree distribution of the generated social media network

6.8 Agent Synthesis

Once individuals and households have been generated and connected to one another in appropriate structures, the populations are reduced into meaningful actor units. For each employed person in a household, an Agent is created. All other household members are grouped into a unit and, in as much as the model evaluates the evacuation of homes, assumed to act together, receiving and sharing information among themselves (see McConnell et al., 2010). This partitioning of the population is important, as modeling employed individuals as independent units is necessary in order to capture the commuting behavior of different family members. In future work, it would be interesting to simulate the merging and partitioning of the household unit, but such efforts are beyond the scope of this work. Having distributed the population of individuals among appropriate actor units, the social

ties that exist among these new Agents are calculated, potentially aggregating up based on the social networks of the individual household members. The functional units are then passed to the simulation itself, and it is these created Agents which support the simulation itself.

6.9 Summary

This chapter has outline the steps necessary to create a synthetic population for the model presented in this work. From these combined processes emerges a demographically realistic population with reasonable social clustering and social media usage patterns, as well as data-driven homes and workplaces. By combining information about various aspects of community and household structures with research into human social networks, it is possible to generate a synthetic population which can serve as a basis for testing counterfactual worlds. This technique is not strictly necessary in parts of the world like Colorado Springs, for which precise, validated, and publicly available data exist (see Wheaton et al., 2009), but hopefully it can be utilized elsewhere. Arguably, it is for that reason that it is helpful to test the population synthesis technique on this well-known population: by comparing how the real population and the synthetic population match up, it is possible to gauge the effectiveness of the technique and draw conclusions about how it may perform in situations where less rigorously collected data is available. The aim of this chapter was to create a simple but representative agent population with appropriate and associated social networks, an effort which is to the author's knowledge without precedent and one of the major contributions of this work. Chapter 7 utilizes the information generated here in order to create a rich, interconnected world of individuals with meaningful patterns of communication and interaction.

Chapter 7: Modeling The 2012 Colorado Wildfires

7.1 Introduction

During June and July of 2012, record heat and droughts across the western United States gave rise to some of the largest wildfires in Colorado state history (Oldham et al., 2012). The Waldo Canyon fire began on June 23 in the mountains west of Colorado Springs and eventually prompted the evacuation of over 34,500 citizens (Udell, 2012). Two individuals died in the fire, and the crisis ultimately cost the area over \$125 million (Minschew & Schneider). It burned 18,247 acres over 18 days, and was at the time the most destructive wildfire in Colorado history (City of Colorado Springs, 2013). Figure 7.1 shows the Colorado Springs area relative to the wildfire progression over time.

Given the fast-moving nature of the fire, the situation changed rapidly. Evacuation orders were broadcast via an Emergency Notification System (ENS), social media including Twitter, and door-to-door notification from response personnel in the evacuation areas (City of Colorado Springs, 2013). News reports focused on the extensive backup of traffic as individuals fled the evacuation zones or tried to reach family members, homes, or pets in the affected areas (Udell, 2012). Extensive conversations about the situation in Colorado went on via social media, with individuals posting pictures of the view of the fire from their locations to platforms like Flickr. The flow of people and of information - and the dependencies between these - made the situation complex, and the need of response personnel for timely, accurate, and actionable information and for planning makes the situation a valuable case study. Traditional sources of information could not support the information needs of responders, as the situation changed rapidly and individuals made choices based on the information they were themselves consuming. Thus, an approach which synthesizes information from a variety of different sources and uses them to project the development

of the crisis provides an understanding of the situation that would otherwise be totally inaccessible to responders.

The model of the Waldo Canyon wildfire evacuation presented in this chapter was developed in order to explore the feasibility of projecting citizen movement in response to the fire. The model attempts to capture the movement of the wildfire itself, the communication among individuals, the movements of individuals both in the absence of knowledge of a threat and in response to it, and the stress level to which they were subject. If the simulation accurately captures the dynamics of the system, based on the processes presented in Chapter 2, it should produce reasonable approximations of the road usage patterns, especially blockages and traffic, as well as patterns of evacuation and stress. To that end, the results in Chapter 8 are presented in terms of the heatmap - that is, road usage - and speed of movement over the course of the simulation. The emotional valence of the terminology utilized on Twitter in conjunction with wildfire-specific hashtags, presented in Chapter 5, is compared with the simulation-generated rates of stress, using social media to attempt to capture the intangible aspects of the evacuation experience.

Section 7.2 describes the data utilized to support and validate the model environment, presenting the sources as well as their cleaning and manipulation into a usable format. Next, Section 7.3 presents a specification of the model itself, with a complete rundown of the methodologies employed and the way these processes are linked together. The section also explains how each of the data sources created in Chapters 5 and 6 are utilized in the model. The final section provides a transition between the model specification presented in this chapter and the exploration of a variety of model results presented in Chapter 8. The model and code described in this chapter is available online at www.css.gmu.edu/swise/thesis.

7.2 Incorporating Data

Given the number of different phenomena described by the model, the range of data used to support the analysis is relatively broad. Being a spatially explicit model, GIS data was required; being a social model, information about demographics, social interaction,

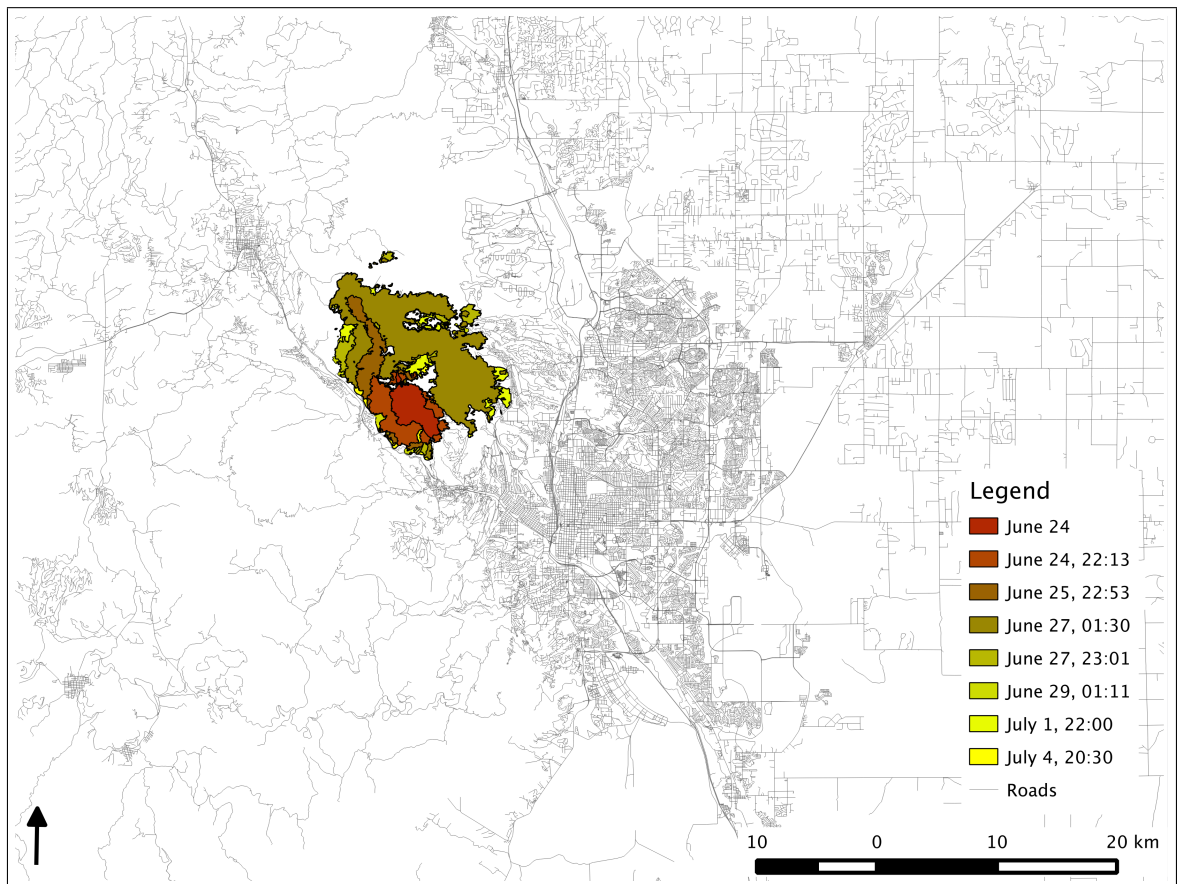


Figure 7.1: Map of study area with overlaid fire progression

and sentiment were utilized. The data are diverse not only in their type, but also with regard to their sources. Official government records of administrative boundaries, maps of the spread of the wildfire over time from news organizations, volunteered geographic information pertaining to roads, and harvested social media information are synthesized with academic studies of the properties of social networks in order to populate the model with the necessary components. This section will introduce other additional data sources and describe how both these and the information collected, processed, and synthesized in the previous chapters are incorporated into the model itself.

In some cases, the simplest path to gathering data was not taken - while demographic information about communities in the United States is largely available, a synthetic population was generated as a proof of concept for areas where such information is less available or significantly out of date. For example, Pakistan has not conducted a census since the 1980s when it was struck by floods in 2010 (Khan & Salman, 2012). The lack of availability of spatial or demographic information in less developed parts of the world makes reliance on this kind of well-cleaned, well-structured data a hazardous proposition. Responders may find themselves targeting population centers that disappeared years ago and ignoring vast refugee camps, or focusing all of their attention on the well-mapped parts of the country, which may correspond to the most developed and therefore least vulnerable populations. Further, even in situations where such information is available, spatial and demographic data can rapidly go out of date, making the ability to quickly generate reasonable approximations of a local population an important point (Crooks & Wise, 2013). The realization of this is described in detail in Chapter 6. As a further benefit, the workplaces for individuals in Chapter 6 also incorporate the movement of individuals during the day into the model, which addresses the frequently overlooked problem of tracking residences rather than daily activity spaces. A terrorist attack on the transportation network might look relatively unthreatening if citizens are assumed to stay in their residential suburbs all day, but by including daily patterns of movement between work and home, the problem becomes apparent.

In other cases, no simple method existed to measure the phenomenon at work. Measures of sentiment at a population level are rare, although marketers and researchers alike are studying the problem even now (e.g. Dodds et al., 2011). Some of these difficulties are elaborated upon further in Chapter 5. Even when explicit structures like social media networks exist, the way individuals engage with their networks varies. This varying engagement was discussed in Chapter 4, and essentially suggests that users may vary in their frequency of both consumption and production of media. Yet despite these difficulties, the features in question impact the choices made by individuals, and it is important to attempt to simulate them. The synthesis of different kinds of information to attempt to support these needs is decidedly an emerging field, as Chapters 4 and 5 explored. Given the extensive treatment of the demographic and social data sources utilized in the creation of agents in Chapter 6, those data are not discussed further here.

In addition to the spatial qualities which influence human movement on a day-to-day basis, the geography of an environment in a crisis situation influences the development of the crisis the resulting human movement patterns. The importance of the spatial element of an evacuation model was made clear above, but it is worthwhile to reiterate that the goal of including GIS data is to explain the spatial relationships and interactions among elements. Tobler’s first law of geography, “Everything is related to everything else, but near things are more related than distant things”, is particularly relevant here (1970). Because the environment shapes the way agents move, observe, and interact with one another, it is important to have meaningful and representative information about the structures which influence these behaviors (Crooks & Heppenstall, 2012). As discussed in Chapter 3, old data, bad data, and variations in quality within different mapped regions are all potential pitfalls; it is with these concerns in mind that the following data sources were synthesized.

Wildfire location data - The locations of the Waldo Canyon Wildfire over the course of its development were mapped by the Denver Post, a local news organization. The data were transformed from their original KML format into shapefiles for inclusion in this work, and the different shapefiles are included in the model with their temporal information,

rounding the time of record to the nearest hour (Denver Post, 2012b). This data was used for comparison against the generated wildfire boundaries to give a sense of the effectiveness of the model.

Road network data - The locations of roads were taken from the CloudMade extract of OpenStreetMap (CloudMade, 2012). The road network was cleaned in order to reduce the number of roads to the Colorado Springs and Woodland Park area, maintaining the highways which led out of the area to ensure that evacuation traffic could be accurately captured. All footpaths, service roads, and other non-car-accessible roads were removed from the dataset, and during the course of the simulation itself, road nodes within the resolution distance (5 meters) of one another are merged together. The goal of this cleaning was to only include accessible, existent roads in the path-planning of the Agents, which were assumed to move through the environment by car. The roads themselves are partitioned where they intersect with other road nodes, to further ease the process of Agent path-planning. Finally, during the initial loading process, all components which are not connected to the main road network are removed from the simulation. This creates a network which reflects the possible movement patterns of vehicles rather than individuals, as befits the evacuation scenario. More importantly, this code is extendable and able to work in conjunction with data sources that have been less carefully curated: for example, intersections which have been imperfectly joined will be merged, making useable data that would otherwise require extensive effort to perfect before it could be utilized. This extensibility is crucially important in situations where data has been rapidly gathered and recorded, as described in Chapter 3.

Physical environment data - A number of layers were utilized in concert to construct the physical environment relevant to the wildfire itself. From the United States Geological Survey, information about the elevation (National Elevation Dataset: see Homer et al., 2007), land cover type (MODIS's normalized difference vegetation index datasets: see Huete et al., 1999), and permeability of the environment was downloaded from the USGS's EarthExplorer website and clipped to the relevant area. These raster data sources


were reprojected to the relevant coordinate system and sampled to a resolution of approximately $30m^2$ for use with the wildfire submodel, that being the lowest resolution raster data available. Figure 7.2 shows each of the layers used to construct the physical environment.

Evacuation data - The location of the evacuation areas that were imposed on Colorado Springs by the City of Colorado Springs. This information was taken from the Colorado Springs Waldo Canyon Fire After Action Review (2013) and manually digitized. The layer, shown in Figure 7.3, was utilized to explore how the evacuation orders influenced individual behavior, prompting Agents whose homes were within the progressively evacuated areas to leave.


Processing and cleaning this data posed a number of challenges, briefly elaborated upon here as an aid to others. In all of the data handling associated with this project, ensuring consistency between formats and the projection of the geographic data was difficult. The data originally posted by the Denver Post was structured as KML data formatted for GoogleEarth, so that it was difficult to read and reproject the data using a program like QuantumGIS; it was ultimately necessary to extract the bundled KML files from the document and then extract the polygon info from these. In addition to formatting problems there were issues of sheer data size: the number of roads in Colorado is computationally challenging, with 283,118 unique geometries before road partitioning was considered. In the end, it was necessary to restrict the roads in the model to the Colorado Springs area, including only the major highways running out of the city.

7.3 Model Overview


Figure 7.4 gives an overview of the data sources and processes that work together to create the simulation described in this chapter. In the remainder of this chapter, the model is presented after the style of Grimm et al. (2006) in order to aid comparison with other agent-based models. The rest of Section 7.3 will present the purpose of the model, a brief review of the state variables and scales involved in the simulation, and the processes and scheduling to which they are subject. Next, Section 7.4 will describe how stochasticity,



A



B



C

Figure 7.2: Geographic data representing A) elevation, B) land use type, and C) permeability of surfaces

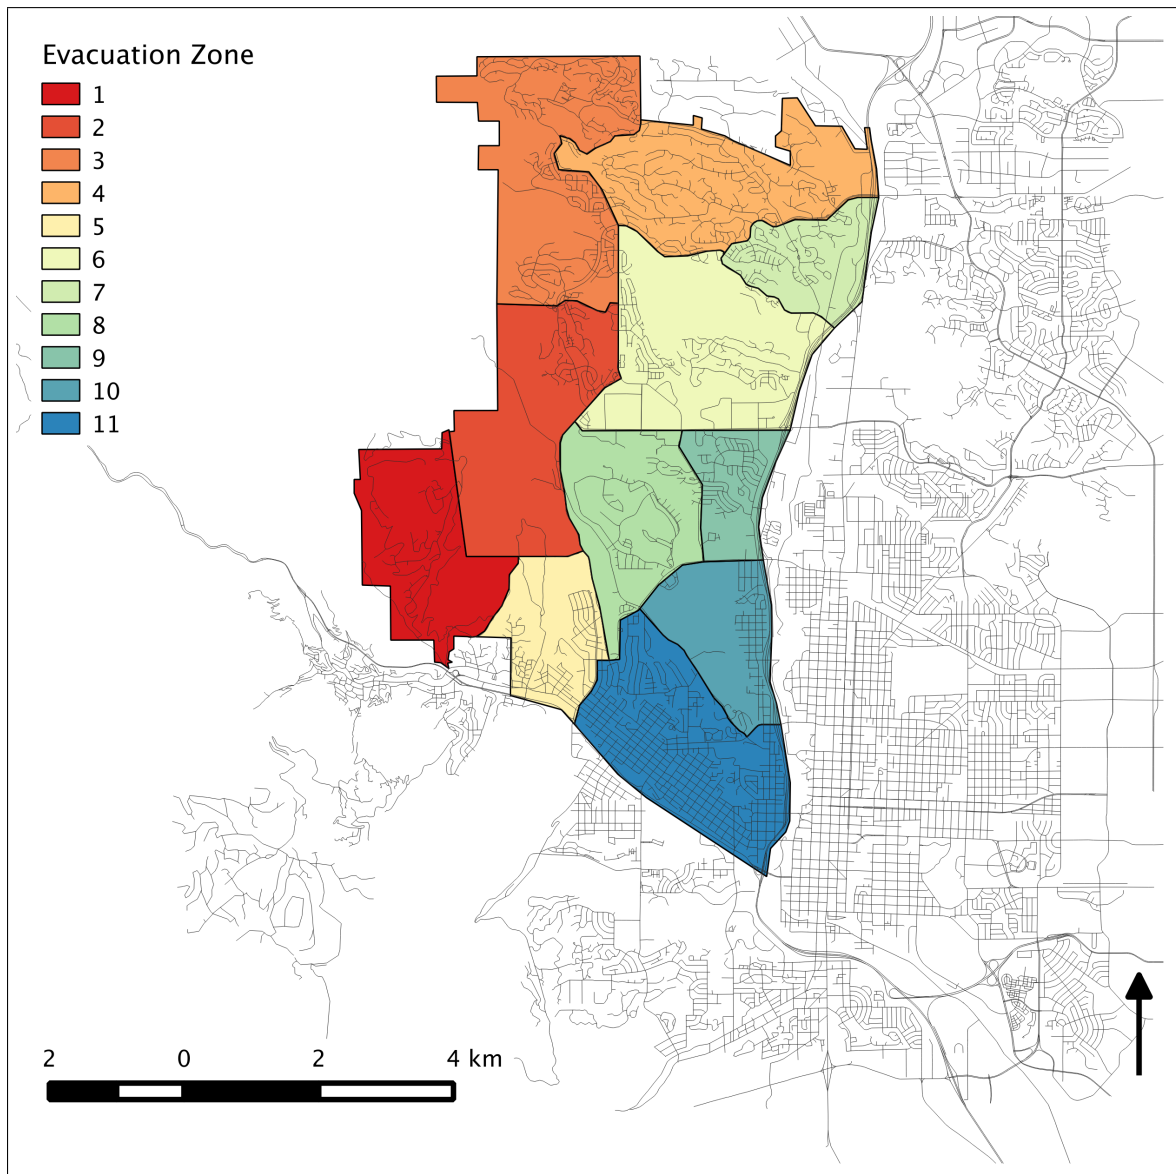


Figure 7.3: Map of study area with overlaid evacuation zones

emergence, and other complexity-related ideas influence the simulation. Finally, Section 7.5 will present the specifics of the implementation of the entities and processes introduced in the overview.

7.3.1 Purpose

The purpose of the model is to simulate the evacuation of individuals from Colorado Springs in light of information, stress, and the physical reality of the environment. The resulting patterns of movement can give insight into how the spread of information or the state of the transportation network could be influenced, ideally in pursuit of a faster, safer evacuation. The model showcases how social media, VGI, synthesized populations, and other data sources can be incorporated into a simulation in order to better inform and guide the processes being studied. Figure 7.4 specifically shows how the different processes work together, and the data sources upon which they draw, reflecting the synthesis of diverse data sources into the physical environment and agents that interact to produce the simulation itself.

7.3.2 Entities, state variables, and scales

The model operates on three different entities: the individual, the population, and the environment. While individuals make decisions, move, observe, and communicate among themselves, the ultimate unit of observation is the population and the success of the population overall at avoiding the wildfire. The environment, too, is subject to its own dynamics and merits observation: the wildfire’s progress across the environment is an important emergent finding. Figure 7.5 shows the entities and how they interact with one another.

With regard to scales, the simulation is updated on a timestep of 5 minutes per tick, following the timescales for action introduced in Chapter 2 which seek to capture the pace of higher-level human decision-making. While the simulation of Agents occurs over a 20 km by 20 km area, the wildfire is simulated on an approximately 90 by 80 km area at a 30m cell resolution. This larger area for the wildfire reflects its larger scale, and it included in

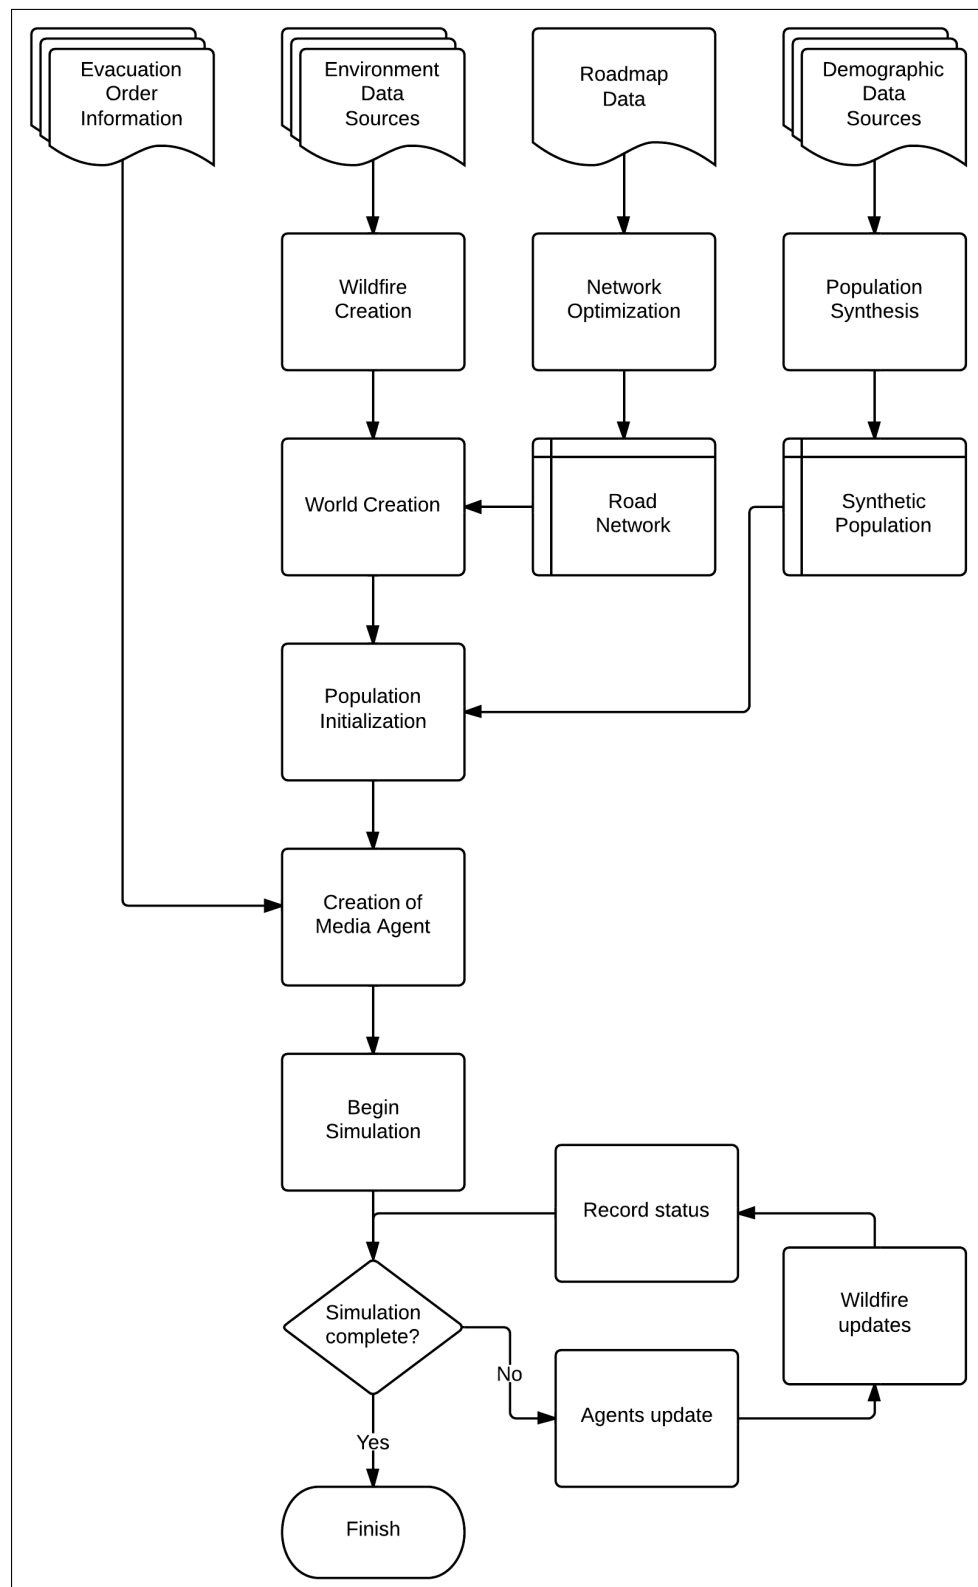


Figure 7.4: Overview of the model structure, detailing how data is used within subprocesses which underlie the simulation

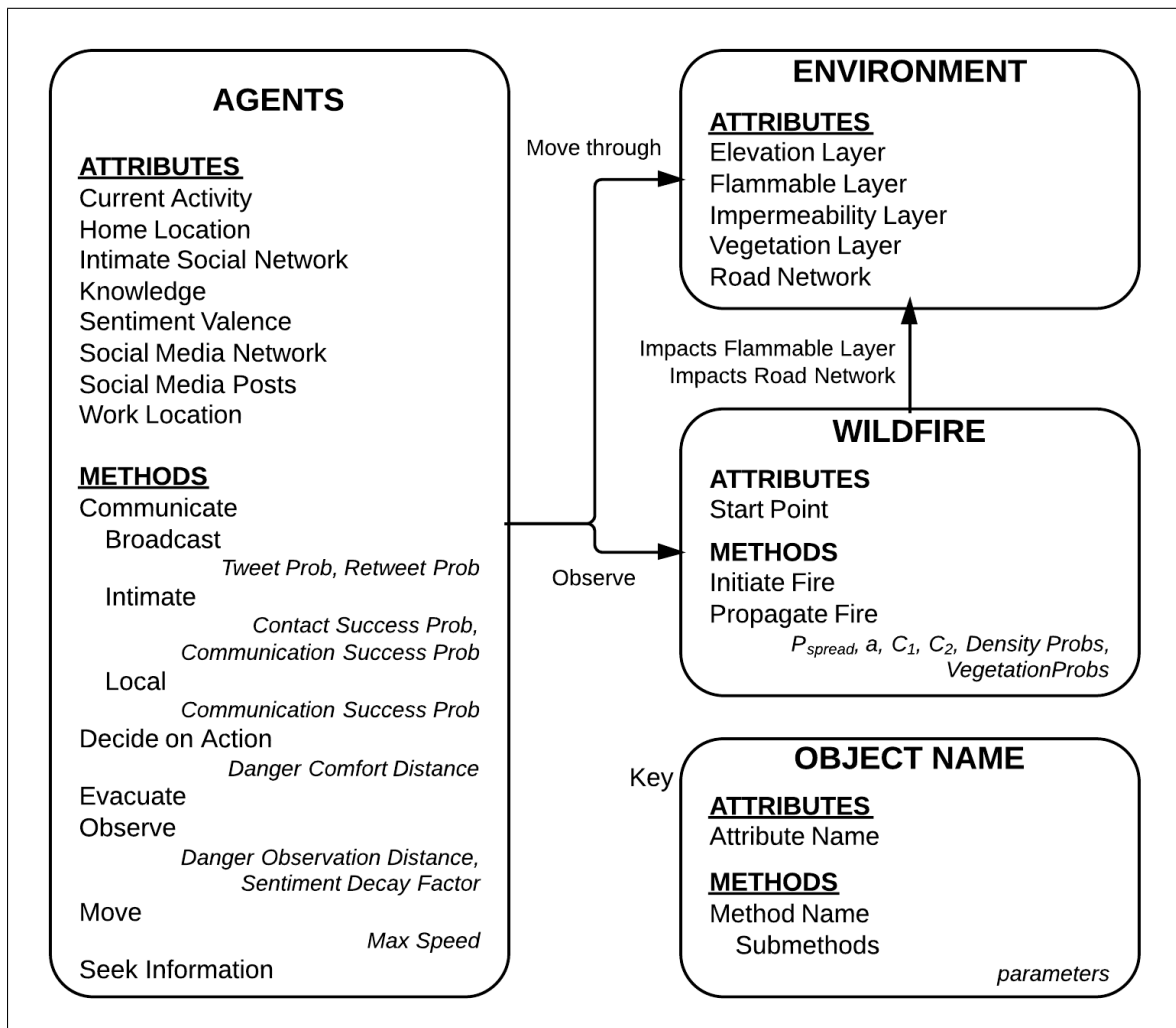


Figure 7.5: Overview of entities, their relationship to on another, and the parameters which effect them

order to prevent boundary effects (see Parker & Meretsky, 2004 for a discussion of boundary effects in ABM). The resolution of the simulation is set at $5m^2$, so that Agents within 5m of one another are considered to be colocated and roads whose ends are within 5m of one another are understood to be connected.

The low-level entities active in the model are Agents and Wildfires. There is additionally a trivially simple Media agent which injects information into the simulation in order to emulate the real-world evacuation announcements that went out over the Emergency Notification System; however, its behavior is limited to pushing information and it should more properly be considered a part of the model framework than an entity with agency.

Table 7.1: Default parameters for Agents.

Parameters	
Contact Success Probability	50%
Communication Success Probability	80%
Tweet Generation Probability	10%
Retweet Probability	10%
Danger Comfort Distance	10000 m
Danger Observation Distance	1000 m
Decay Parameter	0.5
Speed	2000 (m/5min)

Agents

Within the simulation, agents represent households or parts of households making decisions about whether and how to evacuate. They are the operationalization of the members of the synthetic population created in Chapter 6, modeled at the household level to capture the group behaviors described in Chapter 2 as well as the transportation reality of most individuals evacuating via car (Udell, 2012). Agents are embedded with a social network and endowed with an emotional stress level, individual knowledge of the transportation network described in Section 7.2, and individual knowledge of the location of threats, all

of which are updated as they receive and process new information. An Agent also has a home location and potentially a work location. Agents communicate with one another and move around as a function of their knowledge and stress levels. The range of activities and behaviors of Agents are described in greater detail below, but are based upon the literature reviewed in Chapter 2. The default parameters employed in this work are a result of both the findings of Chapters 2 through 4 and the experimentation detailed in Chapter 8: they are shown in Table 7.1.

Wildfires

Within the simulation, wildfires move over the environment burning flammable areas. Functionally, a Wildfire is a cellular automata which underlies the rest of the simulation: Wildfires spread according to the direction of the wind, the type and density of vegetation of both the burning location and its neighbor, and the difference in elevation between the two location. The parameters which define Wildfire development are shown in Table 7.2. As a Wildfire grows, it may overtake segments of roads, rendering them unserviceable and potentially stranding Agents on newly unconnected parts of the fragmenting transportation network. While the Wildfire itself depends only on the elevation, vegetation, and impermeability data described above, in future work it would be possible for fire-fighting Agents to modify the impermeability or vegetation of land and thereby impact the spread of the fire - a behavior which does not appear in other published models at the time of writing.

7.3.3 Process overview and scheduling

The scheduling process has been carefully constructed in order to ease the computational burden of simulating thousands of Agents over a broad extent of space, as well as across social networks and complex geometries. In the results presented in Chapter 8, the model is run for half of a week to coincide with the period of mandated evacuation, although it could easily be shortened or extended according to desire. All of these interactions impact the model, and it is important to consider how the scheduling of these processes impacts

Table 7.2: Parameters utilized by the Wildfire object.

Parameters	
P_{spread}	0.58
a	0.078
C_1	0.045
C_2	0.131
Probability given Density	
Density	$P_{density}$
Sparse	-0.4
Normal	0.0
Dense	0.3
Probability given Vegetation	
Type	$P_{vegetation}$
Agricultural	-0.3
Thickets	0.0
Pine	0.4

the development of patterns overall. Some of these scheduling patterns are simple, such as the scheduling of Wildfires at a regular interval. Agents, however, have complex activation processes which depend on their current activity and knowledge. These processes will be discussed extensively below.

To begin with the simplest case, Wildfires are updated at the beginning of the step on an hourly interval. That is to say, while the simulation captures the behavior of individuals at the resolution of 5 minutes, the wildfire is updated every hour, following the work of Alexandridis et al. (2008), who present their resulting fire-development patterns in one hour intervals. The justification for this selection of wildfire model is discussed in Section 2.6. When activated, the Wildfire goes through each of the currently burning locations and probabilistically ignites the neighboring locations based on their status as of the last simulation step. Roads covered by the wildfire are marked as “closed” and inaccessible from the rest of the network at this point as well.

The activation of Agents is more complex. Agents are scheduled to run their full behavioral module on different timesteps depending on their status; an Agent embarking on its workday will schedule its next mandatory decision check-in for the end of the work day, and

an Agent about to go to sleep will schedule its next mandatory decision check-in for when it wakes up. Traveling or evacuating Agents consider their full range of options every tick until they reach their destination, but sleeping or working Agents usually only make choices at transition points (e.g. leaving for or from work, waking up, etc). However, Agents check their surroundings on a more regular timestep than they make decisions, and can be spurred into action either by their observations or by successfully being contacted by another Agent bearing new information, reflecting the crucial role of information discussed in Chapter 2. This minimization of unnecessary activation substantially speeds up the simulation.

Upon activation of its decision-making process (discussed further in Section 7.5.3), an individual Agent considers its position, may choose based on its set of knowledge whether to attempt to communicate with others, makes choices based on its knowledge of the presence of danger, decides whether/where to move, and then attempts to navigate through the environment, an effort in which it may not be completely successful. As mentioned above, this process is run whenever Agents are transitioning between activities, or when Agents are out in the world navigating. However, an Agent can activate its decision-making process no more than once per tick.

More frequently called is the Agent’s observation behavior, which happens on an hourly basis and comes before the media consumption behavior. All of the Agents are scheduled to complete their activations after the Wildfire has already been updated, so that Agents are working with the current state of the world. Individual Agents are activated in random order to prevent any one Agent from consistently having the advantage of moving first. The specifics of this activation scheme will be elaborated upon in the specific sections that deal with the decision-making (Section 7.5.3), observing (Section 7.5.3), and media participation behaviors (Section 7.5.3).

7.4 Design Concepts

A number of complexity-oriented design concepts feature in the model as specified. **Emergence** is perhaps the most crucial and obvious phenomenon - the simulation relates the

emergence of a huge conflagration to the emergence of an evacuation effort, not to mention emergent information cascades and regular commuter traffic. **Adaptation** is another central concept: Agents plan and replan their fitness-maximizing escape routes as they learn about the location of threats, which roads are inaccessible, and which areas are being evacuated. This dynamic updating of goals also plays into the measurement of **fitness** Agents utilize when they make choices about the fitness of various targets and paths. Agents do not **predict** where the fire will move next or when an evacuation order will be announced; although such behavior would be an interesting addition, it is beyond the scope of this research. Therefore, they do not take action based on projected future events, although they do use **sensing** to observe their environment, including wildfires. **Interaction** between Agents occurs both in terms of communication patterns and via traffic dynamics, with individuals blocking each other in terms of movement. **Stochasticity** plays a role as well, both in the Wildfire (as in the probability any given tile has to burn) and in the Agents (even for a given synthetically derived population, the probability of success in making contact, observing the Wildfire, and so forth). However, for the purposes of the simulation runs presented here, the Wildfire maintains an independent random number generator which is seeded with a precise digit in order to ensure that the Wildfire develops consistently across runs and that the dependent variable being tested is the Agent behavior. In a sense, **collectives** are built into the very structure of the Agents, as Agents themselves may represent collections of individuals. Finally, **observation** of the simulation is accomplished through recording a heatmap of agent locations when traveling, tracking the speeds of moving Agents at a regular interval, tracking the stress of Agents over time, and finally recording the number of deaths and evacuations, respectively.

7.5 Details

This section will elaborate upon the details that were omitted in the overview of scheduling, agent behaviors, interactions, and so forth above.

7.5.1 Initialization

The simulation begins with reading in layers of GIS data: the underlying road networks and evacuation zones create the environment in which the Agents move, and landcover, elevation data, and permeability support the Wildfire. The Wildfire Creation step shown in Figure 7.4 consists of this synthesis of data. During the Network Optimization stage, the road network is parsed into a format usable by Agents and cleaned to remove unconnected components and to merge road points within 5m of one another into single units. These two different processes create the world, shown in the World Creation step. Next, the synthetic population generated in Chapter 6 is then read in from a file. It is important to remember that before the simulation itself is run, a synthetic population is constructed as described in Chapter 6: this is indicated in Figure 7.4 by the Population Synthesis step. From these records, Agents are created and made to schedule themselves; they then set up personal road networks based on their home and work locations as well as the core roads which are familiar to all drivers. They are contextualized within the created environment during the Population Initialization step. Within this world, the Media object mentioned briefly above is created and scheduled to inject information into the system at appointed times. Finally, data structures to record the heatmap trace are initialized, and reporters for stress and speed are created and scheduled.

7.5.2 Input

As described above and in the previous chapter, the model takes a large quantity of data as input. Firstly and most obviously is the synthetic population, along with its social ties. This has been discussed extensively in the previous chapter and will not be elaborated upon further here. The wildfire location is taken as a given, and because the wildfire is initialized with its own seeded random number generator for the purposes of this chapter, the wildfire itself can be understood as input here. The evacuation areas, drawn from the real evacuation areas utilized in the Colorado Springs evacuation, are taken as input, as are the times at which the orders were issued through mass media and emergency notification

systems. The final piece of input information is the road network itself.

7.5.3 Submodels

A number of complex and interconnected submodels support the simulation. The Wildfire object behavior stands essentially by itself. The various Agent procedural submodels, however, interact in various ways, and will be introduced both in unity and as separate parts.

Wildfire

The wildfire submodule draws extensively on the work of Alexandridis et al. (2008). The Wildfire object contains fields representing the elevation, vegetation cover, and impermeability of the entire space, as well as a record of the status of the individual locations in the environment with regard to whether a given location has burned, is burning, or is vulnerable to burning. Every parcel of space is initially assumed to be vulnerable to burning. The Wildfire can be triggered at any point in the environment, at which point it creates a lower-level process called a FireTile which is associated with the given location. When a FireTile is created, its location transitions from being vulnerable to being on fire. At the next tick of the simulation, the FireTile may trigger the creation of FireTiles in its immediately neighboring locations. Specifically, if a neighboring location is both permeable and vulnerable, it will catch fire with a probability based on the current wind speed, the location's landcover, and the difference in elevation between it and the triggering FireTile. If the location catches fire, a new FireTile is generated for it.

Certain parameters constrain the spread of the fire, including the impact of wind, slope, and the landcover type. Alexandridis et al. (2008) call the baseline probability of the fire spreading P_{spread} , and factor the impact of different landcover types using the probabilities $P_{vegetation}$ and $P_{density}$. The mapping of vegetation types to values of $P_{vegetation}$ and $P_{density}$ are taken by merging the values used by Alexandridis et al. (2008) with the vegetation types given in the landcover data. Other parameters include the windspeed W and the empirically

derived wind-related constants C_1 and C_2 . Finally, the constant a describes the impact of slope on the spread of fire. Certain equations differ slightly from the Alexandridis et al. (2008) equations, in that simplifying assumptions have been made about the wind speed and direction. In the results presented in Chapter 8, the wind direction is set according to the direction reported by the After Action Report (City of Colorado Springs, 2013) on the highest day, with low speed to avoid overstating the effect; in the future, dynamic wind speeds and directions could be included, but the variation described by the After Action Report (City of Colorado Springs, 2013) is intensive and beyond the scope of this work. The equations which dictate the probability of the fire spreading from tile to tile are as follow.

$$\begin{aligned}
P_{wind} &= e^{W(C_1 + .5C_2)} \\
P_{slope} &= e^{a\theta_{slope}} \\
P_{burn} &= P_{spread} * (1 + P_{vegetation}) * (1 + P_{density}) * P_{slope} * P_{wind} \quad (7.1)
\end{aligned}$$

As the wildfire expands, a simplified polygon representing the hull of the burned area is calculated for the use of the Agents. This simplification drastically decreases the computation costs of comparing complex geometries. The wildfire is initialized as shown in Figure 7.6. A review of the performance of the Wildfire submodel is given in the verification and validation efforts addressed in Section 8.2. As it grows, the Wildfire updates the road network by making roads which are within the boundary of the fire impassible. Agents who attempt to move along these impassible roads will not make any progress, and will observe that the roads are functionally closed. This phenomena is explored in greater detail in the next section.

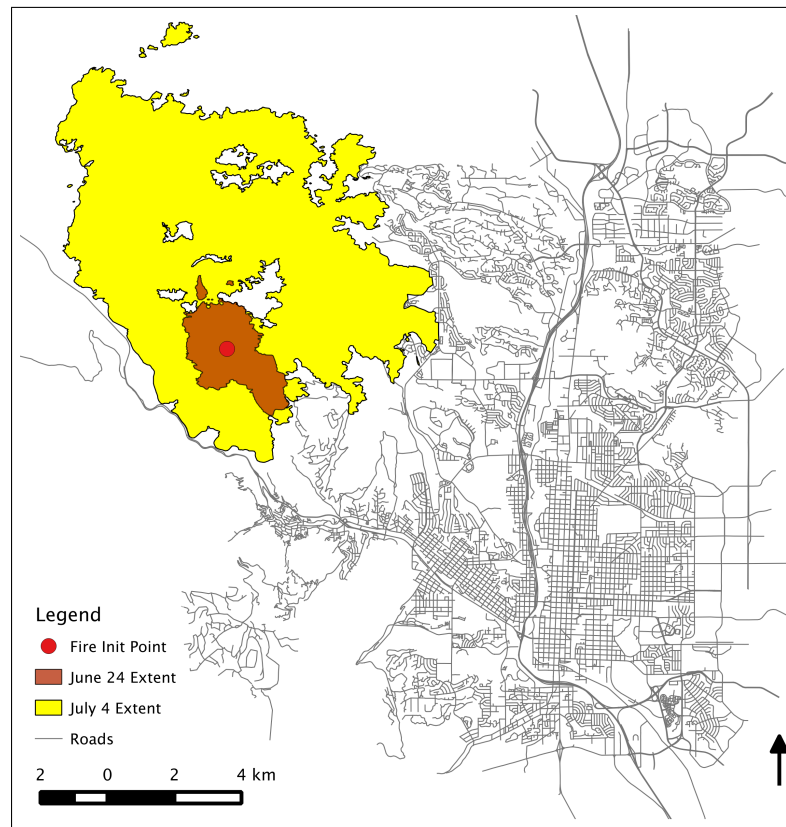


Figure 7.6: Close-up of study area with the point at which the fire is initialized in the model versus the true extent of the wildfire on July 24 and June 4

Agent Submodels

As described above, Agents are explicitly spatial and mobile, contextualized within the physical world and aware of their surroundings. They are also social, enmeshed in relationships with one another which provide them with information and influence their emotions. Chapter 2 addresses the role these features of individual experience play in the decision to evacuate, and an Agent makes choices based on each of these spaces of experience. Absent interruptions and threats, Agents follow a default set of behaviors, a basic course of action upon which the other behaviors build: Figure 7.7 shows the activation process of the Agent, with its various submodels being activated in turn. Observation, movement, communication, and decision-making are the other four major components of Agent behavior, although lower-level supportive behaviors such as risk assessment are also present. These modules interact among one another within any given Agent, and sometimes bridge between different Agents (as in obviously the case of the communication module). In this model, the only emotion being modeled is essentially stress, and is referred to in the following as the agent's (emotional) valence, following the use of the term as introduced in Chapter 5.

Agent Module - Default Behavior

The Agent's typical pattern of activity is structured around the default daily behavior, shown in Figure 7.8. Agents select their wakeup time from a uniform distribution between 6 and 8am. If they are employed, they select their departure for work from a uniform distribution of times between 6:45 and 8:45am. These times were selected based on the most frequent departure times nationally (American Association of State Highway and Transportation Officials, 2013), and were compared with Google Map's reported traffic conditions on Mondays to ensure that they were reasonable for the area in question. Figure 7.9 shows the expected traffic at 7:45am and 5:30pm, solidly during the middle of the commuting period, while Figure 7.10 shows the expected traffic at midnight as a point of comparison. After they arrive at work, Agents will schedule themselves to leave work at a time uniformly

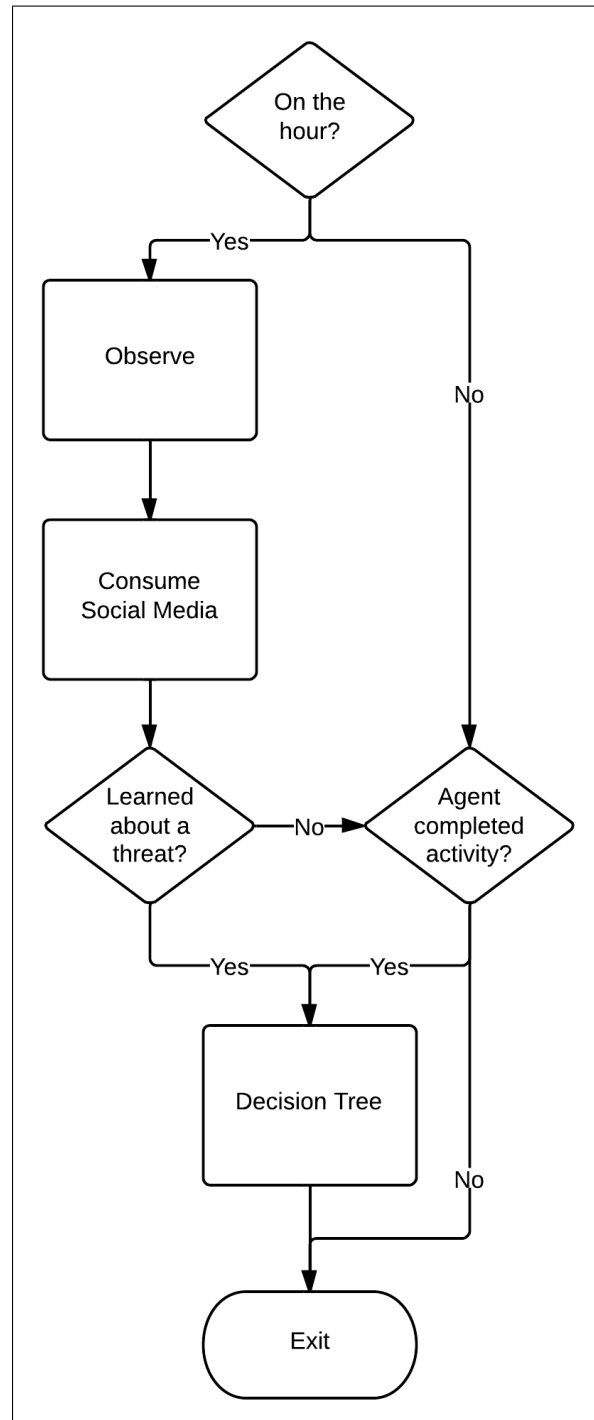


Figure 7.7: Overview of an agent's activation of its subprocesses

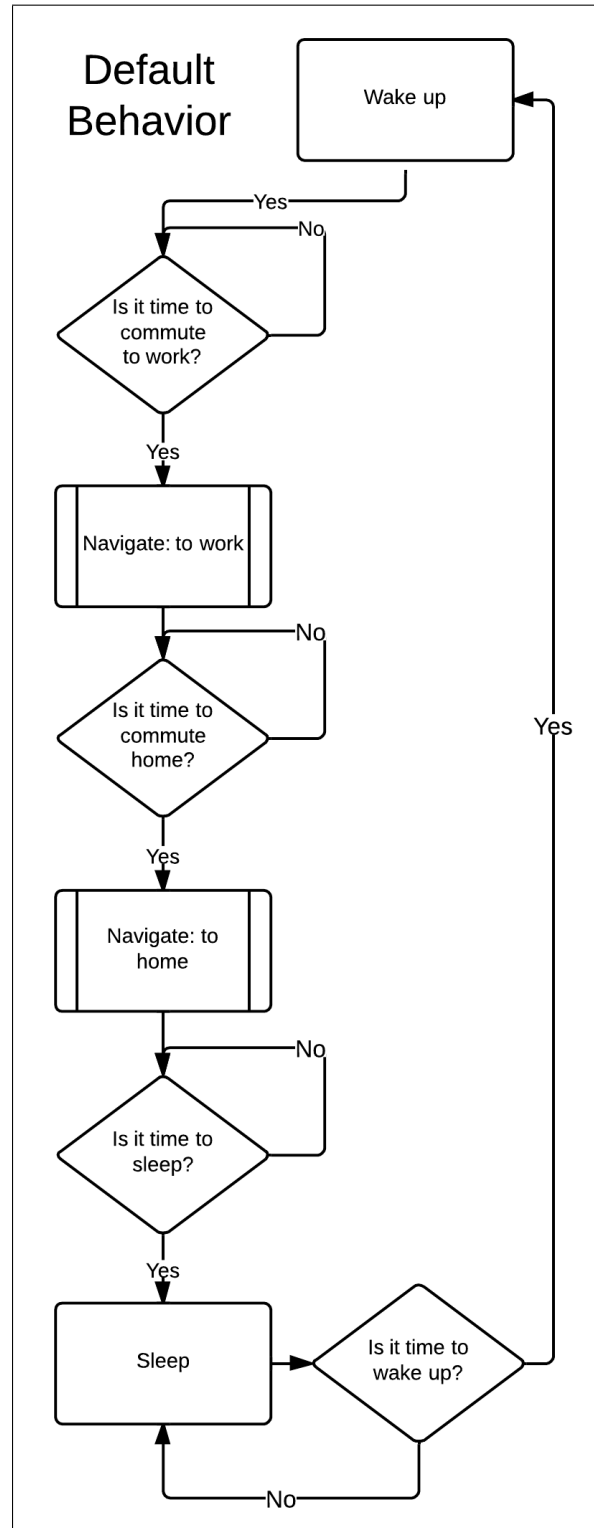


Figure 7.8: The default behavioral subprocess of an Agent

distributed between 4:30 and 6:30pm, based on standard working hours and the traffic patterns shown in Figure 7.9. After commuting home from work, Agents will relax until they go to sleep at a time uniformly distributed between 9 and 11pm (values drawn the National Sleep Foundation, 2008). This schedule structures the Agent's time, so that an Agent who is disrupted from their workday by calls from intimate network members will deal with the calls but be able to resume their pattern of working, commuting, and sleeping if they are in no danger. Behaviors like navigation support this higher-level behavior, while observation and communication function in concert with it.

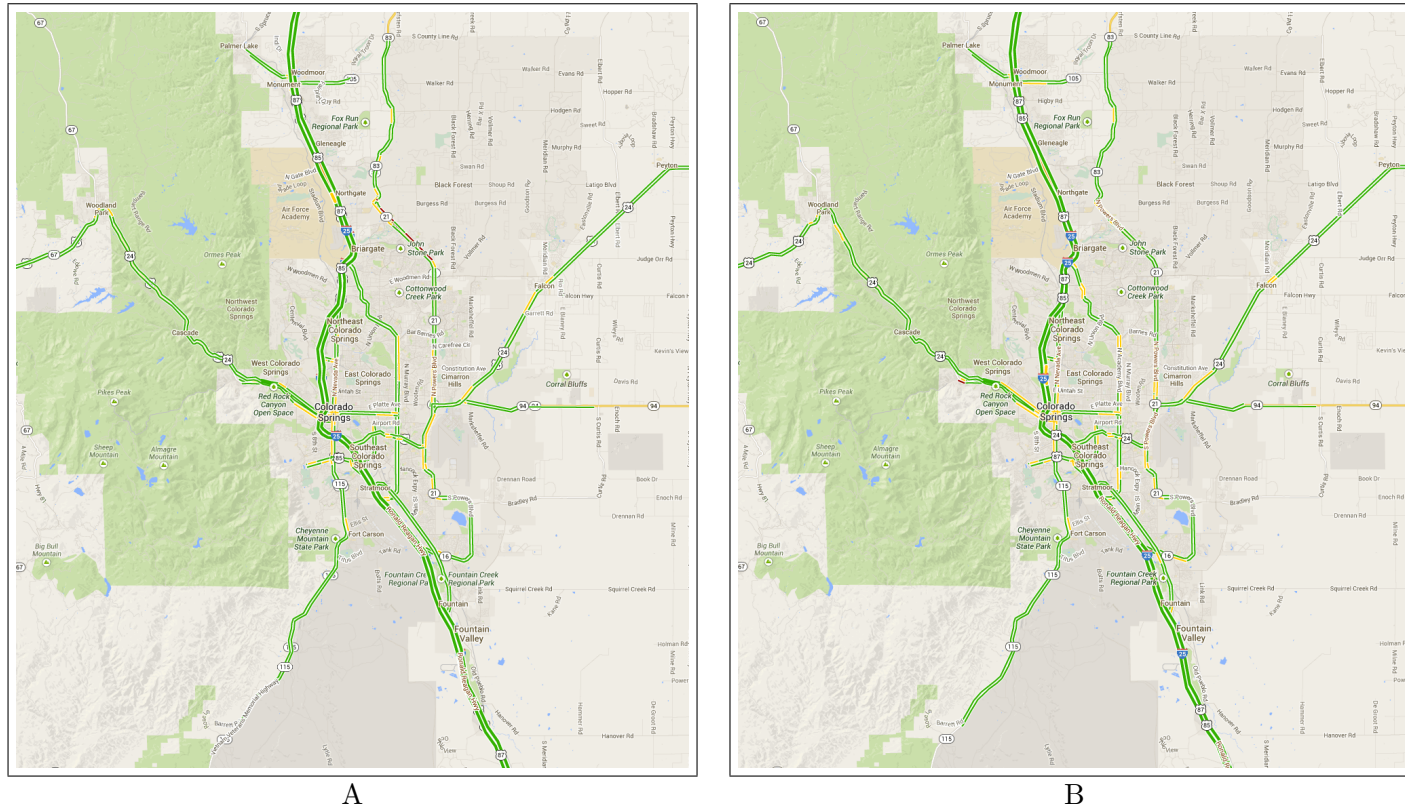


Figure 7.9: GoogleMap's recorded average traffic data for Colorado Springs at A) 8am and B) 5pm on a Monday. Red areas indicate heavy traffic while green areas indicate relatively free-flowing traffic

Agent Module - Observation

Agents directly observe the environment around them, and can recognize the presence of a Wildfire if they are within a given distance of it. This process is visualized in Figure 7.11. Based on these observations, they make assumptions about road closures around the Wildfire, so that they do not plan paths that utilize roads which are within the extent of the Wildfire itself. If the Agent is traveling when it observes the Wildfire, it will check to make sure that its path is still valid, even if it decides not to evacuate. Regardless of whether they have observed the Wildfire or not, Agents check to see whether they've been encompassed by the wildfire and are therefore deceased, removing themselves from the simulation if so. Finally, the observation submodule updates the Agent's emotional valence based on the various pieces of emotional stimuli it has observed. The valence V given decay parameter d and valence-bearing event v_i at time i is given by the equation

$$V = \log \sum_{i=0}^t \left(\frac{t-i}{v_i} \right)^d \quad (7.2)$$

Agent Module - Navigation

Agent movement and therefore evacuation happen along networks of connected road elements. The structure of the process which controls their navigation is shown in Figure 7.12. After being initialized with knowledge of the existing road network, an Agent maintains a personal record of the parts of the road network about which it knows, route planning based on this known network if possible. Agents use a simple A* path planning algorithm on their set of known, available roads, avoiding roads that they know to be inaccessible. If an Agent cannot find a path to its target destination through its personal road network, it will revert to obtaining a GPS-derived route from the road network as it exists. This assumes that an Agent can get access to a GPS unit or to mapping software via a smartphone or some other




Figure 7.10: GoogleMap's recorded average traffic data for Colorado Springs at 12am on a Monday. Red areas indicate heavy traffic while green areas indicate relatively free-flowing traffic

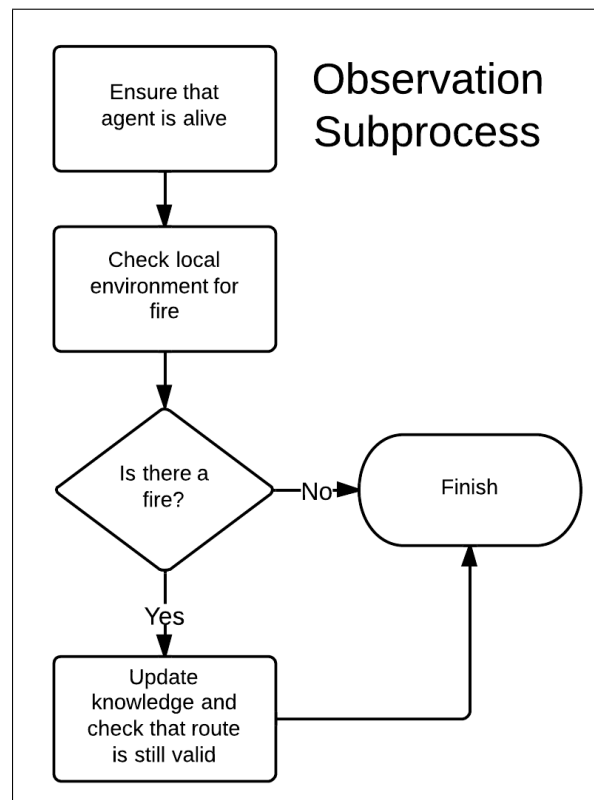


Figure 7.11: The observation subprocess of an Agent

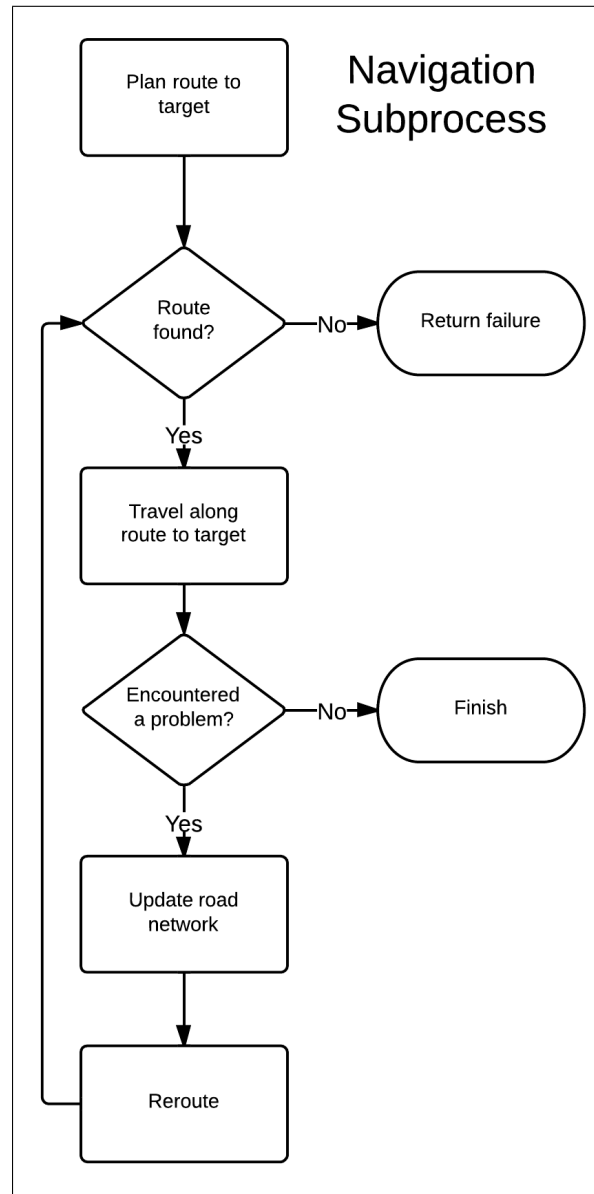


Figure 7.12: The default behavioral subprocess of an Agent

route-planning technology. The Pew Research Center reports that over 50% of American cell phone owners have a smartphone, and are therefore capable of accessing this kind of route-planning information from wherever they are (Duggan & Rainie, 2012). Having set in a route, the Agent proceeds to move through the space.

The Agent's progress can be slowed by the presence of other Agents in its path or stopped by closed roads. If a road is closed, the move step will fail and the Agent will update its known road network, replan a route, and attempt to move again. The newly acquired road closure information becomes a piece of knowledge that the Agent can share with other Agents during the communication subprocesses. If the Agent cannot find a path to its desired target at all, it will stay in place until it receives new information.

If a road is open, the Agent will move along it at a speed which depends on how dense the traffic is on the road link. Agents are assumed to attempt to keep 20 meters between themselves and the other cars on the road when traveling at top speed in order to allow the Agents sufficient stopping distance in the event of the car in front of them suddenly braking. If there are enough Agents along the length of road that leaving this vacant space is impossible, their speed is scaled down proportionate to the available space per Agent. Agents are assumed to travel at a speed of at least 5 miles per hour in order to prevent complete blockage when movement is possible but slow. This simple way of calculating traffic flow is imprecise in capturing the lower-level interactions between individual vehicles, but it allows for the model to handle otherwise prohibitively large numbers of navigating Agents. In future work, it would be possible to add simple probabilities like the chance of an Agent's vehicle breaking down.

Agent Module - Communication

Agents are embedded within social networks derived from social media and implied family relations, as described in Chapter 6. Because humans interact with one another in any given context depending on their relationships and available modes of communication in the ways discussed in Chapter 4, the model attempts to capture some of these differences

within the communication module. In the real-world, individuals choose whether or not to interact with one another in a specific situation based on their relationship or lack thereof: talking to a friend versus a stranger on the bus might be very different things, and the acceptability of such communication might vary across different cultures depending on the circumstance.

In reality, the success two individuals enjoy in their attempts to communicate can be modulated by the means of communication a social media message might be lost among many other messages, spoken conversation may be impossible over the noise of traffic, or a caller might get a busy signal instead of their friend on the line. Thus, relationships, varying probabilities of success in exchanging information, and context should all factor into a reasonable model of human communication. Further, when Agents receive new information, the information contains a valence value of its own, so that learning certain kinds of information will impact the Agent's emotional valence (a road closure is a mildly stressful event, while receiving an evacuation order is arguably much more stressful). In this model, three kinds of communication are simulated:

Intimate intentionally contacting members of an intimate social network, as by phone, text, email, personal message, etc. Wellman and Wortley (1990) note that the strength of a relationship is correlated with frequency of telephone contact and does not vary with distance, in contrast to less intimate relationships. These intimate relationships are important channels in the spread of information when individuals are deciding to evacuate Drabek (1992) indicates that in a wide range of crises many individuals get their information directly from family members. Drabek further notes that individuals tend to regard information from friends or family members as being more credible than information from mass media channels, although less credible than information from officials (Perry & Greene, 1983).

In the model, Agents are endowed with undirected intimate social networks that indicate their closest and most intimate relationships. A link in the network from one Agent to another means that the originating Agent would contact the receiving Agent in the event

of a disaster to ensure that she was aware of the threat and alright. Thus, individuals who have frequent but not intimate contact are left out of this network - perhaps an individual might see her coworker every weekday, but would probably not call his home to ensure that he was evacuating unless the two had a relationship outside of work. The structure of this network is constructed based on data about intimate social networks, drawing especially on Berg et al. (2010), Watts (2003), and Albert and Barabási (2002), and is detailed in Chapter 6.

Given this social structure, an Agent who becomes aware of a crisis will contact its intimate contacts in order to share information about the crisis. Any given attempt at contact is not guaranteed to be successful - if an Agent attempts to call another Agent who is already in the process of communicating, the call will not go through and the calling Agent will be free to try to call someone else. An Agent may attempt to call five other Agents per tick of the simulation, and will call another Agent no more frequently than every 10 minutes. These numbers were selected to capture the process of dialing the phone and waiting for the contact to pick up the phone - the individual is assumed to spend at least a minute picking a contact, dialing their number, and waiting for the contact to pick up the phone, resulting in the limitation of five contacts per tick. The 10-minute interval prevents the individual from continuously calling the top five most intimate contacts. Even if the other Agent is not occupied, the call has a certain percent chance of failing to be picked up, as the Agent receiving the call may not hear the phone ring or may be otherwise occupied. Agents keep a record of when they last contacted any other given intimate contact, to ensure that they don't repeatedly call the same individuals.

Local individuals interacting with others in their immediate vicinity. Other people in one's immediate physical vicinity - coworkers or residents of the same apartment building - can often be a source for important information in crisis scenarios, as described in the milling, information-gathering behavior described earlier (e.g. Drabek, 1992; Kuligowski, 2011; Kuligowski & Mileti, 2009; Sherman et al., 2011). Sherman et al. (2011), discussing

the evacuation of the World Trade Center in 2001, note that once an individual perceives an immediate threat to her life, she will cease this milling, information-gathering process and take action. However, when “context cues are ambiguous” she may continue to seek information rather than evacuate, so this process is extremely important in accurately capturing the dynamics of communication and behavior. Obviously, local communication is common in low-stakes environments as well, as the popular trope of water-cooler chatter at the office suggests. This is one case in which weak ties are important – near-strangers who would otherwise not make contact will exchange information because of physical proximity. Based on context, individuals may choose to communicate with others in their home or workplace. This is as simple as exchanging information with all other Agents within a given distance.

Broadcast the use of technological and media platforms to quickly and broadly share and consume information. As mentioned in Chapter 5, communication through social networks has been referred to in some contexts as “masspersonal” (Wu et al., 2011) – a system that drastically increases the potential audience of any one individual and massively amplifies their signal. Platforms such as Facebook and Twitter have this capability, although Pew Research (Duggan & Brenner, 2013) indicates that the specific audiences they reach differ. Regardless, individuals both consume and propagate information, making information cascades even easier. Sutton et al. (2008) note that in many cases, individuals turn to social media platforms for information because media sources have information that is inappropriate for the specific location (incorrect roads, lack of regional context, etc). However, there is also a great deal of shared information that is off-topic, so that it is important to incorporate this dynamic of noise in the signal.

In addition to intimate social networks, selected Agents are embedded in social network platforms that allow them to communicate with one another masspersonally. In this simulation, the platform being emulated is Twitter, and the population of connected individuals is constructed in order to match the usage profile of Twitter users. Individuals have a certain base frequency of checking their social media accounts, during which time they may

introduce a new piece of information into the system or propagate existing information. They acquire new information by taking at most 30 of the new posts generated on their network since their last check-in and, with some probability, reposting the information on their own account. The 30-item limit is imposed based on the average amount of time spent on socializing and communication: the American Bureau of Labor Statistics (2012) reports that in 2012 American adults who socialized spent almost 2 hours per day on socializing and communicating outside of social events, so that for each of 16 waking hours a bit more than 5 minutes are dedicated to social media consumption, giving the Agent about ten seconds to consume each piece of information. Thanks to the limit which emulates the bounds of attention and time, Agents may fail to perceive a specific piece of information. While not all Agents utilize social media, all Agents are assumed to have access to television, radios, public announcement systems, and so forth, so that news of an emergency might reach them that way. This aspect of broadcast information is captured by the dummy “Media” data stream which injects official information into the environment, as mentioned above.

All of these different types of communication can take place on vastly different timescales. Crooks et al. (2013) note that information about the mild earthquake that hit the East Coast of the United States in 2011 actually traveled faster than the earthquake itself (fulfilling the prophecy of Munroe, <http://xkcd.com/723/>). However, obviously an individual can make only so many phone calls during a 5 minute period. While the process of information distribution can be very fast, Kuligowski and Mileti (2009) note that information gathering and communication take time and can delay action in disaster situations. Thus, it is important to ensure that communication takes time but also proceeds on a reasonable timetable.

The communication processes are therefore split up and activated in a variety of different ways. During the decision step, if a problem exists, Agents will initially attempt to discuss the situation with others in the immediate area. Local communication, being so important to crisis situations, is the first response (see Drabek, 1992). If this is infeasible because no one else is around, an Agent will move on to intimate communication, attempting to reach

out to specific other Agents in search of more information (e.g. Sherman et al., 2011). In situations where an Agent is aware of a problem but is not itself threatened, it may choose to reach out to other Agents in its intimate social network in order to ensure that they too are aware of the problem (Carey, 2002). Independent of these, Agents are regularly and constantly consuming media, either through their social media network or their access to media information in general, so that Agents consume information on a regular timestep throughout the day unless they are asleep. The activation of these processes will be clarified further in the next section.

Decision-making

Finally, all of the submodels must be combined and synthesized into a greater structure. The Agent needs to have a way to select its course of action during any given step. The decision-making process is implemented as a cognitive model rather than an entire cognitive architecture both for ease and to better emulate the processes described in the crisis-situation behavioral literature presented in Chapter 2. Cognitive models deal with the decisions an agent faces given the stimuli and inputs to which it is being subjected, selecting a course of action based on all the information it has been given. This stands in contrast to a cognitive architecture, which is “a specification of the structure of the brain at a level of abstraction that explains how it achieves the function of the mind” (Anderson, 2007). Rather than focusing on decision-making at the level of brain processes, the model utilizes the Fast and Frugal heuristic-based approach (see Goldstein & Gigerenzer, 2002) after the example of Kennedy and Bassett (2011).

Following Kennedy and Bassett (2011), the model utilizes a decision tree based on the search, stopping, and decision rules for heuristics described by Gigerenzer and Todd’s (1999) adaptive toolbox. The decision tree implemented in this model is given by Figure 7.13, and the specific behaviors are broken down into the actions described in Sections 7.5.3–7.5.3 and shown in Figure 7.8. In essence, there are three behavioral options: the default workday-commuting-sleeping process, targeted information-gathering, and evacuating. The details

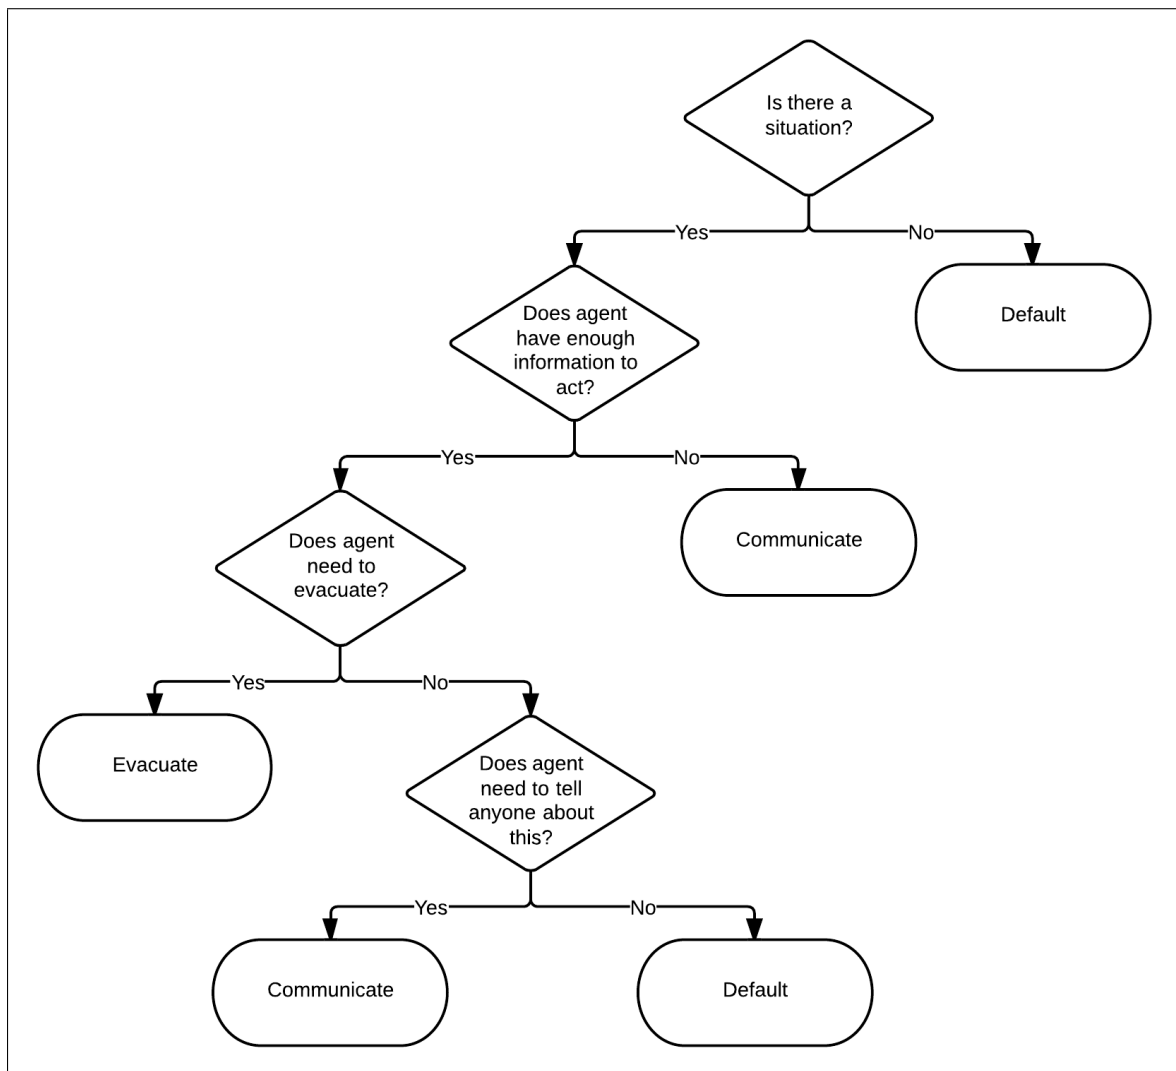


Figure 7.13: A Fast and Frugal decision tree for Agents

of these behaviors are elaborated upon below.

Default as described in Section 7.5.3, the default behavior of the agent is to proceed with its daily tasks, such as working, commuting, and relaxing. Obviously the default behaviors of individuals could vary based on their careers, ages, and the time of day, so this behavior could be more or less specified. The basic structure, however, is that employed individuals have a morning commute, a work day, and an evening commute, followed by a sleep schedule. The time at which these begin and end vary uniformly around set commuting peak times.

Seek Information an Agent who perceives a need for more information based on the information she already has may choose to communicate, that is, actively seek out contact with individuals within its social network. An Agent will communicate first with the other Agents in its home or workplace, and then later reach out to others within its social network. Agents have a certain probability of success of contacting another Agent and sharing information with them. As mentioned in the communication section, emotional valence is also communicated in this fashion.

Evacuate if the information an Agent has gathered suggests that it itself is in danger, an Agent proceeds into evacuation behavior. More precisely, if the Agent is aware of a Wildfire and the distance between the Wildfire and the Agent or its home is within the Agent's comfort distance divided by the Agent's stress level, the Agent will decide to evacuate. To be precise, when the distance between the Agent and the Wildfire or the Agent's house and the Wildfire is less than the Agent's current stress distance S , given by $S = \frac{\text{comfortdistance}}{\text{valence}}$, the Agent will choose to evacuate. This evacuation process consists of going into a movement process involving selecting an evacuation target and trying to reach it, selecting a new target whenever the current target becomes unavailable or unreachable.

Given how interwoven the different aspects of Agent behavior are, it may appear that calling the different aspects of Agent behavior "modules" is inappropriate. The observation, movement, and communication submodules are connected in various places to these behavioral modes so that it can be difficult to disentangle them enough to refer to explicitly

separate modules of behavior. However, they all combine together to form a complex system with interlocking parts, and presenting them in this way is perhaps the most straightforward way to understand such a system.

Despite the complexity of these interdependencies, the behaviors as they exist are quite simple. It would be particularly interesting in future work to expand upon the default behavior. For example, it would be fascinating to explore whether the timing of evacuation announcements might impact the system differently depending on the time of day or day of the week.

7.6 Summary

This chapter has presented the structure, input, and assumptions associated with a model of the evacuation of the Colorado Springs area in the context of the 2012 Waldo Canyon wildfire. The chapters that follow present a brief survey of the validation and verification efforts applied to the model as well as a survey of the results of the model, evaluating the impact on the population along a variety of metrics. As designed, the model can serve as a tool for generating expected patterns of evacuation behavior as a function of evacuation order timing, wildfire position and development, and intentional road closure. Taken together with the crisis informatics efforts of researchers such as Vieweg et al. (2010) or that of Starbird and Palen (2010), such a tool will enable responders to design their response effort, test it *en silico*, compare the generated social media postings and locally observed dynamics, and so forth to be sure that the evacuation is proceeding as planned. By combining the model of expected behavior with real data of the type presented in Chapter 3, it should be possible to detect incongruities between the simulation - that is, the situation as the planners understand it should be developing - and the situation on the ground. Adding projective technology to the toolbox of crisis informatics would be a powerful step toward helping to organize response efforts, and could draw on the emerging technologies that offer such promise of new insight. Chapter 8 will address the success of the model in projecting the movement of citizens in response to the threat.

Part III

Results

Chapter 8: Verification, Validation, and Results

There terms verification and validation (sometimes collectively abbreviated as V & V) are a pair of processes which lie at the heart of agent-based modeling. North and Macal (2007) draw the distinction between the two thusly: “Verification is the process of making sure that an implemented model matches its design. Validation is the process of making sure that an implemented model matches the real-world.” In this work, the term “verification” is used to refer to the process of ensuring that the implemented model matches the designed model. It involves checking that the components of the model behave as expected, a feature which is often taken for granted. Performing this type of verification is sometimes referred to as testing the “inner validity” of the model (Brown, 2006). In this work, the model was developed in an iterative fashion, with each new module being tested both individually and in conjunction with the other modules existing at the time of its development. Through this iterative verification process, sometimes referred to as unit testing, it was possible to ensure that the unique parts functioned as intended. Further, walkthroughs of the code ensured that the functions at each step were in line with the expectations with which they were designed, giving further confidence to the results. The major module tests will be explored in this section in order to demonstrate how changing or disabling various parts of the model impacts the overall system.

Supplementing this is the process of validation, which describes the extent to which the model represents the system being modelled (Casti, 1997). This comparison is not a binary judgement of valid or invalid, but a measure of the degree of fitness of the model to capture the relevant dynamics (Law & Kelton, 1991). The process of collecting data from a real-world system for comparison against the generated results can be quite difficult, and the selection of statistics of comparison has generated some debate (see Crooks et al., 2008; Pontius & Malanson, 2005). In terms of qualitative assessments of the results, Mandelbrot

(1983) argued that spatial or physical predictions generated by models need to “look right”, and Axelrod (2007) suggested the evaluation of the modeled process as the simulation itself progressed, qualitatively considering whether the development of the process seemed reasonable. The model presented here has been subject to these processes at every level of development, as specified in previous chapters, and the following sections describe the process of exploring the interaction between the different parameters and processes of the simulation, culminating in the assessment of the overall model.

First, to give a sense of the behavior of Agents in the absence of an evacuation situation, Section 8.1 presents a run of the model with no wildfire or evacuation information. That is, the wildfire submodel was disabled and Agents do not receive evacuation orders. Next, a simple demonstration of the wildfire module itself is shown in Section 8.2. In Section 8.3, a parameter sweep was carried out in order to ensure that the behavior of the model was not dependent on a specific and fragile combination of parameters. Both the control case in Section 8.1 and the results of the parameter sweep in Section 8.3 are compared with the cumulative results of the fully operational model with the default values shown in Tables 7.1 and 7.2 in order to demonstrate the variation they reflect relative to the final product. The cumulative results of these default, completely operational model results are presented in Figure 8.1, and repeated in the relevant sections for ease of contrast and comparison before being subject to their own analysis and comparison with real-world data in Section 8.4. Throughout this chapter, the results for each run of the model will be shown in terms of the average speed of all of the moving Agents per tick, the average stress of all Agents in the simulation, and a heatmap tracking the positions of Agents as they move. For each combination of parameters discussed in this chapter, the speed and stress values are visualized for each of the 15 runs associated with that combination, to give a sense of their development over each of the generated worlds. For the purposes of comparing heatmaps, the differentials between the normalized default heatmap and the mean normalized heatmap of the comparison case results are shown.

8.1 Control Case

In order to determine the impact of the wildfire and evacuation orders compared to the way the population would move normally, the model was run with the wildfire submodule and the Media agent disabled. Thus, the Agents spent the entire duration of the model pursuing their default behaviors. The results of this test case - referred to here as the no-wildfire case - are shown in Figure 8.2.

In the absence of road closures, evacuation orders, and the wildfire itself, stress is predictably baseline low. Likewise, commuting is relatively simple and uninterrupted - there are the daily rush hours in the mornings and evenings, with traffic later on in the morning commute slowing down the average rate of speed. The difference between the normalized heatmaps of the default and the no-wildfire cases is striking: Figure 8.2a shows that Highway 24, the highway that would be nearest to the wildfire, experiences far more traffic in the no-wildfire case. Similarly, many of the major roads in the evacuated areas saw more traffic in the case with the wildfire enabled, as individuals utilized major roads to evacuate town rather than commuting at least partially along backroads. The areas near homes and workplaces are also more frequently utilized in the no-wildfire case, because Agents were carrying out their normal work patterns. Thus the model performs reasonably, and in ways that suggest that the wildfire significantly influences evacuation behavior in meaningful and realistic ways.

8.2 Wildfire

The Wildfire module is presented here both as it exists at the end of the simulation run and as it exists at the end of the peak wildfire period. Comparing the generated values in Figure 8.3 with the true wildfire situation given in Figure 8.4, it is clear that the model falls short of reality. The failure to include realistic wind data for every day accounts for part of this problem. By incorporating information about wind patterns and speeds, as well as a robust model of the impact of canopy on the system, it should be possible to improve

upon the current performance. Regardless, the model generates a wildfire that falls short of operational needs but is certainly sufficient for exploring the case study presented in this work.

8.3 Parameter Sweep

Table 8.1: The range of parameters utilized during the parameter sweep. Default values are shown in bold.

Parameter	Sweep Values		
Communication Success Probability	10%	50%	90%
Contact Failure Probability	10%	50%	90%
Tweet Probability	10%	50%	90%
Retweet Probability	10%	50%	90%
Comfort Distance (m)	1000	10000	100000
Observation Distance (m)	100	1000	10000
Decay Parameter	10%	50%	90%
Max Speed (m/5 min)	1000	2000	8000

Each of the parameters was varied in turn in order to give a sense of the impact of the variable on the overall behavior of the system. Table 8.1 gives the values utilized in the sweep, with the default values in bold to give a sense of comparison. The averaged results of tuning the parameters are presented separately below.

8.3.1 Communication Success Probability


The Communication Success Probability represents the chance that an Agent will succeed in its attempts to exchange information with others. This parameter can be varied to capture things like phone connection quality, probability of shared language, and so forth. By default this is quite high - the assumption is that an individual succeeds in his attempt to share information with others in his environment. Figure 8.5 shows how road usage was affected by varying the parameter, while Figures 8.6 and 8.7 give the stress and speed

profiles. When the Communication Success Probability is lowered to 50%, the rate of usage of Highway 24 increases slightly and individuals spend less time on the roads near their homes; when communication success dips even further, the impact is more extreme, and highway usage in particular increases substantially. All other things being equal, successfully communicating the existence of threats and road closures increased evacuation. While stress varies only slightly as more stressful news propagates through the social network, the impact is more clear in the records of speed: lower rates of communication success led to fewer evacuations initially, leaving more commuting cars on the road to interact with evacuation traffic during the later rush hours. This impact is less severe once Agents have a larger change of successfully communicating, at which point the parameter ceases to influence the outcome so strongly.


8.3.2 Contact Failure Probability

Contact Failure Probability gives the probability that an Agent who is attempting to contact another, distant Agent will fail to make contact on any given call. This parameter attempts to capture the chance that the other Agent will be out of signal range on their cell phone, unavailable to take a call, or otherwise inaccessible at that moment in time. The default assumption gives the Agent an even chance of making contact with an available other Agent. By decreasing the failure rate, communication improves and marginally more people evacuate, with fewer individuals taking Highway 24, as shown in Figure 8.8. Interestingly, there also seem to be fewer individuals on the other major highways outside of the evacuation area - indeed, the evacuation effort seems to proceed more quickly than in the default case, as the mean speed dips more significantly during the evacuation period. Similarly, when contact becomes less likely, the flow of information is impeded and the usage of Highway 24 is substantially greater. Figure 8.10 shows that the mean speed is significantly higher throughout the study period after the initial reports of the wildfire spread. The stress profile shown in Figure 8.9 indicates that stress levels rise relatively more slowly, but Agents apparently spend so much time attempting to contact one another that they delay leaving


work or home and end up creating massive traffic jams.



A




B




C


Figure 8.1: The normalized heatmap (A), stress profile (B), and speed profile (C) generated by the default parameter settings of the complete model. Time is measured in 5 minute intervals, or 288 to a day



A



B



C

Figure 8.2: The normalized heatmap differential (A), stress profile profile (B), and speed profile (C) generated by the default parameter settings of the model with the wildfire submodel disabled. Time is measured in 5 minute intervals, or 288 to a day.



A



B

Figure 8.3: The development of the wildfire submodel: the yellow area indicates the extent of the wildfire at the end of the third day (A) and at the end of the peak wildfire period, after 11 days (B)




Figure 8.4: The real Waldo Canyon wildfire as it existed in the real-world at the end of the third day: the yellow area indicates the extent of the wildfire

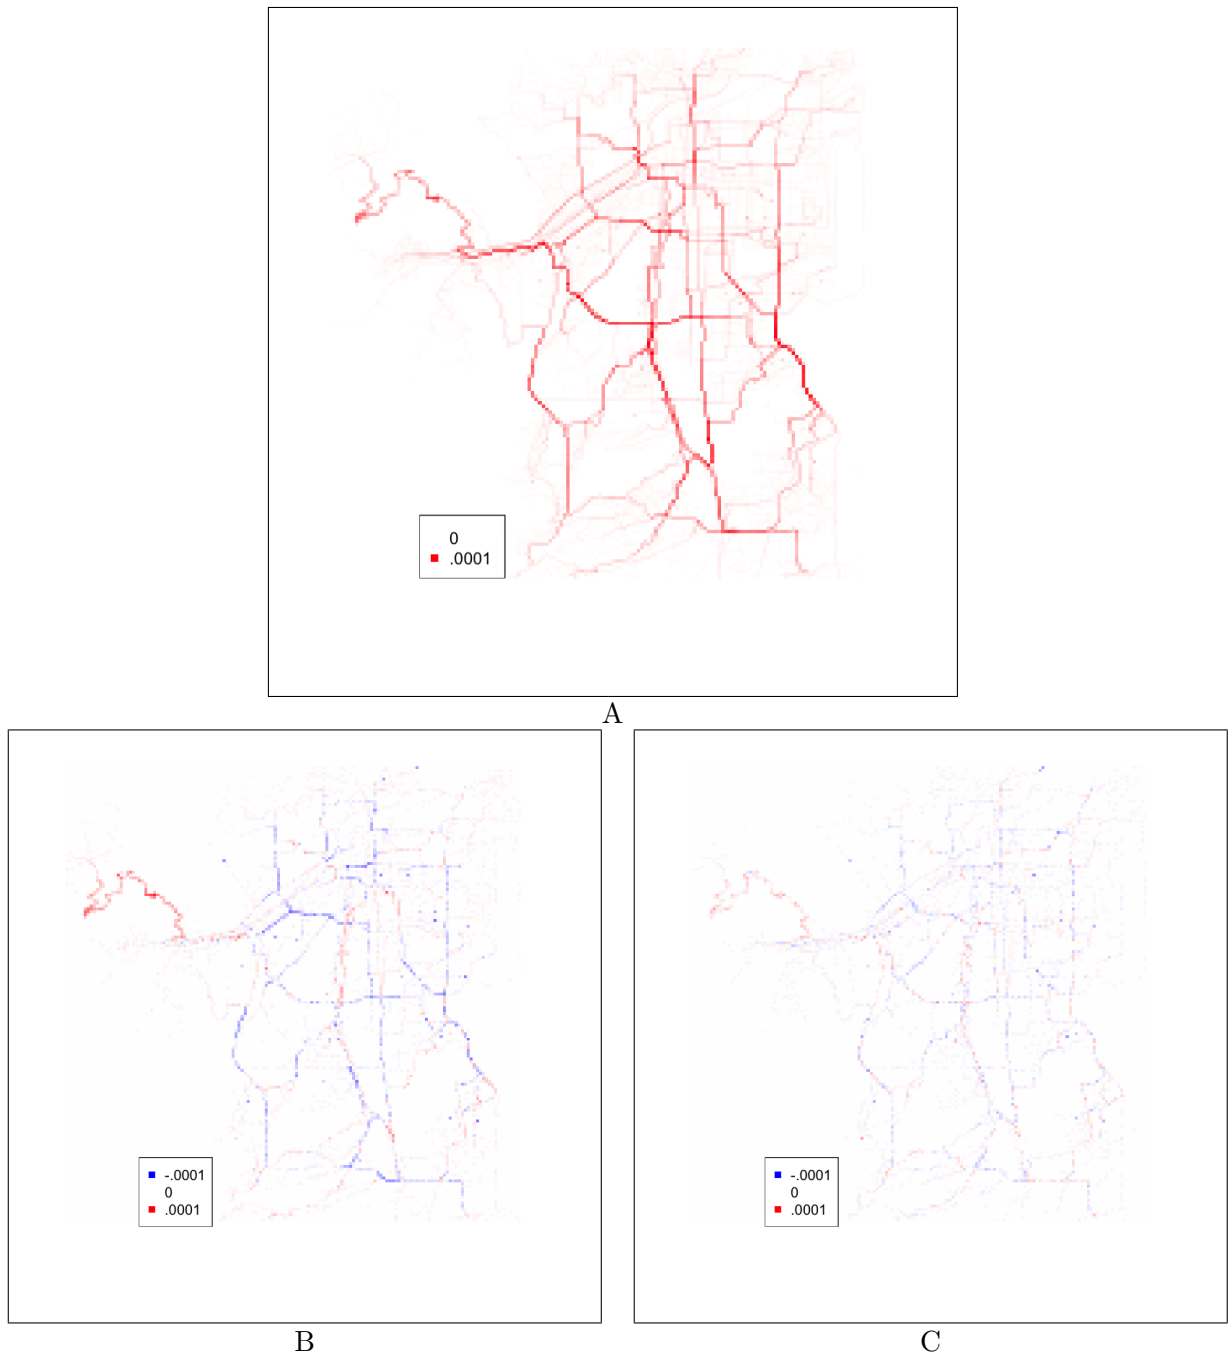





Figure 8.5: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Communication Success Probability set at 10% (B) and 50% (C).



A




B




C


Figure 8.6: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Communication Success Probability set at 10% (B) and 50% (C).



A



B




C


Figure 8.7: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Communication Success Probability set at 10% (B) and 50% (C).




Figure 8.8: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Contact Failure Probability set at 10% (B) and 90% (C).



A




B




C


Figure 8.9: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Contact Failure Probability set at 10% (B) and 90% (C).



A



B



C

Figure 8.10: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Contact Failure Probability set at 10% (B) and 90% (C).

8.3.3 Tweet and Retweet Probability

The Tweet and Retweet Probabilities deal with, respectively, the likelihood of an Agent generating information and the likelihood that it will propagate information it has found in the environment. The Retweet Probability in particular clearly influences the heatmap shown in Figures 8.11 and 8.14. Further, as shown in Figures 8.12, 8.13, 8.15, and 8.16, neither of these parameters seems to significantly impact the speed or stress profiles of either parameter setting substantially. By increasing the chance that an Agent will push information into the social network, regardless of the relevance of the information, some number of Agents fail to consume relevant information about road closures. Thus, the evacuation still occurs on the timetable shown in the default case, but it happens along roads that would otherwise not be utilized.

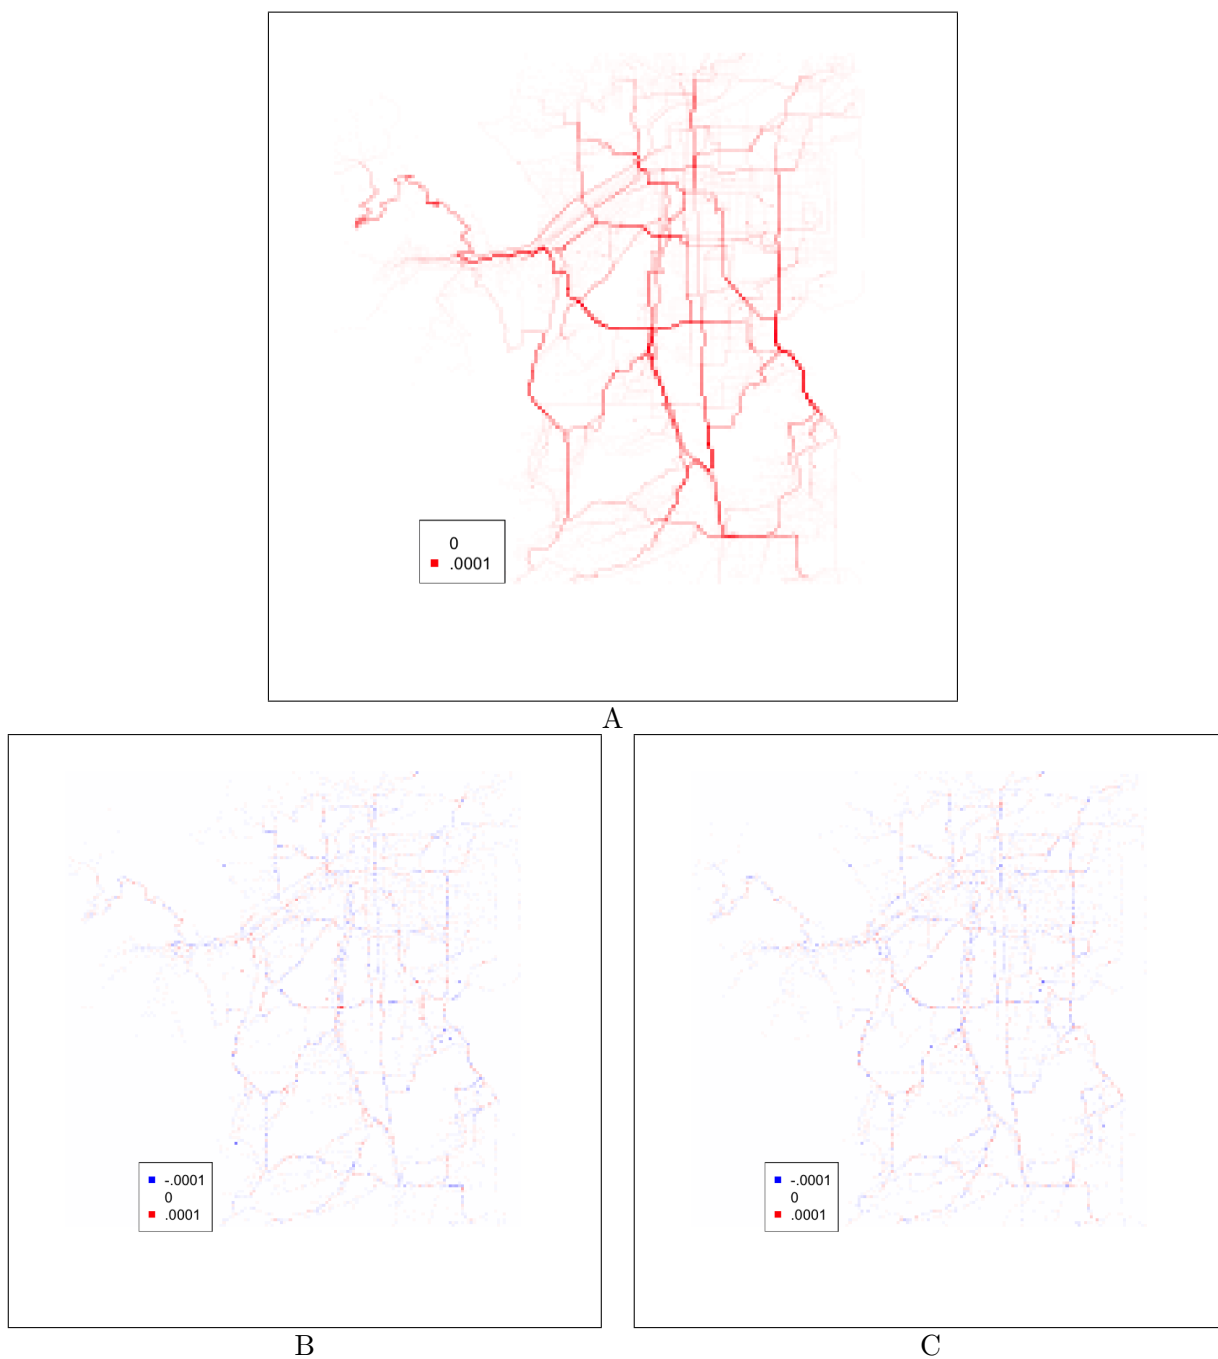





Figure 8.11: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Tweet Probability set at 50% (B) and 90% (C).



A




B




C


Figure 8.12: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Tweet Probability set at 50% (B) and 90% (C).



A



B



C

Figure 8.13: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Tweet Probability set at 50% (B) and 90% (C).

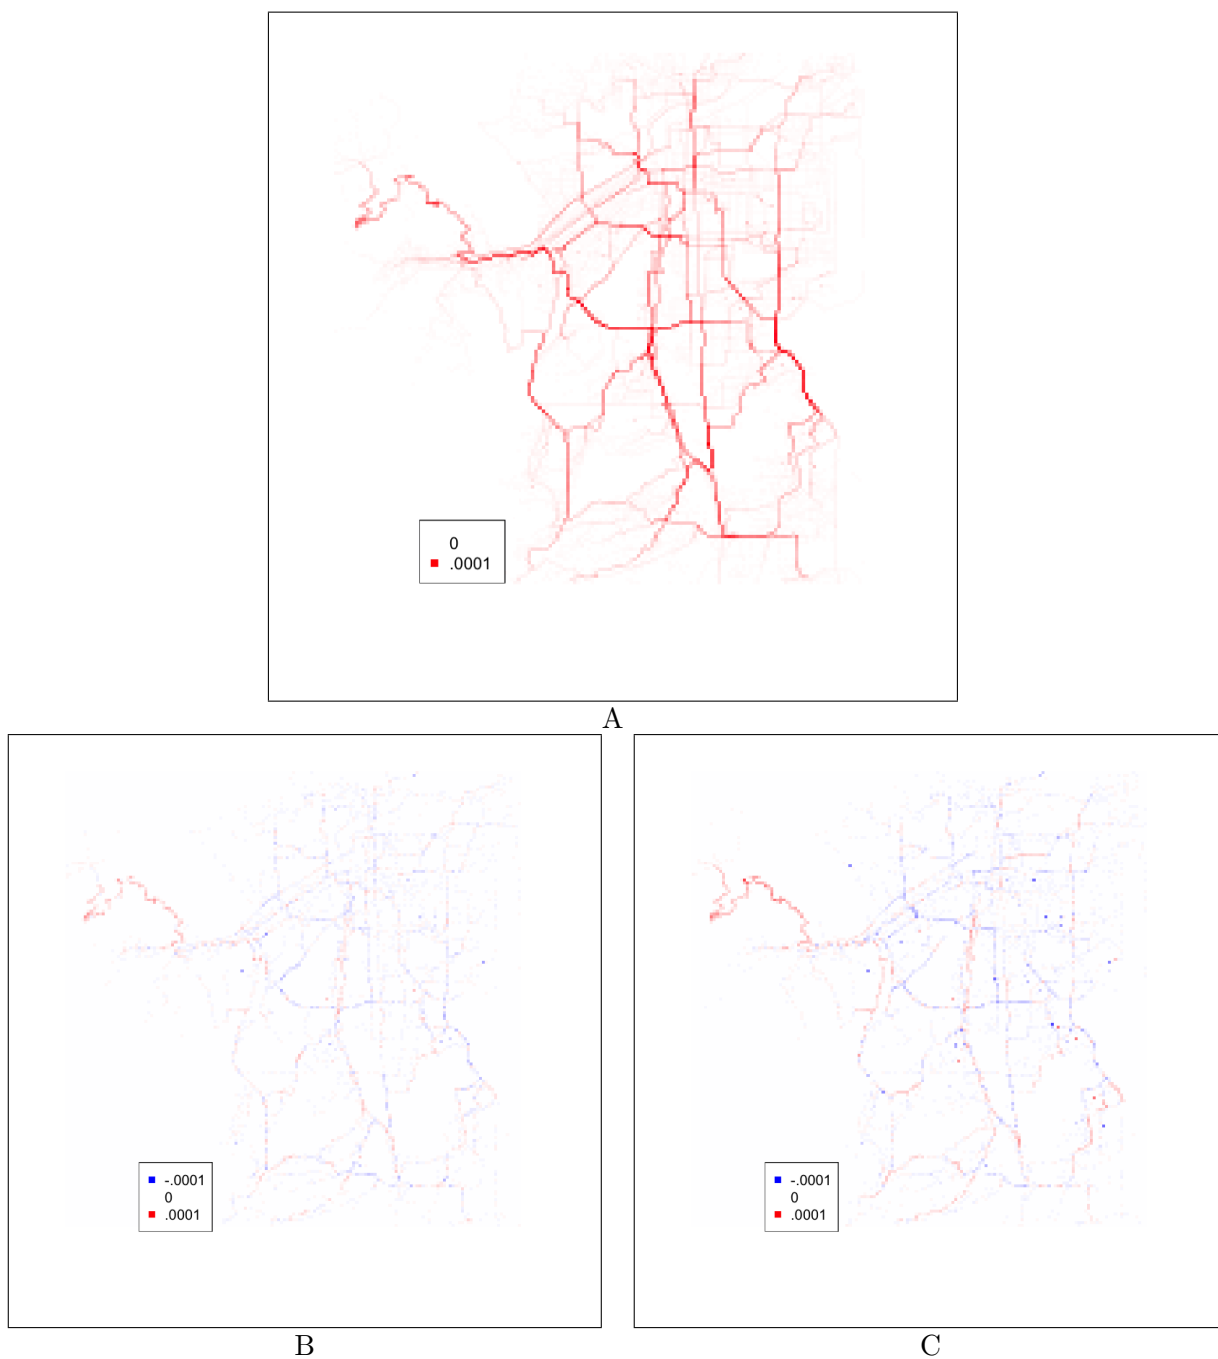





Figure 8.14: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Retweet Probability set at 50% (B) and 90% (C).



A




B




C


Figure 8.15: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Retweet Probability set at 50% (B) and 90% (C).



A



B



C

Figure 8.16: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Retweet Probability set at 50% (B) and 90% (C).

8.3.4 Comfort Distance


The Comfort Distance parameter attempts to capture the distance within which an Agent feels that it or its home is threatened. Thus, if the Agent or its home is within a range influenced by the comfort distance, it will decide to evacuate. Varying the parameter has almost no impact on the parameters in the context of evacuation orders - the heatmap, sentiment, and speed profiles all deviate from the default insignificantly, as shown in Figures 8.17, 8.18, and 8.19. As a result of the existence of other forces which supersede the influence of individual comfort, the parameter has little effect.

8.3.5 Observation Distance


The Observation Distance parameter dictates how close an individual Agent must be to the Wildfire object in order to perceive it. Varying this parameter produces some of the most dramatic influences on the population, as it determines whether Agents have the extremely stressful experience of personally experiencing the threat. While Figure 8.20 shows that the relatively short-sighted case produces a heatmap similar to the default, Agents endowed with extreme vision universally experience the Wildfire quite early on within the simulation, not needing to be informed of it by the media or their peers. Because each individual experiences it themselves and find the experience to be stressful, there is an increase in stress (seen in the stress profile in Figure 8.21) and a massive evacuation far earlier than in any of the other cases. The speed profiles in Figure 8.22 shows that because almost everyone begins to evacuate during what would otherwise be the morning commute, they leave the simulation area, leaving others with less traffic and the ability to evacuate much more easily. For the few that stay, the daily commute becomes much quicker, further pushing the heatmap toward the extreme of little activity.




Figure 8.17: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Comfort Distance set at 100 m (B) and 100000 m (C).



A




B




C


Figure 8.18: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Comfort Distance set at 100 m (B) and 100000 m (C).



A



B



C

Figure 8.19: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Comfort Distance set at 100 m (B) and 100000 m (C).

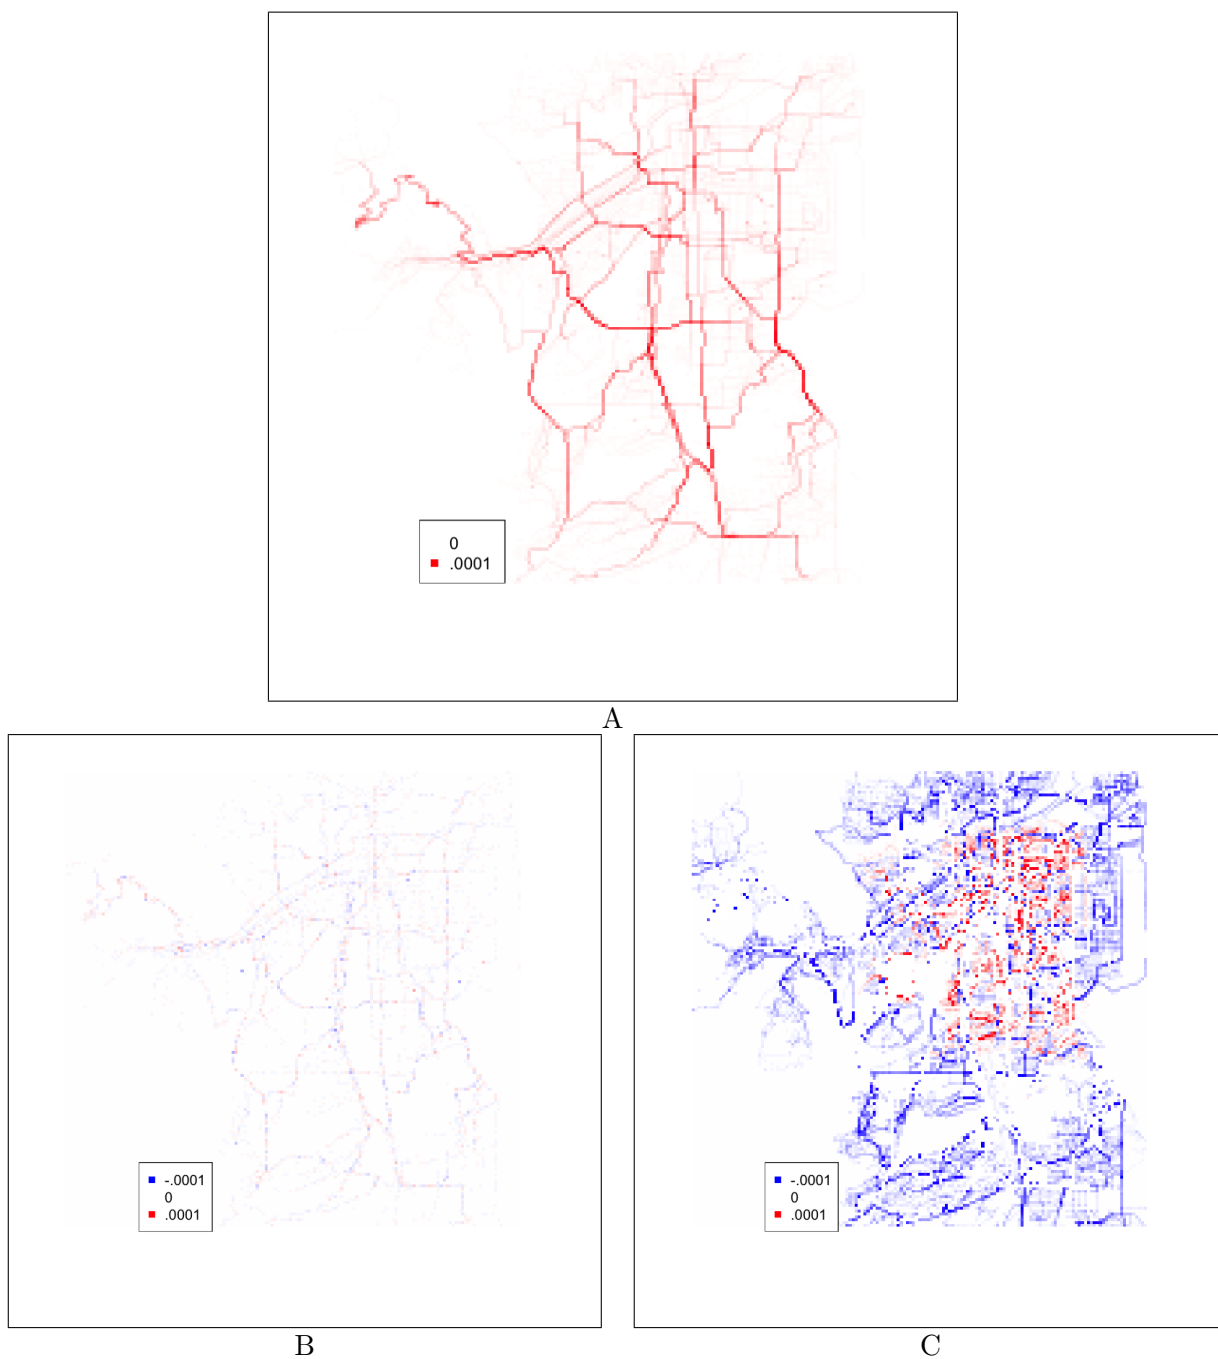





Figure 8.20: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Observation Distance set at 100 m (B) and 10000 m (C).



A




B




C


Figure 8.21: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Observation Distance set at 100 m (B) and 10000 m (C).



A



B



C

Figure 8.22: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Observation Distance set at 100 m (B) and 10000 m (C).

8.3.6 Decay Parameter

The Decay Parameter refers to how strongly Agents discount the stress value of information they have previously consumed. Agents calculate their valences as a function of exponentially decaying spikes in stress, and the Decay Parameter influences how quickly that decay occurs. While varying it has small but ultimately unremarkable impacts on the heatmap and speed profiles shown in Figures 8.23 and 8.25, it predictably seriously impacts how Agents experience their stress valences. Figure 8.24 indicates that decreasing the decay parameter (that is, weighting temporally distant spikes lower) produces lower levels of stress activation; increasing the parameter produces higher levels of stress across the population. Thus, the parameter functions as designed.

8.3.7 Maximum Speed

The Max Speed parameter influences the maximum possible speed Agents can achieve when network conditions permit. By varying the maximum permitted speed, the Agents could in theory move through the network slower or faster, easing or complicating the evacuation and commuting efforts. In fact, however, given the traffic conditions, the maximum speed has little impact on the simulation. The stress and speed profiles shown in Figures 8.28 and 8.27 vary imperceptibly from the default, and the heatmaps shown in Figure 8.26 imply only limited variation. With so few Agents able to achieve maximum speed to begin with, the parameter might influence the behavior of Agents only in situations where there were fewer Agents on the road for some reason.

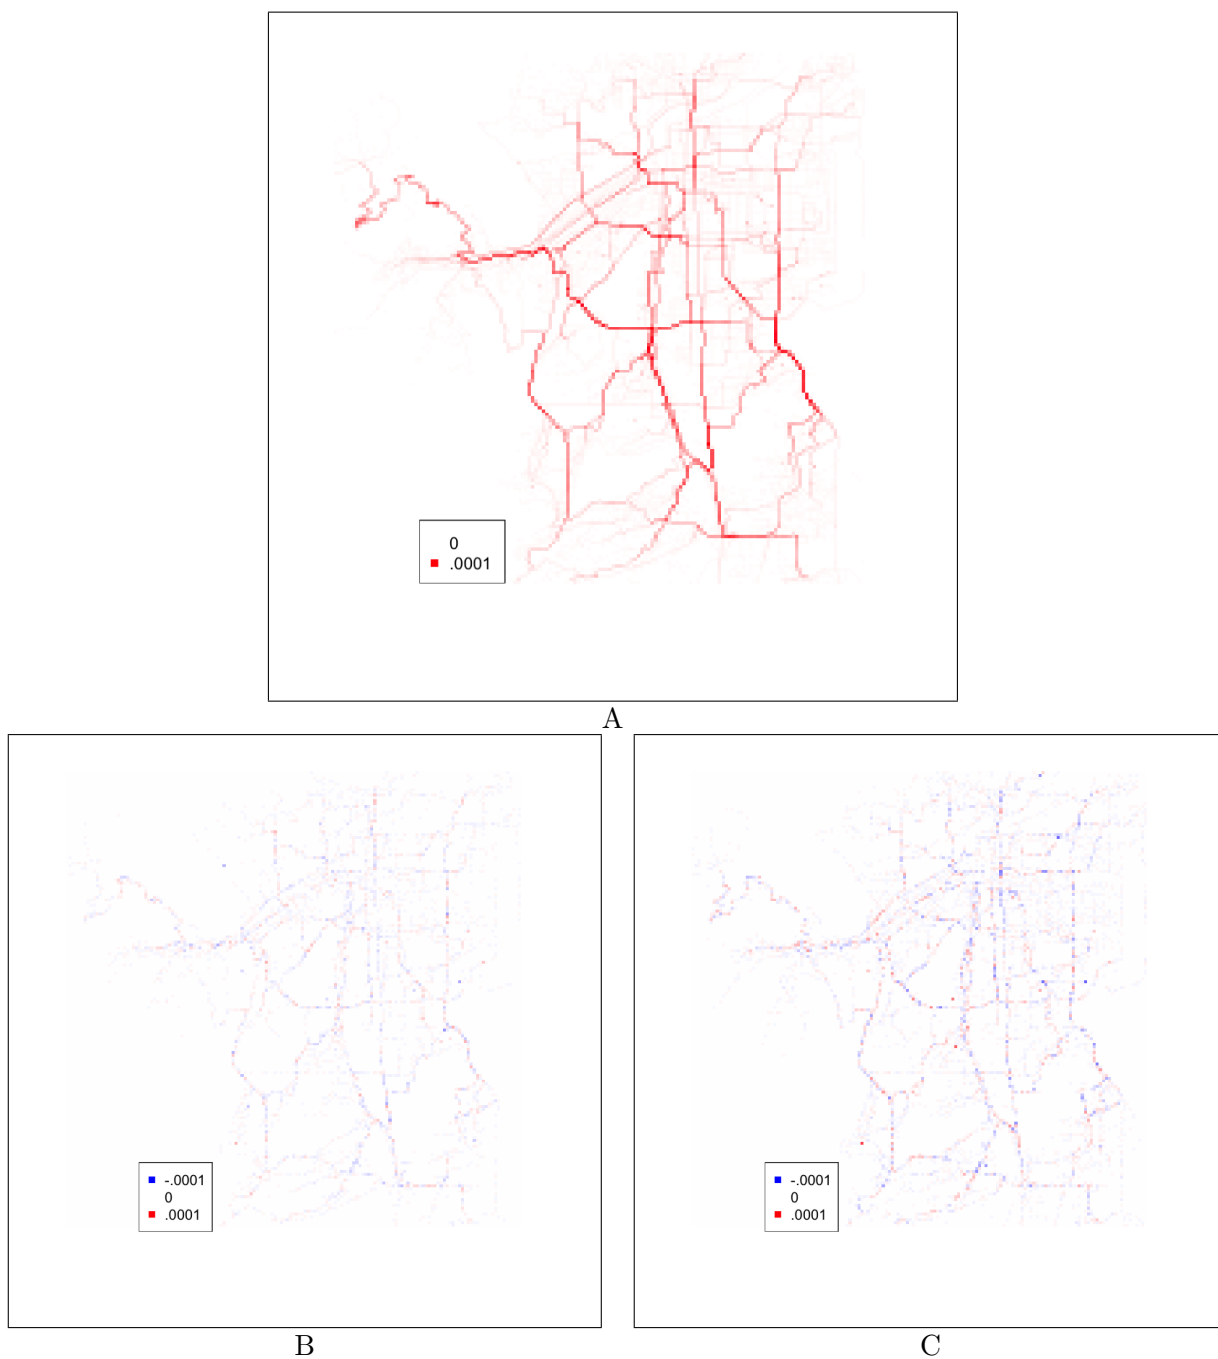





Figure 8.23: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Decay Parameter set at 10% (B) and 90% (C).



A




B




C


Figure 8.24: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Decay Parameter set at 10% (B) and 90% (C).



A



B



C

Figure 8.25: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Decay Parameter set at 10% (B) and 90% (C).

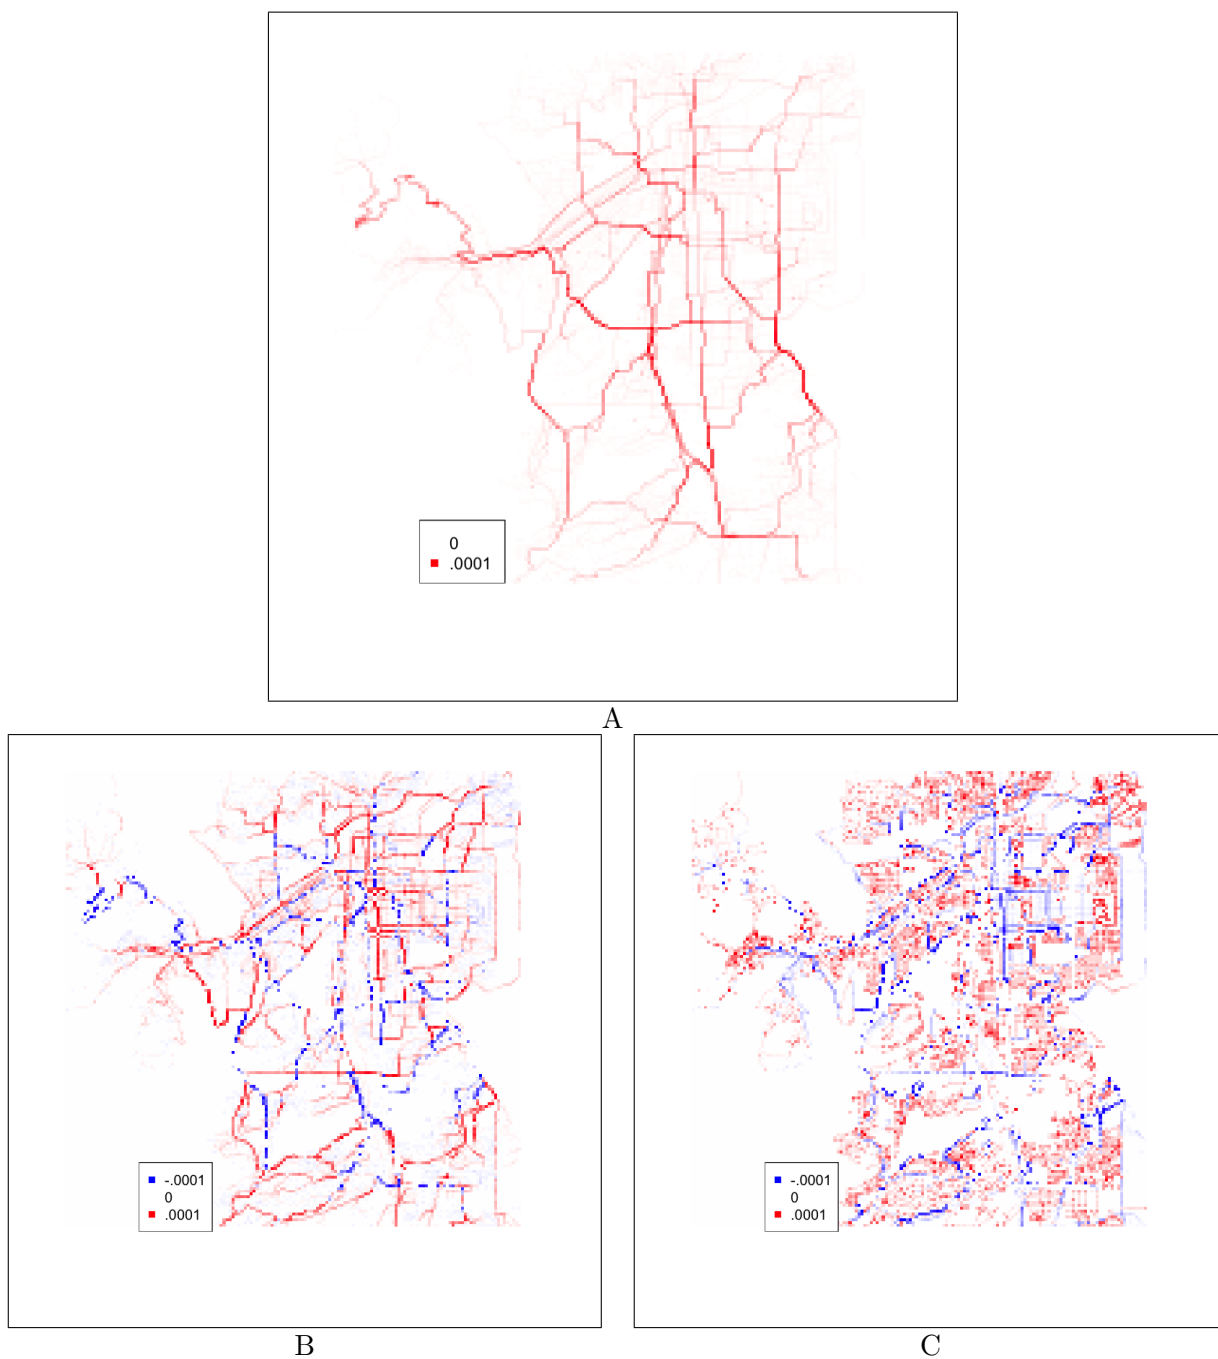





Figure 8.26: Comparison of heatmaps generated under the default parameters (A) and the parameter sweep runs with Max Speed set at 1000 m/5 min (B) and 8000 m/5 min (C).



A




B




C


Figure 8.27: Comparison of stress profiles generated under the default parameters (A) and the parameter sweep runs with Max Speed set at 1000 m/5 min (B) and 8000 m/5 min (C).



A



B



C

Figure 8.28: Comparison of speed profiles generated under the default parameters (A) and the parameter sweep runs with Max Speed set at 1000 m/5 min (B) and 8000 m/5 min (C).

8.4 Results

In order to compare the results generated by the model with the real-world evacuation patterns, it is necessary to consider reports of the situation as it truly developed in the city over the course of the simulation period - from midnight on the morning of June 23 through June 26. By reviewing the available information regarding real-world road usage during the crisis and comparing these with the speed and heatmap records produced by the simulation, it is possible to gain insight into the effectiveness of the model in capturing these patterns. The stress data can be compared against the information presented in Chapter 5 as a point of reference. These two avenues of comparison are explored in the remainder of this section.

8.4.1 Sentiment

Figure 8.1B shows the development of stress over the course of 30 different runs of the model with default parameters. There is extremely little variation between runs, and the general pattern shows that stress levels remain nonexistent until after the first mandatory evacuation efforts were ordered. After the first evacuation is ordered a little after 1pm, both the maximum and the average stress levels rise throughout the population. As time goes on, the fire remains unconstrained, and the conversation about the situation continues, stress continues to rise slightly, although the greatest influence is obviously the initial evacuation orders. By the end of the period, stress levels have effectively plateaued. This tracks well with the results of the stress detection analysis presented in Chapter 5; the wildfire triggers extreme emotions, but the overwhelming population response is one of relative emotional neutrality, with high-valence messages dropping off in the aftermath of the major evacuations. The generated stress levels therefore match reasonably well with the data.

8.4.2 Road Usage

The physical progress of the evacuation is another point of possible comparison, and a rich source of validation material. Of particular interest are reports of traffic backups, as these were highly likely to be documented by the media. Both in text and images,

local news organizations documented and reported on heavy traffic on Interstate 25 (I-25) and “all major eastbound roadways” (Udell, 2012). The Colorado Springs Final After Action Report (2013) notes that Pikes Peak Highway and 30th Street below Garden of the Gods Road were closed on June 23, with Highway 24 being closed on June 24. On June 26, westbound roads off of I-25 were closed to the public, and southbound I-25 was opened in both directions to traffic. The Denver Post captured photographs of heavy traffic on June 26 along Centennial Boulevard (Figure 8.30A), Garden of the Gods Road (Figure 8.30B), I-25 (Figure 8.30C), and Woodmen Road (Figure 8.30D), all contrasting with the relatively deserted neighborhoods shown in Figure 8.30E. Figure 8.29 maps traffic and closures against the overall road network, giving a sense of both the road closures and the areas of particularly heavy road usage. Because the images utilize the exchangeable image file format (Exif), it is possible to determine the date on which they were captured, allowing confidence in the validity of the texts associated with the posted images. By combining these reports from local news outlets such as CBS Denver (Hillan, 2012) and the Denver Post (2012a), it is possible to construct a sense of the movement of people, and to use this information to attempt to validate the results of the model.

In the real-world, the mass of evacuations took place on June 26, as documented in the images shown in Figure 8.30. In the simulation results shown in Figures 8.1A and C, however, Agents respond to the fire rapidly, essentially undertaking the full-scale evacuation effort on the night of the 23 rather than waiting for the evacuation orders of the 26. Because information spreads rapidly through the system, many individuals learn about the wildfire situation on the first day and judge it to be a threat worthy of evacuation. Thus, their threat evaluation metric leads them to respond to the threat more quickly than they did in reality, suggesting that in future work it would be interesting to explore precisely how individuals assess risk. From the response of the population to the threat in the model, patterns of road usage and speed appear similar between the model and the 26 of June. Generally speaking, the global speed profile indicates that after the first evacuation order goes out and everyone returns from their evening commute, they begin the evacuation process, with the first to

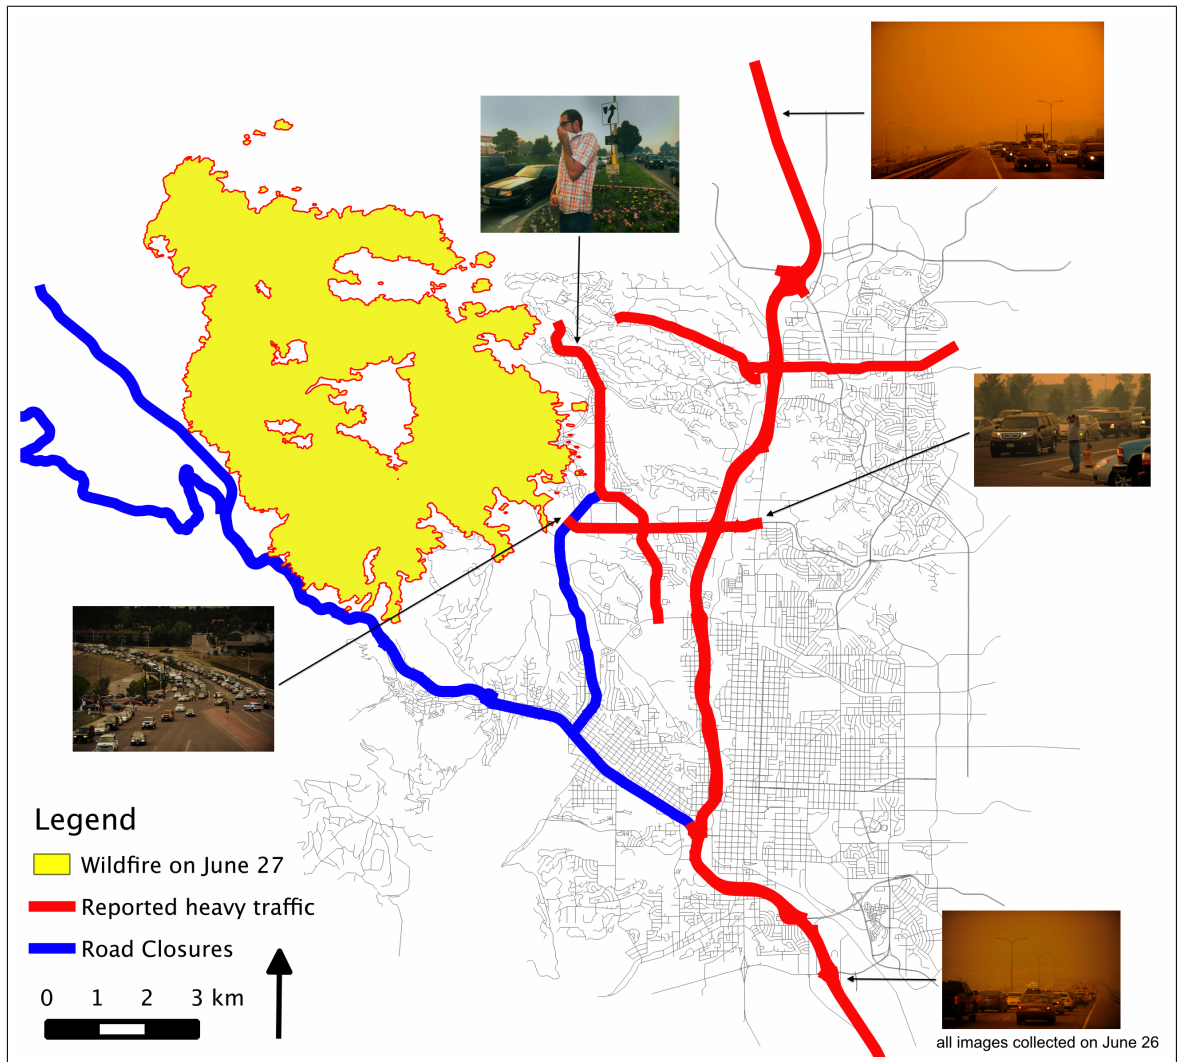


Figure 8.29: Visualization of the closed and heavy-traffic roads during the evacuation period. All images taken from the Denver Post (2012a); all road closure data taken from the Colorado Springs After Action Report (City of Colorado Springs, 2013)



A. Centennial Boulevard



B. Garden of the Gods Road



C. Northbound Interstate 25



D. Woodmen Road



E. A residential neighborhood

Figure 8.30: A collection of images taken on June 26 in Colorado Springs, documenting the relative traffic levels in various parts of the city. All images taken from the Denver Post (2012a)

evacuate interfering with the commuting patterns of those still on the roads. The traffic congestion continues throughout the next day, with the final evacuees leaving the local road network just as the evening commute the next day begins.

These efforts manifest in the road usage patterns in clear ways. Specifically, a number of roads which saw heavy usage in the wildfire scenario also appear in the set of congested roads in the real evacuation case - Woodmen Road and Garden of the Gods Road both show up in both the generated and the real sets, as shown in both Figures 8.1A and 8.29. The influence of the fire on the road usage is highlighted by Figure 8.2, where blue areas indicate that there was relatively less traffic under “default” circumstances. The comparison also highlights I-25 as an interesting case, being more heavily utilized for commuting by a larger population during the no-fire scenario but still briefly heavily utilized in the evacuation scenario. Other major roads in the evacuation areas appear to have been used more in the evacuation case than in the no-fire case, although they were not documented by major news organizations. As mentioned previously, the no-fire scenario generates far greater usage of Highway 24, even without needing to build the fact that the road was closed into the model. Overall, the road usage and speed patterns track well with the road usage patterns of June 26, suggesting that once prompted, the evacuation proceeds in a fashion that captures the relevant dynamics.

8.5 Summary

The influence of the various modules and parameters presented here support the conclusion that the simulation produces a satisfactory model of population movement in light of communication, observation, and the environment. The parameters interact as designed and expected, although the extremes of parameter settings can produce interesting system-level results. Based on this verification of the parameters and validation of the generated wildfire, Chapter 9 will present the results of the simulation itself, along with some thoughts regarding how the model might be employed in the future.

Chapter 9: Discussion

9.1 Overview

The work presented here represents a step toward utilizing new data sources and agent-based modeling in a humanitarian context in order to gain insight into how a crisis situation may develop. The thesis highlighted how VGI, AGI, and social media data in general can be synthesized with authoritative data sets in order to create information about the crisis and the affected population, as well as demonstrating how that information can support an agent-based model of a crisis scenario. These methodologies promise to fill informational gaps in knowledge of the location, attitudes, needs, and behaviors of crisis-affected populations, knowledge that responders sorely need. Given the focus of this work on developing methodologies which can be applied to areas with imperfect or limited data, the techniques discussed here are particularly applicable to situations in developing countries, a particularly important and underserved context.

Structurally, the thesis can be broken down into three overarching parts. The first introduced principles of complexity, human behavior, social network theory, and new data sources, reviewing the theory and existing applications of these lines of research. The second part put the theories introduced in the first into practice, synthesizing data sources and incorporating them into an agent-based model informed by real studies of behavior. Finally, the third portion presented the results generated by the processes developed in the second section, reviewing the success of these processes in capturing the dynamics described in the first section.

The introduction contextualized the importance of taking a complexity-based approach to the simulation of crisis situations, highlighting the importance of understanding human behavior and having accurate data in such research. Chapter 2 built upon this beginning,

exploring the research on behavior and the way it can be modeled. It further presented a range of models that explored crisis situations. These understandings of behavior and their operationalizations are explored relative to the capabilities of the various types of implemented simulations, highlighting the relative strengths and weaknesses of each of the methodologies. An overview of other agent-based models which address human behavior in high-stress situations was given, along with a survey of some agent-based approaches to specific processes implemented in the work here. Chapter 3 introduced new data sources which can be used to synthesize information, creating new information to support simulations. Adding to and expanding upon this line of research, Chapter 4 explored human social networks and how these shape interactions, especially in light of new forms of communication. The chapter presented an example of how information drawn from the kinds of platforms discussed in Chapter 3 can be combined with the information about social network structure and usage in order to gain insight into how humans are responding to and communicating about an emerging event. This example provides a bridge between the theoretical and the applied parts of the thesis.

The second section dealt with the methodologies directly employed in building the framework which is the core of this thesis. Chapter 5 reviewed how sentiment is understood by researchers and how it has been explored in the past before proposing a sentiment detector designed for usage in situations where a body of texts are poorly cleaned, possibly not in English, have little context, and must be rapidly processed. This sentiment detector was employed upon an extension of the example presented in Chapter 4, enriching the understanding of the situation and further highlighting the power of these combined techniques. Chapter 6 went further in reviewing existing methodologies for generating synthetic populations in order to create a synthetic population generation engine capable of dealing with similarly relatively uncleaned data and a need for rapid turn-around time. Fusing information from Chapters 3 and 4, the synthetic population generator created a group of individuals who resembled the target population at the aggregate statistical level, and who were connected to one another by intimate social networks and patterns of social media

interaction which match real-world network structures. This set of individuals was turned into units which populate the agent-based model of evacuation presented in Chapter 7. Chapter 7 detailed the construction of the model, highlighting how sources of information like those introduced in Chapter 3 underlie a model of a wildfire in a world populated by behavioral units taken from the population generated in Chapter 6 who interact with one another according to the dynamics described in Chapter 4 and Chapter 5. It reflected the final synthesis of all of these methodologies, the combination of authoritative as well as crowdsourced or ambient data sources with theoretical principles and behavioral patterns.

Based on the framework specified in the second part of the thesis, the third part carried out verification and validation efforts across a series of parameter sweeps, control cases, and the ultimate product generated by the model. Chapter 8 described each of these efforts in detail, reviewing the iterative process of construction and the way the interaction of different parts of the model were tested. The ultimate results of the full running model were compared with data drawn from real-life descriptions of emergent traffic and sentiment patterns, helping to assess the success of the model in capturing the real dynamics of the system. This chapter summarizes the entire process, reviewing what has been done and how it contributes to research overall.

9.2 Research Contributions

The model and supporting processes developed in this work represent a step toward an extensible framework for the simulation of crisis situations. The code comprising the simulation described in this thesis is available at www.css.gmu.edu/swise/thesis. By drawing upon a diverse range of data sources, the simulation shows how authoritative data can be synthesized with crowdsourced, volunteered, ambient, or synthesized data sources to gain a deeper understanding of the situation and to support simulations which incorporate spatiality, sentiment, decision-making, and heterogeneity in the population. To answer the research question posed in Chapter 1, it is empirically possible to synthesize data from a wide range of sources within an agent-based model, and thereby to project how a crisis

situation might evolve.

The thesis contributes a number of developments across a range of fields, especially agent-based modeling, computational social science, and geography. The construction of a spatially explicit ABM with realistic social networks utilizing new sources of information represents a hybrid contribution to all of these fields. By combining realistic agent behaviors with the otherwise static contributions of GIS, it is possible to explore why individuals might make decisions in different contexts; by combining VGI and AGI with traditional sources of data, new kinds of information can be used to support and validate these behaviors. Computational simulation represents a powerful way of combining the respective strengths of each of these fields, and this model pushes the boundaries of existing research in each accordingly.

To address the contributions more specifically, agent-based modeling is advanced by the creation of synthetic population generation methods which can create the kinds of data they increasingly require. The process of synthetic population generation usually does not include the creation of a social network, let alone the establishment of a network of social media connections, making the method presented here an important contribution. Agent-based models are often paired with explicit spatial data, but rarely with social networks. By ensuring that the structure of the information network allowed individuals to communicate with one another based on these important features of their social networks, a level of realism that is usually lacking from disaster simulations is introduced to the model. Given the importance of information to the behavior of individuals in crisis situations described in Chapter 2, ensuring a realistic pattern of communication and flow of information is a crucial contribution of the model.

Further, computational social science in general benefits by the designing of the population synthesis process around a highly generalizable structure. The model anticipates and allows for data of varying formats and qualities to be incorporated into the population generation process. Given the paucity of good, validated data capable of supporting a simulation of the magnitude presented in Chapter 7 - and especially the unlikelihood of

acquiring such data for historically marginalized areas such as Haiti or rural Pakistan - the combination of these tools allows for the simulation to be much more broadly applicable. By combining information from across a range of data sources as in Chapters 4 and 5, these gaps in information can be patched and projections run with reasonable confidence that the data reflect the current situation on the ground. These manipulations of data into new formats contributes to geography in interesting ways as well, advancing the boundaries of understanding of the intersection of geographic and social phenomena.

Finally, by combining these rich data sources with meaningful social networks, the model has the ability to include meaningful decision-making processes for the individual agents. Agents are embedded not only in physical space but in social space, informed by their heterogeneous characteristics and informed by their own personal social networks. Not only can agents therefore be designed with meaningful decision-making methods, but the data sources can both validate and inform the decision agents make. By comparing the activities of agents with real-world, individual-level data, the quality of the model's decision trees can be validated; by training the weights agents assign to different options on that real-world data, the model can be optimized to better project the true decisions in the future. Over time, the quality of the projections will only improve, making the tool better and better.

9.3 Limitations

It has been noted that researchers have an excellent model of the real-world in the world itself. Incorporating all of the complexity of a situation into the situation necessarily requires that aspects of the world be omitted or simplified, by the nature of simulation. Miller and Page (2007) argue that any simulation must leave out aspects of reality, just as any map must leave out unnecessary details in favor of capturing key features of the world that pertain to the map's intended use. Understanding the limitations and constraints of a model, therefore, is an important aspect of model design and usage.

Perhaps the most obvious limitation of agent-based modeling is the way detail is handled. While a model could conceivably incorporate any range of agent attributes, any number

of agents, any suite of behaviors, and do so at any combination of spatial and temporal scales, the realities of information availability and hardware constrain the feasibility of some efforts. The portion of this thesis dedicated to synthesizing information sources to support simulations does a great deal to alleviate questions of information availability, and the agent behaviors are based on as much firm evidence as is available at the time of writing. Toward the end of addressing constraints in terms of modeled behavioral complexity, it is necessary to understand behaviors and attributes as abstractions, and to make simplifying assumptions about how individuals perceive, communicate, feel, judge risk, and plan. Part of the goal of modeling is to ensure that the abstractions of these processes produce a signal based on the stimuli to which they are subject which resembles the real-world responses of individuals; constructing these diverse but mutually interactive processes is arguably the central challenge of agent-based modeling.

Given the complexity of the code involved in creating this model, the potential exists for errors to exist despite the extensive unit testing and verification efforts specified in Chapter 8. Further, running the simulation requires a reasonable quantity of memory - roughly 30G of RAM - which is frequently not available on regular computing hardware. Technology therefore does limit the number of agents that can be simulated in a reasonable amount of time, although Moore's law regarding the ever-increasing processing power of computers suggests that this constraint will ease as time goes on. In further technical issues, the way geographic information is transformed between projections by converting and resampling the scale of certain datasets can introduce mild distortions to the environmental data which underlies the wildfire model. These errors, however minor they may be, have been minimized to the greatest extent possible; regardless, caution should be taken if the framework is extended to deal with coarser resolutions or more extreme latitudes. Further, the absence of precise wind information limits the effectiveness of the wildfire model in its current implementation.

9.4 Further Work

A wide range of further work is suggested by the material presented here. Because the focus of the work has been the exploration of how a variety of tools can be combined to better explore structures and systems, nearly all of the individual processes could be profitably expanded upon and explored in greater detail. In certain cases, it would be interesting to add or interchange aspects of the system - rather than focusing on Twitter, it would be fascinating to utilize information about communication via Facebook, if the data were available. It would also be interesting to modify the risk assessment mechanism of the agents so that the evacuation proceeded on a timetable more in keeping with the real-world evacuation temporal patterns. More generally, exploring interaction through other social media platforms would be an interesting addition to the research done here. That being said, it could be profitable to further investigate the structure and influence of Twitter itself - fine-tuning the local social media networks of the individuals to include different sources of authority, or to capture the variation in the way different kinds of information spread, would enrich the communication network, and adopting personalized or attribute-dependent levels of interaction with social network would improve the way information flowed. Including variation in the rate of social media usage based on interest would be another possible extension. Finally, intentionally modeling subcommunities or relatively locally isolated persons with many connections outside the local network could give further insight into how evacuation orders or other information could be targeted to groups who might otherwise be isolated from the social network.

Leaving social networks aside, the process of sentiment detection is a rich field with constant developments. In the interest of studying and automatically identifying sentiment-bearing social media posts, further exploring sentiment detection methods could be extremely useful. Especially in the context of Twitter usage, the inclusion of amplifiers, explicit language, and a wider range of emoticons could potentially improve substantially on the existing method. Experimentation with different stemmers and lexicons might also prove fruitful. Within the context of the model, simulating emotions other than stress could

add an interesting dimension to the decision-making process, as would allowing agents to influence one another emotionally without actually communicating information.

The population synthesis process, because of its many steps and complexity, could be expanded upon in a number of ways. Validation studies comparing the generated houses against the United States Census Bureau’s Public Use Microdata Areas data set (2012) and the generated population against the RTI United States Synthesized Population Database (Wheaton et al., 2009) would give more insight into the quality of the generated population, potentially indicating correlated phenomena that are visible at the level of the distribution but not the aggregate statistics. The simple way in which workplaces are generated could be replaced by data drawn from the Yellow Pages, or perhaps from sources of volunteered information. Exploring more nuanced family generation methods, and especially tailoring these methods to be easily extensible in areas with very different family structures, could also play into different methods for partitioning the family into agents. Different breakdowns might allow for other agent activity processes - children could attend school, adding a further wrinkle to agent evacuation planning. Incorporating other kinds of housing structures such as group housing into the simulation would add similarly to the nuance of evacuation, as retirement homes and prisons tend to be evacuated in very different ways. Even without added these new structures, however, including a socioeconomic attribute into the agent’s design might extensively impact decision-making, as it could influence the location of agent’s homes and workplaces, the likelihood with which they participate in certain social media platforms, the hours during which they work and sleep, and the ease with which they can forgo a day of work in order to prepare for an evacuation.

Delving further into the behaviors of agents, incorporating richer inter-family behaviors into the process of evacuation would allow for the model to capture dynamics like families converging at home to pick up the family dog and consolidate into one car. The number of cars on the road could be significantly influenced by dynamics like these, as well as by richer activity scheduling in general. Exploring the ways in which agents carry out their pathfinding efforts could both improve upon the realism of individual wayfinding behavior

and substantially speed up the simulation, as pathfinding is one of the most expensive operations involved in the simulation. Not only the way that agents travel but the way they select their destinations could also be expanded upon: a richer evacuation target selection method, perhaps informed by the previously suggested socioeconomic attribute, might influence the proportion of individuals who find themselves on any given highway. Such a subprocess would be well-served by allowing agents to predict how a wildfire might develop, permitting richer decision-making methods. Along those lines, the wildfire model itself could be made to incorporate other features, or improved upon in a variety of ways: the work of Coen and Schroeder (2013) or Kochanski et al. (2013) on predicting wildfire development in near real-time suggests an interesting direction.

Finally, exploring other case studies could give interesting insights into how future evacuations could be designed. What might have happened if certain roads had been closed or restricted to traffic flow in one direction, and how would the timing of those traffic-control measures have impacted the evacuation? Evacuation orders issued at different times of day and over different sources of media might prompt different responses, and might influence individuals very differently depending on the local sentiment profile: as shown by Hurricane Katrina, an evacuation order which comes too late to a population largely complacent in the face of such threats does not prompt rapid responses (see Simerman et al., 2005). Further, what influence might misinformation, injected into the system either deliberately or accidentally, have on the population's response to the situation? The popularity of an image of a shark swimming around submerged streets in Manhattan in the aftermath of Hurricane Sandy, and the widespread acceptance of this image as credible, suggests that responders will have to deal with incorrect information in the system in the future (Gupta et al., 2013). Any of these modifications would represent an interesting new avenue of exploration, and the variety of options suggests that the work has a great deal more to offer.

As the framework developed here exists now, it is suitable for the exploration of crisis situations and can support efforts to understand the dynamics present under different scenarios. It can therefore serve as a tool to develop possible worlds for researchers and

planners to consider when assessing likely outcomes. Perhaps the most exciting direction in which the framework could move would be toward forecasting of a situation, hopefully to the point of being about to present responders with short-range projections of the likely development of a crisis situation given a variety of possible interventions. This obviously represents the most ambitious suggestion, but could significantly change the way crisis response is handled.

9.5 Conclusion

This thesis has sought to explore how diverse data sources could be synthesized to support agent-based models of complex and varying situations toward the end of projecting the development of crisis situations. By reviewing emerging sources of data and how they compare with other, more authoritative sources of information, the thesis advances the field of research into knowledge extraction. As these sources of information are used to inform richer and more powerful models of complex systems, agent-based modeling is advanced as well. The successes of the model in projecting meaningful dynamics suggests that this is a valuable line of research upon which to continue, and this chapter has discussed a number of logical extensions upon the existing work.

References

- Aaby, K., Abbey, R. L., Herrmann, J. W., Treadwell, M., Jordan, C. S., & Wood, K. (2006). Embracing Computer Modeling to Address Pandemic Influenza in the 21st Century. *Journal of Public Health Management and Practice*, 12(4), 365–72.
- Achrekar, H., Gandhe, A., Lazarus, R., Yu, S., & Liu, B. (2011). Predicting flu trends using Twitter data. In *holder2011 Conference on Computer Communications Workshops* (p. 702-707). IEEE.
- Afshar, A. M., & Haghani, A. (2008). Heuristic Framework for Optimizing Hurricane Evacuation Operations. *Transportation Research Record: Journal of the Transportation Research Board*, 2089(1), 9–17.
- Albert, R., & Barabási, A.-L. (2002). Statistical Mechanics of Complex Networks. *Reviews of Modern Physics*, 74(1), 47–97.
- Albinet, G., Searby, G., & Stauffer, D. (1986). Fire Propagation in a 2-D Random Medium. *Journal de Physique*, 47(1), 1–7.
- Albini, F. A. (1979). *Spot fire distance from burning trees: A predictive model*. Inter-mountain Forest and Range Experiment Station, Forest Service, US Department of Agriculture.
- Alexandridis, A., Vakalis, D., Siettos, C., & Bafas, G. (2008). A Cellular Automata Model for Forest Fire Spread Prediction: The Case of the Wildfire that Swept Through Spetses Island in 1990. *Applied Mathematics and Computation*, 204(1), 191–201.
- American Association of State Highway and Transportation Officials. (2013). *Commuting in America 2013* (Tech. Rep.). Retrieved from <http://traveltrends.transportation.org/Documents/CA11-4.pdf>

- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, 111(4), 1036.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Psychology Press.
- Anderson, P. (2007). What is Web 2.0? Ideas, Technologies and Implications for Education. *Horizon Scanning Report*. Retrieved from <http://www.jisc.ac.uk/whatwedo/services/techwatch/reports/horizonsscanning/hs0701.aspx>
- Andrews, P., Bevins, C., & Seli, R. (2005). *BehavePlus Fire Modeling System, Version 3.0* (Tech. Rep.). Ogden, UT: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station.
- Antweiler, W., & Frank, M. Z. (2004). *Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards* (Tech. Rep.). Sauder School of Business. Retrieved from <http://www.ssrn.com/abstract=282320>
- Aptekar, L. (1990). A Comparison of the Bicoastal Disasters of 1989. *Cross-Cultural Research*, 24(1-4), 73–104.
- Arentze, T., Timmermans, H., & Hofman, F. (2007). Creating Synthetic Household Populations: Problems and Approach. *Transportation Research Record*, 2014(1), 85–91.
- Autant-Bernard, C., Billand, P., Frachisse, D., & Massard, N. (2007). Social Distance versus Spatial Distance in R&D cooperation: Empirical Evidence from European Collaboration Choices in Micro and Nanotechnologies. *Papers in Regional Science*, 86(3), 495–519.
- Axelrod, R. (1997). Advancing the Art of Simulation in the Social Sciences. In R. Conte, R. Hegselmann, & P. Terno (Eds.), *Simulating Social Phenomena* (pp. 21–40). Berlin, Germany: Springer.
- Axelrod, R. (2007). Simulation in the Social Sciences. In J.-P. Reynard (Ed.), *Handbook of Research on Nature Inspired Computing for Economy and Management* (pp. 90–100). Hershey, CA: Idea Group.

- Axtell, R. L. (2000). *Why Agents? On the Varied Motivations for Agent Computing in the Social Sciences* (Tech. Rep. No. 17). Washington, D.C.: Center on Social and Economic Dynamics (The Brookings Institute).
- Axtell, R. L., & Epstein, J. M. (1994). Understanding Our Creations. *Bulletin of the Santa Fe Institute, Winter*, 28–32.
- Bailey, N. T. J. (1975). *The Mathematical Theory of Infectious Diseases and its Applications* (2nd ed.). London: Charles Griffin & Co.
- Bakshy, E., Hofman, J. M., Watts, D. J., & Mason, W. A. (2011). Identifying ‘Influencers’ on Twitter. In *Fourth ACM International Conference on Web Search and Data Mining WSDM*. Hong Kong, China.
- Balbi, J. H., Morandini, F., Silvani, X., Filippi, J.-B., & Rinieri, F. (2009). A Physical Model for Wildland Fires. *Combustion and Flame*, 156(12), 2217–2230.
- Balcik, B., & Beamon, B. M. (2008). Facility location in humanitarian relief. *International Journal of Logistics*, 11(2), 101–121.
- Ballas, D., Clarke, G., Dorling, D., Eyre, H., Thomas, B., & Rossiter, D. (2005). SimBritain: a Spatial Microsimulation Approach to Population Dynamics. *Population, Space and Place*, 11(1), 13–34.
- Balog, K., Mishne, G., & De Rijke, M. (2006). Why are they excited?: Identifying and explaining spikes in blog mood levels. In *holderProceedings of the Eleventh Conference of the European Chapter of the Association for Computational Linguistics: Posters & Demonstrations* (p. 207–210). Association for Computational Linguistics.
- Banos, A., Godara, A., & Lassarre, S. (2005). Simulating Pedestrians and Cars Behaviours in a Virtual City: an Agent-Based Approach. In *Proceedings of the European Conference on Complex Systems* (pp. 2–5). Paris, France.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439), 509–512. (PMID: 10521342)
- Barbarosoglu, G., & Arda, Y. (2004). A Two-Stage Stochastic Programming Framework

- for Transportation Planning in Disaster Response. *Journal of the Operations Research Society*, 55, 43-53.
- Barbosa, L., & Feng, J. (2010). Robust Sentiment Detection on Twitter from Biased and Noisy Data. In *Proceedings of the International Conference on Computational Linguistics* (pp. 36-44). Beijing, China.
- Barthelemy, J., & Toint, P. L. (2012). Synthetic Population Generation Without a Sample. *Transportation Science*, 47(2), 266-279.
- Batty, M. (2003). *Agent-Based Pedestrian Modelling* (Tech. Rep. No. 61). London, UK: Centre for Advanced Spatial Analysis (University College London).
- Batty, M., Desyllas, J., & Duxbury, E. (2003). Safety in Numbers? Modelling Crowds and Designing Control for the Notting Hill Carnival. *Urban Studies*, 40(8), 1573-1590.
- Batty, M., Hudson-Smith, A., Milton, R., & Crooks, A. T. (2010). Map Mashups, Web 2.0 and the GIS Revolution. *Annals of GIS*, 16(1), 1-13.
- Batty, M., Jiang, B., & Thurstain-Goodwin, M. (1998). Local movement: Agent-based models of pedestrian flows. *Centre for Advanced Spatial Analysis, University College, London*.
- Beckman, R. J., Baggerly, K. A., & McKay, M. D. (1996). Creating Synthetic Baseline Populations. *Transportation Research Part A: Policy and Practice*, 30(6), 415-429.
- Berg, P., Arentze, T., & Timmermans, H. (2010). A Multilevel Path Analysis of Contact Frequency Between Social Network Members. *Journal of Geographical Systems*, 14(2), 125-141.
- Berger, J., & Iyengar, R. (2013). *How Interest Shapes Word-of-Mouth Over Different Channels* (Tech. Rep.). University of Pennsylvania.
- Berger, J., & Milkman, K. L. (2012a). *Social Transmission, Emotion and Virality of Online Content* (Tech. Rep.). University of Pennsylvania.
- Berger, J., & Milkman, K. L. (2012b). What Makes Online Content Viral? *Journal of Marketing Research*, 49(2), 192-205.

- Berjak, S. G., & Hearne, J. W. (2002). An Improved Cellular Automaton Model for Simulating Fire in a Spatially Heterogeneous Savanna System. *Ecological Modelling*, 148(2), 133–151.
- Biewald, L., & Janah, L. (2010). TechCrunch: Crowdsourcing Disaster Relief. *TechCrunch*.
- Birkin, M., & Clarke, M. (1988). SYNTHESIS - A Synthetic Spatial Information System for Urban and Regional Analysis: Methods and Examples. *Environment and Planning A*, 20(12), 1645–1671.
- Birkin, M., & Wu, B. M. (2012). A Review of Microsimulation and Hybrid Agent-Based Approaches. In A. J. Heppenstall, A. T. Crooks, L. M. See, & M. Batty (Eds.), *Agent-based Models of Geographical Systems* (pp. 51–68). New York, NY: Springer.
- Blendon, R. J., Benson, J. M., DesRoches, C. M., Lyon-Daniel, K., Mitchell, E. W., & Pollard, W. E. (2007). The Public's Preparedness for Hurricanes in Four Affected Regions. *Public Health Reports*, 122, 167–176.
- Blood, D. J., & Phillips, P. C. B. (1995). Recession Headline News, Consumer Sentiment, the State of the Economy and Presidential Popularity: a Time Series Analysis 1989-1993. *International Journal of Public Opinion Research*, 7(1), 2–22.
- Bollen, J., Mao, H., & Zeng, X. (2011). Twitter Mood Predicts the Stock Market. *Journal of Computational Science*, 2(1), 1–8.
- Bollobás, B. (2003). Mathematical results on scale-free random graphs. In S. Bornholdt & H. G. Schuster (Eds.), *Handbook of Graphs and Networks* (pp. 1–38). Weinheim, FRG: Wiley-VCH Verlag GmbH & Co. KGaA.
- Bremmer, I. (2010). Democracy in Cyberspace: What Information Technology Can and Cannot Do. *Foreign Affairs*, 89(6), 86–92.
- Brouwers, L. (2005a). MicroPox: A Large-Scale and Spatially Explicit Microsimulation Model for Smallpox Planning. In *International Conference on Health Sciences Simulation* (pp. 70–76). New Orleans, Louisiana.
- Brouwers, L. (2005b). *Microsimulation Models for Disaster Policy Making*. Unpublished

- doctoral dissertation, Stockholm University.
- Brown, D. G. (2006). Agent-based Models. In H. Geist (Ed.), *The Earth's Changing Land: An Encyclopedia Of Land-Use And Land-Cover Change* (pp. 7–13). Westport, CT: Greenwood Publishing Group.
- Brownstein, J. S., Freifeld, C. C., Reis, B. Y., & Mandl, K. D. (2008). Surveillance Sans Frontieres: Internet-Based: Internet-Based Emerging Infectious Disease Intelligence and the Health-Map Project. *PLoS Medicine*, 5(7), 1019–1024.
- Bureau of Labor Statistics. (2012). *American Time Use Survey 2012: Civilian Noninstitutional Population Age 15 and Over* (Tech. Rep.). Bureau of Labor Statistics, United States Department of Labor. Retrieved from http://www.bls.gov/tus/tables/a1_2012.pdf
- Burt, R. S. (1984). Network Items and the General Social Survey. *Social Networks*, 6(4), 293–339.
- Caragea, C., Mcneese, N., Jaiswal, A., Traylor, G., Kim, H.-w., Mitra, P., ... Yen, J. (2011). Classifying Text Messages for the Haiti Earthquake. In *Proceedings of the 8th International International Conference on Information Systems for Crisis Response and Management Conference* (pp. 1–10). Lisbon, Portugal.
- Carey, J. (2002). Media Use During a Crisis. *Prometheus*, 20(3), 201–207.
- Casti, J. L. (1997). *Would-Be Worlds: How Simulation is Changing the Frontiers of Science*. New York, NY: John Wiley and Sons.
- Castle, C., & Longley, P. (2008). Emergency Pedestrian Evacuation Analysis: A Review and Interpretation of Software Applications. In D. Sui & S. Cutter (Eds.), *Geospatial Technologies and Homeland Security: Research Frontiers and Challenges* (pp. 209–228). Berlin, Germany: Springer.
- Castle, C. J. (2007). *Guidelines for Assessing Pedestrian Evacuation Software Applications* (Tech. Rep. No. 115). London, UK: Centre for Advanced Spatial Analysis (University College London).
- Cha, M., Mislove, A., Adams, B., & Gummadi, K. P. (2008). Characterizing Social Cascades

- in Flickr. *Proceedings of the first workshop on Online social networks*, 13.
- Cha, M. C., Benevenuto, F., Haddadi, H., & Gummadi, K. P. (2010). Measuring User Influence in Twitter: the Million Follower Fallacy. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*. Washington, DC.
- Chen, F., Zhai, Z., & Madey, G. (2010). Dynamic Adaptive Disaster Simulation: Developing a Predictive Model of Emergency Behavior Using Cell Phone and GIS Data. In *2010 Computational Social Science Society Conference*. Phoenix, AZ.
- Chen, X. (2008). Microsimulation of Hurricane Evacuation Strategies of Galveston Island. *The Professional Geographer*, 60(2), 160–173.
- City of Colorado Springs. (2013). *Waldo Canyon Fire* (Tech. Rep.). Colorado Springs, CO: City of Colorado Springs.
- CloudMade Downloads*. (2012). Retrieved Oct 23, 2012, from <http://cloudmade.com/downloads>
- Coen, J. L., & Schroeder, W. (2013). Use of Spatially Refined Satellite Remote Sensing Fire Detection Data to Initialize and Evaluate Coupled Weather-wildfire Growth Model Simulations. *Geophysical Research Letters*, 40(20), 5536–5541.
- Cohn, R., Carley, K. M., Harrauld, J. R., & Wallace, W. A. (2000). Emotions in Crisis Management: an Analysis of the Organisational Response of Two Natural Disasters. *International Journal of Technology Management*, 19(3), 313–335.
- Cole, J. W., Sabel, C. E., Blumenthal, E., Finnis, K., Dantas, A., Barnard, S., & Johnston, D. M. (2005). GIS-Based Emergency and Evacuation Planning for Volcanic Hazards in New Zealand. *Bulletin of the New Zealand Society for Earthquake Engineering*, 38(3), 149–164.
- Corley, C. D., Cook, D. J., Mikler, A. R., & Singh, K. P. (2010). Text and Structural Data Mining of Influenza Mentions in Web and Social Media. *International Journal of Environmental Research and Public Health*, 7(2), 596–615.
- Cormack, G. V., Gómez Hidalgo, J. M., & Sáenz, E. P. (2007). Spam Filtering for Short Messages. In *Proceedings of the Sixteenth ACM Conference on Information and Knowledge*

- Management* (p. 313). New York, New York, USA: ACM Press.
- Cosgrave, J. (1996). Decision Making in Emergencies. *Disaster Prevention and Management*, 5(4), 28–35.
- Cova, T. J., & Church, R. L. (1997). Modelling Community Evacuation Vulnerability Using GIS. *International Journal of Geographical Information Science*, 11(8), 763–784.
- Cova, T. J., & Johnson, J. P. (2003). A Network Flow Model for Lane-based Evacuation Routing. *Transportation Research Part A: Policy and Practice*, 37(7), 579–604.
- Crane, R., & Sornette, D. (2008). Robust Dynamic Classes Revealed by Measuring the Response Function of a Social System. *Proceedings of the National Academy of Sciences of the United States of America*, 105(41), 15649–53.
- Croitoru, A., Crooks, A., Radzikowski, J., & Stefanidis, A. (2013). GeoSocial Gauge: A System Prototype for Knowledge Discovery from Social Media. *International Journal of Geographical Information Science*, 27(12), 2483–2508.
- Croitoru, A., Crooks, A. T., Radzikowski, J., Stefanidis, A., Vatsavai, R. R., & Wayant, N. (2014). Geoinformatics and Social Media: A New Big Data Challenge. In H. Karimi (Ed.), *Big Data Techniques and Technologies in Geoinformatics* (pp. 207–232). Boca Raton, FL: CRC Press.
- Croitoru, A., Stefanidis, A., Radzikowski, J., Crooks, A. T., Stahl, J., & Wayant, N. (2012). Towards a Collaborative GeoSocial Analysis Workbench. In *COM-Geo*. Washington, DC.
- Crooks, A. T., Castle, C., & Batty, M. (2008). Key Challenges in Agent-based Modelling for Geo-spatial Simulation. *Computers, Environment and Urban Systems*, 32(6), 417–430.
- Crooks, A. T., Croitoru, A., Stefanidis, A., & Radzikowski, J. (2012). #Earthquake: Twitter as a Distributed Sensor System. *Transactions in GIS*, 17(1), 124–147.
- Crooks, A. T., & Heppenstall, A. J. (2012). Agent-Based Models of Geographical Systems. In A. J. Heppenstall, A. T. Crooks, L. M. See, & M. Batty (Eds.), *Agent-based Models of Geographical Systems* (Vol. 164, pp. 85–108). New York, NY: Springer.

- Crooks, A. T., Hudson-Smith, A., Croitoru, A., & Stefanidis, A. (2014). The Evolving GeoWeb. In R. Abraham & L. See (Eds.), *Geocomputation* (2nd ed.). CRC Press.
- Crooks, A. T., & Wise, S. (2013). GIS and agent-based models for humanitarian assistance. *Computers, Environment and Urban Systems*, 41, 100-111.
- Cuervo, R., Diaz, F., Namen, I., Palacio, C., & Sierra, C. (2010). Humanitarian Crisis: When Supply Chains Really Matter. In *International Conference of the System Dynamics Society* (pp. 1-36). Seoul, Korea: System Dynamics Society.
- Danielsson, M., & Ohlsson, K. (1999). Decision Making in Emergency Management: A Survey Study. *International Journal of Cognitive Ergonomics*, 3(2), 91-99.
- Das, S. R., & Chen, M. Y. (2007). Yahoo! for Amazon: Sentiment Extraction from Small Talk on the Web. *Management Science*, 53(9), 1375-1388.
- Dave, K., Lawrence, S., & Pennock, D. M. (2003). Mining the peanut gallery: Opinion extraction and semantic classification of product reviews. In *Proceedings Of The 12th International Conference On World Wide Web* (pp. 519-528). ACM.
- Davidov, D., & Tsur, O. (2010). Enhanced Sentiment Learning Using Twitter Hashtags and Smileys. In *Proceedings of the International Conference on Computational Linguistics*.
- Dawson, J. (1974). Suffix removal for word conflation. *Bulletin of the Association for Literary & Linguistic Computing*, 2(3), 33-46.
- Dawson, R. J., Peppe, R., & Wang, M. (2011). An Agent-based Model for Risk-based Flood Incident Management. *Natural Hazards*, 59(1), 167-189.
- De Zwart, M., Lindsay, D., Henderson, M., & Phillips, M. (2011). Randoms vs Weirdos: Teen Use of Social Networking Sites and Perceptions of Legal Risk. *Alternative Law Journal*, 36(3), 153-157.
- Deming, W. E., & Stephan, F. F. (1940). On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known. *Annals of Mathematical Statistics*, 11(4), 427-444.

- Denver Post. (2012a). *Colorado Wildfires Photos: The Waldo Canyon Fire in Colorado Springs*. Denver, CO. Retrieved February 26, 2014, from <http://blogs.denverpost.com/captured/2012/06/27/colorado-wildfires-waldo-canyon-fire-colorado-springs/5732/>
- Denver Post. (2012b). *Waldo Canyon Wildfire KML Files*. Retrieved November 8, 2012, from <http://extras.denverpost.com/media/maps/kml/>
- de Silva, F. N. (2001). Providing Spatial Decision Support for Evacuation Planning: a Challenge in Integrating Technologies. *Disaster Prevention and Management*, 10(1), 11–20.
- de Silva, F. N., & Eglese, R. W. (2000). Integrating Simulation Modelling and GIS: Spatial Decision Support Systems for Evacuation Planning. *The Journal of the Operational Research Society*, 51(4), 423–430.
- Ding, X., & Liu, B. (2007). The utility of linguistic rules in opinion mining. In *holder- Proceedings Of The 30Th Annual International ACM SIGIR Conference On Research And Development In Information Retrieval* (p. 811-812). ACM.
- Ding, X., Liu, B., & Yu, P. S. (2008). A holistic lexicon-based approach to opinion mining. In *Proceedings of the 2008 International Conference on Web Search and Data Mining* (p. 231-240). ACM.
- Divol, R., Edelman, D., & Sarrazin, H. (2012). Demystifying Social Media. *McKinsey Quarterly*, 1–11.
- Dodds, P. S., Harris, K. D., Kloumann, I. M., Bliss, C. a., & Danforth, C. M. (2011). Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter. *PloS one*, 6(12), e26752.
- Dombroski, M., Fischhoff, B., & Fischbeck, P. (2006). Predicting Emergency Evacuation and Sheltering Behavior: a Structured Analytical Approach. *Journal of Risk Analysis*, 26, 1675–1688.
- Drabek, T. E. (1986). *Human System Responses to Disaster: An Inventory of Sociological Findings*. New York, NY: Springer Verlag.

- Drabek, T. E. (1992). Variations in Disaster Evacuation Behavior: Public Responses Versus Private Sector Executive Decision-making Processes. *Disasters*, 16(2), 104–118.
- Driskell, J. E., & Salas, E. (1991). Group Decision Making Under Stress. *Journal of Applied Psychology*, 76(3), 473–478.
- Drury, J., Cocking, C., Reicher, S., Burton, A., Schofield, D., Hardwick, A., ... Langston, P. (2009). Cooperation Versus Competition in a Mass Emergency Evacuation: a New Laboratory Simulation and a New Theoretical Model. *Behavior Research Methods*, 41(3), 957–70.
- Duggan, M., & Brenner, J. (2013). *The Demographics of Social Media Users - 2012* (Tech. Rep.). Pew Research Center. Retrieved from http://pewinternet.org//media//Files/Reports/2013/PIP_SocialMediaUsers.pdf
- Duggan, M., & Rainie, L. (2012). *Cell Phone Ownership Hits 91% of Adults* (Tech. Rep.). Washington, DC: Pew Research Center. Retrieved from <http://www.pewresearch.org/fact-tank/2013/06/06/cell-phone-ownership-hits-91-of-adults/>
- Dunbar, R. (1992). Neocortex Size as a Constraint on Group Size in Primates. *Journal of Human Evolution*, 22(6), 469–493.
- Dunbar, R. (1993). Co-evolution of Neocortex Size, Group Size and Language in Humans. *Behavioral and Brain Sciences*, 16, 681–735.
- Dunbar, R. (2011). *The human story*. Faber & Faber.
- Durrett, R. (2007). *Random Graph Dynamics*. Cambridge University Press.
- Easley, D., & Kleinberg, J. (2010). Information Cascades. In *Networks, Crowds and Markets: Reasoning about a Highly Connected World* (chap. 16). Cambridge University Press.
- Elmitiny, N., Ramasamy, S., & Radwan, E. (2007). Emergency Evacuation Planning and Preparedness of Transit Facilities: Traffic Simulation Modeling. *Transportation Research Record*, 1992(1), 121–126.

- Epstein, J., & Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. Brookings Institution Press.
- Epstein, J. M. (2009). Modelling to Contain Pandemics. *Nature*, 460(7256), 687.
- Epstein, J. M., Pankajakshan, R., & Hammond, R. A. (2011). Combining Computational Fluid Dynamics and Agent-Based Modeling: a new approach to evacuation planning. *PloS one*, 6(5), e20139.
- Erdős, P., & Rényi, A. (1959). On Random Graphs. *Publicationes Mathematicae Debrecen*, 6, 290–297.
- Eubank, S. G., Guclu, H., Kumar, V. S. A., & Marathe, M. V. (2004). Modelling Disease Outbreaks in Realistic Urban Social Networks. *Nature*, 429(6988), 180–184.
- Farooq, B., Hurtubia, R., Bierlaire, M., & Flötteröd, G. (2013). *Simulation-based Synthesis of Population* (Tech. Rep.). Transport and Mobility Laboratory, Ecole Polytechnique Fédérale de Lausanne.
- Fawcett, W., & Oliveira, C. S. (2000). Casualty Treatment After Earthquake Disasters: Development of a Regional Simulation Model. *Disasters*, 24(3), 271–87.
- Feng, T., & Keller, L. R. (2006). A multiple-objective decision analysis for terrorism protection: Potassium iodide distribution in nuclear incidents. *Decision Analysis*, 3(2), 76–93.
- Filippi, J.-B., Bosseur, F., Pialat, X., Santoni, P.-A., Strada, S., & Mari, C. (2011). Simulation of Coupled Fire/Atmosphere Interaction with the MesoNH-ForeFire Models. *Journal of Combustion*, 2011, 1–13.
- Fink, C., Kopecky, J., & Bos, N. (2012). Inferring Demographic Attributes and Extracting Political Discourse from Nigerian Social Media. In *Proceedings of the 2nd International Conference on Cross-cultural Decision Making*. San Francisco, California.
- Finney, M. (1998). *FARSITE: Fire Area Simulator - Model Development and Evaluation* (Tech. Rep.). Ogden, UT: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station.
- Fischer, C. (1982). *To Dwell Among Friends: Personal Networks in Town and City*.

Chicago, IL: University of Chicago Press.

Fischer, H. W. (1996). What Emergency Management Officials Should Know To Enhance Mitigation and Effective Disaster Response. *Journal of Contingencies and Crisis Management*, 4(4), 208–217.

Fons, W. (1946). Analysis of Fire Spread in Light Forest Fuels. *Journal of Agricultural Research*, 72, 93–121.

Ford, H. (2012). *Can Ushahidi Rely on Crowdsourced Verifications?* Retrieved March 1, 2014, from <http://hblog.org/2012/03/09/can-ushahidi-rely-on-crowdsourced-verifications/>

Forestry Canada Fire Danger Group. (1992). *Development and structure of the canadian forest fire behavior prediction system* (Tech. Rep. No. Information Report ST-X-3).

Frantzich, H. (2001). Occupant Behaviour and Response Time - Results from Evacuation Experiments. In *Human Behaviour in Fire - Proceedings of the Second International Symposium* (pp. 159–165). Boston, MA.

Fraustino, J. D., Liu, B., & Jin, Y. (2012). *Social Media Use during Disasters: A Review of the Knowledge Base and Gaps* (Tech. Rep.). Final Report to Human Factors/Behavioral Sciences Division, Science and Technology Directorate, U.S. Department of Homeland Security: College Park, MD: START.

Frazier, T. J., & Alfons, A. (2012). Generating a Close-to-Reality Synthetic Population of Ghana. *SSRN Electronic Journal*.

Frederick, C. J. (1980). Effects of Natural vs. Human Induced Violence Upon Victims. In L. Kivens (Ed.), *Evaluation & Change: Services for Survivors* (Vol. 71). Minneapolis, Minnesota: Minneapolis Medical Research Foundation.

Friedkin, N., & Johnsen, E. C. (1999). Social Influence Networks and Opinion Change. In *Advances in Group Processes* (Vol. 16, pp. 1–29). JAI Press Inc.

Friedland, G., & Sommer, R. (2010). Cybercasing the Joint: on the Privacy Implications of Geo-tagging. In *Proceedings of the Fifth USENIX Workshop on Hot Topics in*

- Security*. Washington, DC.
- Friedman, T. L. (2006). *The World is Flat: A Brief History of the Twenty-first Century*. New York, NY: Farrar, Straus, and Giroux.
- Fruin, J. (1971). *Pedestrian Planning and Design*. Ann Arbor, MI: Metropolitan Association of Urban Designers and Environmental Planners.
- Galuba, W., Aberer, K., Chakraborty, D., Despotovic, Z., & Kellerer, W. (2010). Outtweeting the Twitterers - Predicting Information Cascades in Microblogs. In *Proceedings of the 3rd Conference on Online Social Networks*. Berkeley, CA, USA: USENIX Association.
- George, A. L. (1986). The Impact of Crisis Induced Stress on Decision-making. In F. Solomon & R. Q. Marston (Eds.), *Medical Implications of Nuclear War*. Washington, DC: National Academy of Sciences Press.
- Georgopoulos, P. G., Fedele, P., Shade, P., Liroy, P. J., Hodgson, M., Longmire, A., ... Brown, M. A. (2004). Hospital response to chemical terrorism: Personal protective equipment, training, and operations planning. *American Journal of Industrial Medicine*, 46(5), 432-445.
- Gershon, R. R. M., Hogan, P. H. G., Qureshi, K. A., & Doll, L. (2004). Preliminary results from the world trade center evacuation study. *Morbidity and Mortality Weekly Report*, 53(35), 815-817.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the Fast and Frugal Way: Models of Bounded Rationality. *Psychological Review*, 103(4), 650-69.
- Gigerenzer, G., Group, A. R., & Todd, P. M. (1999). *Simple heuristics that make us smart*. Oxford, UK: Oxford University Press.
- Gilbert, E., & Karahalios, K. (2010). Widespread Worry and the Stock Market. In *Proceedings of the International Conference on Weblogs and Social Media*.
- Gilbert, N., & Troitzsch, K. (2005). *Simulation for the Social Scientist*. Milton Keynes, UK: Open University Press.
- Glance, N., Hurst, M., Nigam, K., Siegler, M., Stockton, R., & Tomokiyo, T. (2005).

- Deriving Marketing Intelligence from Online Discussion. In *Proceeding of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining* (p. 419). New York, New York, USA: ACM Press.
- Glasgow, K., Ebaugh, A., & Fink, C. (2012). #Londonsburning: Integrating Geographic Topical , and Social Information during Crisis. In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* (pp. 8–13). Trinity College, Dublin, Ireland.
- Go, A., Huang, L., & Bhayani, R. (2009). *Twitter Sentiment Analysis*.
- Gode, D. K., & Sunder, S. (1993, February). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1), 119–137.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of Ecological Rationality: the Recognition Heuristic. *Psychological Review*, 109(1), 75–90.
- Gómez Hidalgo, J. M., Bringas, G. C., Sáenz, E. P., & García, F. C. (2006). Content-based SMS Spam Filtering. In *Proceedings of the 2006 ACM Symposium on Document Engineering* (p. 107). New York, New York, USA: ACM Press.
- González-Bailón, S., Wang, N., Rivero, A., Borge-Holthoefer, J., & Moreno, Y. (2012). Assessing the Bias in Communication Networks Sampled from Twitter. *SSRN Electronic Journal*, 44.
- Goodchild, M. F. (2007). Citizens as Sensors: the World of Volunteered Geography. *GeoJournal*, 69(4), 211–221.
- Goodchild, M. F. (2008). Assertion and Authority: the Science of User-generated Geographic Content. In *Proceedings of the Colloquium for Andrew U Frank's 60th Birthday, GeoInfo 39* (pp. 1–18). Vienna, Austria.
- Goodchild, M. F., & Glennon, J. A. (2010). Crowdsourcing Geographic Information for Disaster Response: a Research Frontier. *International Journal of Digital Earth*, 3(3), 231–241.
- Graham, M., Poorthuis, A., & Zook, M. (2012). Digital trails of the uk floods: How well

- do tweets match observations? *The Guardian Datablog*.
- Granovetter, M. S. (1973). The Strength of Weak Ties. *American Journal of Sociology*, 78(6), 1360.
- Granovetter, M. S. (1978). Threshold Models of Collective Behavior. *The American Journal of Sociology*, 83(6), 1420–1443.
- Grimm, V., Berger, U., Bastiansen, F., Eliassen, S., Ginot, V., Giske, J., ... Huse, G. (2006). A standard protocol for describing individual-based and agent-based models. *Ecological Modelling*, 198(1), 115–126.
- Gruhl, D., Liben-Nowell, D., Guha, R., & Tomkins, A. (2004). Information Diffusion Through Blogspace. *ACM SIGKDD Explorations Newsletter*, 6(2), 43–52.
- Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013). Faking sandy: Characterizing and identifying fake images on twitter during hurricane sandy. In *holderProceedings of the 22nd international conference on World Wide Web* (p. 729–736). International World Wide Web Conferences Steering Committee.
- Gupta, R., & Ratnov, L.-A. (2008). Text categorization with knowledge transfer from heterogeneous data sources. In *AAAI* (p. 842–847).
- Haklay, M. (2010). How Good is Volunteered Geographical Information? A Comparative Study of OpenStreetMap and Ordnance Survey Datasets. *Environment and Planning B*, 37(4), 682–703.
- Haklay, M., Singleton, A., & Parker, C. (2008). Web Mapping 2.0: The Neogeography of the GeoWeb. *Geography Compass*, 2(6), 2011–2039.
- Haklay, M. M., Basiouka, S., Antoniou, V., & Ather, A. (2010). How Many Volunteers Does it Take to Map an Area Well? The Validity of Linus’ Law to Volunteered Geographic Information. *The Cartographic Journal*, 47(4), 315–322.
- Hampton, K. N., Wellman, B., Ronald, T., Behrens, D., Bello, N., Erickson, B., ... Shipley, A. (2000). Examining Community in the Digital Neighborhood: Early Results from

- Canada's Wired Suburb. In T. Ishida & K. Isbister (Eds.), *Digital Cities* (pp. 194–208). Springer Berlin Heidelberg.
- Hannon, B., & Ruth, M. (1994). *Dynamic modeling*. Springer-verlag New York.
- Harland, K., Heppenstall, A. J., Smith, D., & Birkin, M. (2012). Creating Realistic Synthetic Populations at Varying Spatial Scales: A Comparative Critique of Population Synthesis Techniques. *JASSS*, 15(1).
- Harvard Humanitarian Initiative. (2011). *Disaster Relief 2.0: The Future of Information Sharing in Humanitarian Emergencies* (Tech. Rep.). Washington, D.C. and Berkshire, UK: UN Foundation & Vodafone Foundation Technology Partnership.
- Hatzivassiloglou, V., & Wiebe, J. M. (2000). Effects of adjective orientation and gradability on sentence subjectivity. In *Proceedings Of The 18Th Conference On Computational Linguistics* (p. 299-305). Association for Computational Linguistics.
- Healy, M., Delany, S. J., & Zamolotskikh, A. (2004). An Assessment of Case Base Reasoning for Short Text Message Classification. In *The 15th Artificial Intelligence and Cognitive Science Conference* (pp. 9–18). Castlebar, Ireland.
- Heinzelman, J., & Waters, C. (2010). *Crowdsourcing Crisis Information in Disaster-Affected Haiti* (Tech. Rep.). United States Institute of Peace.
- Helbing, D. (1996). Gas-kinetic Derivation of Navier-Stokes-like Traffic Equations. *Physical Review E*, 53(3), 2266–2381.
- Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating Dynamical Features of Escape Panic. *Nature*, 407(6803), 487–90.
- Heppenstall, A., Crooks, A., See, L., & Batty, M. (Eds.). (2012). *Agent-based Models of Geographic Systems*. New York, NY: Springer.
- Hernández Encinas, A., Hernández Encinas, L., Hoya White, S., Martín del Rey, A., & Rodríguez Sánchez, G. (2007). Simulation of Forest Fire Fronts Using Cellular Automata. *Advances in Engineering Software*, 38(6), 372–378.
- Hernández Encinas, L., Hoya White, S., Martín del Rey, A., & Rodríguez Sánchez, G.

- (2007). Modelling Forest Fire Spread Using Hexagonal Cellular Automata. *Applied Mathematical Modelling*, 31(6), 1213–1227.
- Hillan, J. (2012). *Waldo Canyon And Flagstaff Fires Force Road Closures*. Retrieved April 3, 2014, from <http://denver.cbslocal.com/2012/06/27/waldo-canyon-and-flagstaff-fire-forces-road-closures/>
- Hirsch, K. (2003). *Canadian Forest Fire Behavior Prediction (FBP) System: User's Guide* (Tech. Rep.). Northwest Region, Northern Forestry Centre: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station.
- Hoadley, C. M., Xu, H., Lee, J. J., & Rosson, M. B. (2010). Privacy as Information Access and Illusory Control: the Case of the Facebook News Feed Privacy Outcry. *Electronic Commerce Research and Applications*, 9(1), 50–60.
- Hoard, M., Homer, J., Manley, W., Furbee, P., Haque, A., & Helmkamp, J. (2005). Systems Modeling in Support of Evidence-based Disaster Planning for Rural Areas. *International Journal of Hygiene and Environmental Health*, 208(1-2), 117–25.
- Homer, C., Dewitz, J., Fry, J., Coan, M., Hossain, N., Larson, C., ... Wickham, J. (2007). Completion of the 2001 National Land Cover Database for the Conterminous United States. *Photogrammetric Engineering & Remote Sensing*, 73(4), 337–341.
- Howe, J. (2006). The Rise of Crowdsourcing. *Wired*, 161-165. Retrieved September 25th, 2008, from <http://www.wired.com/wired/archive/14.06/crowds.html>
- Hu, M., & Liu, B. (2004). Mining and Summarizing Customer Reviews. In *Proceedings of the 2004 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (p. 168). New York, New York, USA: ACM Press.
- Huang, Z., & Williamson, P. (2001). *A Comparison of Synthetic Reconstruction and Combinatorial Optimisation Approaches to the Creation of Small-Area Microdata* (Tech. Rep.). Department of Geography, University of Liverpool.
- Huberman, B. A., Romero, D. M., & Wu, F. (2009). Social Networks that Matter: Twitter Under the Microscope. *First Monday*, 14(1).

- Hudson-Smith, A., Batty, M., Crooks, A., & Milton, R. (2009). Mapping for the Masses: Accessing Web 2.0 Through Crowdsourcing. *Social Science Computer Review*, 27(4), 524–538.
- Hudson-Smith, A., Crooks, A., Gibin, M., Milton, R., & Batty, M. (2009). NeoGeography and Web 2.0: concepts, tools and applications. *Journal of Location Based Services*, 3(2), 118–145.
- Huete, A., Justice, C., & van Leeuwen, W. (1999). *MODIS Vegetation Index (MOD13): Algorithm Theoretical Basis Document* (Tech. Rep. No. Mod 13). Greenbelt, MD: University of Maryland/NASA.
- Huettnner, A., & Subasic, P. (2000). Fuzzy Typing for Document Management. In *ACL 2000 Companion Volume: Tutorial Abstracts and Demonstration Notes* (pp. 26–27). Hong Kong, China.
- Hughes, A. L., & Palen, L. (2009). Twitter Adoption and Use in Mass Convergence and Emergency Events. *International Journal of Emergency Management*, 6(3/4), 248.
- Hull, D. a. (1996). Stemming Algorithms: a Case Study for Detailed Evaluation. *Journal of the American Society for Information Science*, 47(1), 70–84.
- Hupert, N., Mushlin, a. I., & Callahan, M. a. (2002). Modeling the Public Health Response to Bioterrorism: Using Discrete Event Simulation to Design Antibiotic Distribution Centers. *Medical Decision Making*, 22(5), S17–S25.
- Hutchison, E. D. (2011). *Dimensions of Human Behavior* (4th ed.). SAGE. Retrieved from <http://www.sagepub.com/booksProdDesc.nav?prodId=Book234064>
- Iftene, A., & Ginsca, A. L. (2012). Using Opinion Mining Techniques for Early Crisis Detection. *International Journal Of Computers, Communications & Control*, 7(5), 857–864.
- Ito! (2010). *ITO World at TED 2010 - Project Haiti*. Retrieved March 1, 2014, from <http://itoworld.blogspot.co.uk/2010/02/ito-world-at-ted-2010-project-haiti.html>
- Izard, C. E. (2007). Basic Emotions, Natural Kinds, Emotion Schemas, and a New

- Paradigm. *Perspectives on Psychological Science*, 2(3), 260–280.
- Jackson, S., W., M., Agouris, P., Crooks, A. T., Croitoru, A., & Stefanidis, A. (2013). Assessing Completeness and Spatial Error of Features in Volunteered Geographic Information. *ISPRS International Journal of Geo-Information*, 2(2), 507–530.
- Jager, W., & Mosler, H. J. (2007). Simulating Human Behavior for Understanding and Managing Environmental Resource Use. *Journal of Social Issues*, 63(1), 97–116.
- Jakob, N., & Gurevych, I. (2010). Extracting Opinion Targets in a Single- and Cross-Domain Setting with Conditional Random Fields. In *Proceedings of Conference on Empirical Methods in Natural Language Processing* (pp. 1035–1045). MIT, Massachusetts.
- Java, A., Song, X., Finin, T., & Tseng, B. (2007). Why We Twitter: Understanding Microblogging Usage and Communities. In *Knowledge Discovery and Data Mining* (pp. 56–65). San Jose, California: ACM Press.
- Jha, M., Moore, K., & Pashaie, B. (2004). Emergency Evacuation Planning with Microscopic Traffic Simulation. *Transportation Research Record*, 1886(1), 40–48.
- Jiang, L., Yu, M., Zhou, M., Liu, X., & Zhao, T. (2011). Target-dependent Twitter Sentiment Classification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics* (pp. 151–160). Portland, Oregon.
- Johansson, A., & Kretz, T. (2012). Applied Pedestrian Modeling. In A. Heppenstall, A. Crooks, L. See, & M. Batty (Eds.), *Agent-based Models of Geographical Systems* (pp. 451–462). New York, NY: Springer.
- John Simerman. (2005, December). Majority of New Orleans deaths tied to floodwalls’ collapse. *The Mercury News*.
- Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review*, 303(3,582-591).
- Kalnikaite, V., Sellen, A., Whittaker, S., & Kirk, D. (2010). Now Let Me See Where I Was. In *Proceedings of the 28th International Conference on Human Factors in Computing Systems* (p. 2045). New York, New York, USA: ACM Press.

- Kantrowitz, M. (2003). *Method And Apparatus For Analyzing Affect And Emotion In Text*.
- Kaplan, A. M., & Haenlein, M. (2010). Users of the world, unite! the challenges and opportunities of social media. *Business Horizons*, 53(1), 59-68.
- Keen, A. (2007). *The Cult of the Amateur*. Nicholas Brealey Publishing Ltd.
- Kennedy, J. F., Kennedy, J., & Eberhart, R. C. (2001). *Swarm intelligence*. Morgan Kaufmann.
- Kennedy, W. G. (2012). Modelling Human Behaviour in Agent-Based Models. In M. Batty, A. Heppenstall, & A. Crooks (Eds.), *Agent-based Models of Geographical Systems* (pp. 167–179). Springer.
- Kennedy, W. G., & Bassett, J. K. (2011). Implementing a "Fast and Frugal" Cognitive Model within a Computational Social Simulation. In *Second Annual Meeting of the Computational Social Science Society of the Americas*. Santa Fe, NM.
- Kerridge, J., Hine, J., & Wigan, M. (2001). Agent-based modelling of pedestrian movements: The questions that need to be asked and answered. *Environment and Planning B*, 28(3), 327-342.
- Khan, F. A., & Salman, A. (2012). A Simple Human Vulnerability Index to Climate Change Hazards for Pakistan. *International Journal of Disaster Risk Science*, 3(3), 163–176.
- Kiel, L. D. (1995). Chaos Theory and Disaster Response Management: Lessons for Managing Periods of Extreme Instability. In G. A. Koehler (Ed.), *What Disaster Response Management Can Learn From Chaos Theory* (pp. 1–22).
- Kilduff, M., & Tsai, W. (2003). *Social Networks and Organizations*. SAGE Publications.
- Kim, S., George, B., & Shekhar, S. (2007). Evacuation Route Planning. In *Proceedings of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems* (p. 1). New York, New York, USA: ACM Press.
- Kim, S., & Hovy, E. (2004). Determining the Sentiment of Opinions. In *Proceedings of the International Conference on Computational Linguistics*.
- Kinater, M., Pauli, P., Müller, M., Krieger, J., Heimbecher, F., Rönau, I., ...

- Mühlberger, A. (2013). Human Behaviour in Severe Tunnel Accidents: Effects of Information and Behavioural Training. *Transportation Research Part F: Traffic Psychology and Behaviour*, 17, 20–32.
- Kleit, R. G. (2001). The role of neighborhood social networks in scattered-site public housing residents' search for jobs. *Housing Policy Debate*, 12(3), 541–573.
- Kobes, M., Helsloot, I., de Vries, B., & Post, J. G. (2010). Building Safety and Human Behaviour in Fire: A Literature Review. *Fire Safety Journal*, 45(1), 1–11.
- Kochanski, A., Jenkins, M., Mandel, J., Beezley, J., & Krueger, S. (2013). Real Time Simulation of 2007 Santa Ana Fires. *Forest Ecology and Management*, 294, 136–149.
- Koppel, M., & Shtrimberg, I. (2006). Good News or Bad News? Let the Market Decide. In J. G. Shanahan, Y. Qu, & J. Wiebe (Eds.), *Computing Attitude and Affect in Text: Theory and Applications*. Springer Netherlands.
- Kouloumpis, E., Wilson, T., & Moore, J. (2011). Twitter Sentiment Analysis: The Good the Bad and the OMG! In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media* (pp. 538–541).
- Kowald, M., & Axhausen, K. W. (2011). Surveying Data on Connected Personal Networks. In *9th International Conference on Transport Survey Methods*. Termas de Puyehue, Chile.
- Krishnamurthy, B., Gill, P., & Arlitt, M. (2006). A Few Chirps About Twitter. In *First Workshop on Online Social Networks* (pp. 19–24). Seattle, Washington.
- Kuligowski, E. D. (2011). Predicting Human Behavior During Fires. *Fire Technology*, 49(1), 101–120.
- Kuligowski, E. D., & Mileti, D. S. (2009). Modeling Pre-evacuation Delay by Occupants in World Trade Center Towers 1 and 2 on September 11, 2001. *Fire Safety Journal*, 44(4), 487–496.
- Kulshrestha, A., Wu, D., Lou, Y., & Yin, Y. (2011). Robust shelter locations for evacuation planning with demand uncertainty. *Journal of Transportation Safety & Security*, 3(4), 272–288.

- Küpper, A. (2005). *Location-Based Services: Fundamentals and Operation*. New York, NY: Wiley.
- Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a Social Network or a News Media? In *Proceedings of the 19th International Conference on the World Wide Web*. Raleigh, North Carolina, USA: ACM Press.
- Kwak, H., Moon, S., & Lee, W. (2012). More of a Receiver Than a Giver: Why Do People Unfollow in Twitter? In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media* (pp. 499–502).
- Latane, B. (1981). The Psychology of Social Impact. *American Psychologist*, 36(4), 343–356.
- Lau, J., Kim, J., Tsui, H., & Griffiths, S. (2007). Perceptions Related To Human Avian Influenza And Their Associations With Anticipated Psychological And Behavioral Responses At The Onset Of Outbreak In The Hong Kong Chinese General Population. *American Journal of Infection Control*, 35, 38–49.
- Lavrenko, V., Schmill, M., Lawrie, D., Ogilvie, P., Jensen, D., & Allan, J. (2000). Mining of Concurrent Text and Time Series. In *Proceedings of the 6th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining Workshop on Text Mining* (pp. 2–9). Boston, Massachusetts.
- Law, A. M., & Kelton, D. (1991). *Simulation Modelling and Analysis* (2nd ed.). New York, NY: McGraw-Hill.
- Lazarsfeld, P. F., & Merton, R. K. (1954). Friendship as a Social Process: A Substantive and Methodological Analysis. In M. Berger, T. Abel, & C. H. Page (Eds.), *Freedom and Control in Modern Society* (pp. 18–66). New York, NY: Van Nostrand.
- Lee, E. K., Chen, C.-H., Pietz, F., & Benecke, B. (2009). Modeling and Optimizing the Public-Health Infrastructure for Emergency Response. *Interfaces*, 39(5), 476–490.
- Lehman, J., Laird, J., & Rosenbloom, P. (2006). *A Gentle Introduction to SOAR, an Architecture for Human Cognition: 2006 Update*. Retrieved March 1, 2014, from <http://ai.eecs.umich.edu/soar/sitemaker/docs/misc/GentleIntroduction-2006.pdf>

- Leonard, G. S., Power, W., Lukovic, B., Smith, W., Langridge, R., Johnston, D. M., & Downes, G. (2008). *Interim Tsunami Evacuation Planning Zone Boundary Mapping for the Wellington and Horizons Regions Defined by a GIS-calculated Attenuation Rule* (Tech. Rep.). GNS Science Report.
- Leow, S. (2009). *You Don't Know Me but Can I be Your Friend? Accepting Strangers as Friends in Facebook*. Unpublished doctoral dissertation, University of North Texas.
- Lerman, K., & Galstyan, A. (2008). Analysis of Social Voting Patterns on Digg. In *Proceedings of the first workshop on Online social networks* (p. 7). New York, New York, USA: ACM Press.
- Lerman, K., & Ghosh, R. (2010). Information Contagion: an Empirical Study of the Spread of News on Digg and Twitter Social Networks.
- Leskovec, J., Backstrom, L., & Kleinberg, J. (2009). Meme-tracking and the Dynamics of the News Cycle. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (p. 497). New York, New York, USA: ACM Press.
- Leskovec, J., McGlohon, M., Faloutsos, C., Glance, N., & Hurst, M. (2007). *Cascading Behavior in Large Blog Graphs* (Tech. Rep.). Retrieved from <http://arxiv.org/abs/0704.2803>
- Liben-Nowell, D., & Kleinberg, J. (2008). Tracing Information Flow on a Global Scale using Internet Chain-letter Data. *Proceedings of the National Academy of Sciences of the United States of America*, 105(12), 4633–8.
- Lindell, M. K. (2011). Disaster Studies. *sociopedia.isa*.
- Linn, R., Reisner, J., Colman, J., & Winterkamp, J. (2002). Studying Wildfire Behavior Using FIRETEC. *International Journal of Wildland Fire*, 11, 233–246.
- Liu, B. (2012). *Sentiment Analysis and Opinion Mining* (Vol. 5) (No. 1). Morgan & Claypool Publishers.
- Liu, H. X., He, X., & Ban, X. (2007). A Cell-Based Many-to-One Dynamic System Optimal

- Model and Its Heuristic Solution Method for Emergency Evacuation. In *Transportation Research Board 86th Annual Meeting*. Washington, DC.
- Liu, S. B., & Palen, L. (2010). The New Cartographers: Crisis Map Mashups and the Emergence of Neogeographic Practice. *Cartography and Geographic Information Science*, 37(1), 69–90.
- Loewenstein, G., & Lerner, J. (2003). The Role of Affect in Decision Making. In R. Davidson, K. Scherer, & H. Goldsmith (Eds.), *Handbook of Affective Science* (pp. 619–642). Oxford, UK: Oxford University Press.
- Longley, P., Goodchild, M., Maguire, D., & Rhind, D. (2010). *Geographical Information Systems and Science* (3rd ed.). New York, NY: John Wiley and Sons.
- Longley, P. A., & Singleton, A. D. (2009). Classification through Consultation: Public Views of the Geography of the e-Society. *International Journal of Geographical Information Science*, 23(6), 737–763.
- Lotan, G., Graeff, E., Ananny, M., Gaffney, D., Pearce, I., & boyd, D. (2011). The Revolutions Were Tweeted: Information Flows During the 2011 Tunisian and Egyptian Revolutions. *International Journal of Communication*, 5, 1375–1406.
- Louch, H. (2000). Personal Network Integration: Transitivity and Homophily in Strong-tie Relations. *Social Networks*, 22(1), 45–64.
- Lyons, K., & Lessard, L. (2012). S-FIT. In *Proceedings of the 2012 iConference* (pp. 263–270). New York, New York, USA: ACM Press.
- Malleson, N. (2008). *The PECS Behavioural Framework - A Brief Summary* (Tech. Rep.).
- Malleson, N., Heppenstall, A. J., & See, L. (2009). *Simulating Burglary with an Agent-Based Model*.
- Mandelbrot, B. B. (1983). *The Fractal Geometry of Nature*. New York, NY: W. H. Freeman.
- Margerit, J., & Séro-Guillaume, O. (2002). Modelling Forest Fires. Part II: Reduction to Two-dimensional Models and Simulation of Propagation. *International Journal of Heat and Mass Transfer*, 45(8), 1723–1737.
- Marsden, P. V. (1988). Homogeneity in Confiding Relations. *Social Networks*, 10(1),

57–76.

- Marwick, A. E., & Boyd, D. (2010). I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience. *New Media & Society*, 13(1), 114–133.
- Maslow, A. (1943). A Theory of Human Motivation. *Psychological Review*, 50, 370–396.
- McCarty, C., Killworth, P. D., Bernard, H. R., Johnsen, E. C., & Shelley, G. A. (2001). Comparing Two Methods for Estimating Network Size. *Human Organization*, 60(1), 28–39.
- McConnell, N. C., Boyce, K. E., Shields, J., Galea, E. R., Day, R. C., & Hulse, L. M. (2010). The UK 9/11 Evacuation Study: Analysis of Survivors’ Recognition and Response Phase in WTC1. *Fire Safety Journal*, 45(1), 21–34.
- McGlohon, M., Gance, N., & Reiter, Z. (2010). Star Quality: Aggregating Reviews To Rank Products And Merchants. In *Proceedings of the International Conference on Weblogs and Social Media*.
- McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a Feather: Homophily in Social Networks. *Annual Review of Sociology*, 27(1), 415–444.
- McPherson, T. N., & Brown, M. J. (2004). *Estimating Daytime and Nighttime Population Distributions in U.S. Cities for Emergency Response Activities* (Tech. Rep.). Los Alamos National Laboratory.
- Mejova, Y. (2009). *Sentiment Analysis: An Overview* (Tech. Rep.). Computer Science Department, University of Iowa.
- Mell, W., Charney, J. J., Jenkins, M. A., Cheney, P., & Gould, J. (2005). Numerical Simulations of Grassland Fire Behavior from the LANL - FIRETEC and NIST - WFDS Models. In *EastFIRE Conference* (pp. 1–10). Fairfax, VA: George Mason University.
- Mell, W., Jenkins, M. A., Gould, J., & Cheney, P. (2007). A Physics-based Approach to Modelling Grassland Fires. *International Journal of Wildland Fire*, 16(1), 1.
- Mellers, B. A., Schwartz, A., & Cooke, a. D. (1998). Judgment and Decision Making.

- Annual Review of Psychology*, 49, 447–77.
- Mihalcea, R., & Liu, H. (2006). A Corpus-based Approach to Finding Happiness. In *AAAI 2006 Symposium on Computational Approaches to Analysing Weblogs* (pp. 271–280). AAAI Press.
- Milgram, S. (1967). The Small-World Problem. *Psychology Today*, 1(1), 61–67.
- Miller, C. C. (2006). A Beast in the Field: The Google Maps Mashup as GIS/2. *Cartographica: The International Journal for Geographic Information and Geovisualization*, 41(3), 187–199.
- Miller, J., & Page, S. (2007). *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press.
- Min, P., & Hong, C. (2011). System Dynamics Analysis for the Impact of Dynamic Transport and Information Delay to Disaster Relief Supplies. In *18th International Conference on Management Science and Engineering* (pp. 1–6). Rome, Italy.
- Minshew, C., & Schneider, D. J. (2012). *2012 Colorado Wildfires - At a Glance*. Retrieved January 14, 2014, from http://www.denverpost.com/wildfires/ci_20998199/2012-colorado-wildfire-overview#
- Muller, G., Grébaut, P., & Gouteux, J.-P. (2004). An Agent-Based Model Of Sleeping Sickness: Simulation Trials Of A Forest Focus In Southern Cameroon. *Comptes Rendus Biologies*, 327(1), 1–11.
- Müller, K., & Axhausen, K. W. (2010). Population Synthesis for Microsimulation: State of the Art. In *Swiss Transport Research Conference*.
- Mulligan, K., & Scherer, K. R. (2012). Toward a Working Definition of Emotion. *Emotion Review*, 4(4), 345–357.
- Munro, R. (2010). Crowd-Sourced Translation For Emergency Response In Haiti: The Global Collaboration Of Local Knowledge. In *Relief 2.0 in Haiti*.
- Munro, R., & Manning, C. (2010). Subword Variation In Text Message Classification. In *The Human Language Technologies: The Annual Conference of the North American*

Chapter of the ACL.

- Mussavi Rizi, S. M., Latek, M. M., & Geller, A. (2013). Fusing Remote Sensing with Sparse Demographic Data for Synthetic Population Generation: an Algorithm and Application to Rural Afghanistan. *International Journal of Geographical Information Science*, 27(5), 986–1004.
- Nagy, A., & Stamberger, J. (2012). Crowd Sentiment Detection during Disasters and Crises. In L. Rothkrantz, J. Ristvej, & Z. Franco (Eds.), *Proceedings of the 9th International International Conference on Information Systems for Crisis Response and Management Conference* (pp. 1–9). Vancouver, Canada.
- Narayanan, R., Liu, B., & Choudhary, A. (2009). Sentiment Analysis of Conditional Sentences. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing* (pp. 180–189). Singapore.
- Nasukawa, T., & Yi, J. (2003). Sentiment Analysis: Capturing Favorability Using Natural Language Processing. In *Proceedings of the Conference on Knowledge Capture*.
- National Sleep Foundation. (2008). *2008 Sleep in America Poll: Summary of Findings* (Tech. Rep. No. 202). National Sleep Foundation.
- Nelson, R. (2000). Prediction Of Diurnal Change In 10-H Fuel Stick Moisture Content. *Canadian Journal of Forest Research*, 30, 1071–1087.
- Nielsen. (2012). *State of the Media: The Social Media Report*. Retrieved March 2, 2014, from <http://slidesha.re/Zrwu6V>
- Nilsson, D., Johansson, M., & Frantzich, H. k. (2009). Evacuation Experiment in a Road Tunnel: A Study of Human Behaviour and Technical Installations. *Fire Safety Journal*, 44(4), 458–468.
- Norén, A., & Winér, J. (2003). *Modelling Crowd Evacuation from Road and Train Tunnels - Data and Design for Faster Evacuations* (Tech. Rep.). Lund, Department of Fire Safety Engineering, Lund University.
- Norheim-Hagtun, I., & Meier, P. (2010). Crowdsourcing for Crisis Mapping in Haiti. *Innovations: Technology, Governance, Globalization*, 5(4), 81–89.

- North, M. J., & Macal, C. M. (2007). *Managing Business Complexity: Discovering Strategic Solutions with Agent-based Modelling and Simulation*. New York, NY: Oxford University Press.
- O'Connor, B., Balasubramanyan, R., Routledge, B. R., & Smith, N. A. (2010). From Tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. In *The International AAAI Conference on Weblogs and Social Media* (p. 8). Washington, DC.
- O'Donnell, I., Smart, K., & Ramalingam, B. (2009). *Responding to Urban Disasters* (Tech. Rep.). ALNAP.
- Oh, O., Agrawal, M., & Rao, H. R. (2010). Information Control and Terrorism: Tracking the Mumbai Terrorist Attack through Twitter. *Information Systems Frontiers*, 13(1), 33–43.
- Okazaki, S., & Matsushita, S. (1993). A Study of Simulation Model for Pedestrian Movement with Evacuation and Queuing. In R. Smith & J. Dickie (Eds.), *Engineering for Crowd Safety: Proceedings of the International Conference on Engineering for Crowd Safety* (pp. 271–280). London, UK.
- Okolloh, O. (2009). *Thoughts on Implementing Ushahidi*. Retrieved July 11, 2009, from <http://blog.ushahidi.com/index.php/2009/06/08/thoughts-on-implementing-ushahidi/>
- Oldham, J., Crawford, A. J., & Jones, T. (2012). *Colorado Wildfire Forces 34,500 to Evacuate as Homes Burn*. Retrieved from <http://www.businessweek.com/news/2012-06-27/colorado-wildfire-forces-32-000-to-evacuate-as-homes-burn>
- OpenStreetMap. (2012). *OpenStreetMap Wiki*. Retrieved May 14, 2012, from <http://wiki.openstreetmap.org/>
- Orcutt, G. H. (1957). A New Type of Socio-economic System. *The Review of Economics and Statistics*, 58, 773–797.
- O'Reilly, T. (2005). *What Is Web 2.0? Design Patterns and Business Models for the Next Generation of Software*. Retrieved March 2, 2014, from

<http://oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>

- Ortony, a., Clore, G., & Foss, M. (1987). The Referential Structure of the Affective Lexicon. *Cognitive Science*, 11(3), 341–364.
- Oxendine, C. E. (2013). *Analysis Of Volunteered Geographic Information For Improved Situational Awareness During No-Notice Emergencies*. Unpublished doctoral dissertation, George Mason University.
- Ozdamar, L., Ekinici, E., & Kucukyazici, B. (2004). Emergency Logistics Planning in Natural Disasters. *Annals of Operations Research*, 129(1-4), 217–245.
- Pang, B., & Lee, L. (2004). A Sentimental Education. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics* (pp. 271–es). Morristown, NJ, USA: Association for Computational Linguistics.
- Pang, B., & Lee, L. (2008). Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, 2(1-2), 1–135.
- Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs Up?: Sentiment Classification using Machine Learning Techniques. In *Proceedings of the ACL-02 conference on Empirical methods in natural language processing* (Vol. 10, pp. 79–86). Philadelphia, Pennsylvania: Association for Computational Linguistics.
- Papazoglou, I., & Christou, M. (1997). A Decision Support System For Emergency Response To Major Nuclear Accidents. *Journal of Nuclear Technology*, 118, 97–122.
- Parker, D. C., & Meretsky, V. (2004). Measuring Pattern Outcomes in an Agent-based Model of Edge-effect Externalities using Spatial Metrics. *Agriculture, Ecosystems & Environment*, 101(2-3), 233–250.
- Pastor, E., Zarate, L., Planas, E., & Arnaldos, J. (2003). Mathematical Models and Calculation Systems for the Study of Wildland Fire Behaviour. *Progress in Energy and Combustion Science*, 29(2), 139–153.
- Paul, J. A., George, S. K., Yi, P., & Lin, L. (2006). Transient Modeling in Simulation of Hospital Operations for Emergency Response. *Prehospital and Disaster Medicine*,

21(4), 223–36.

- Perry, R., & Greene, M. (1983). *Citizen Response to Volcanic Eruptions: the Case of Mount Saint Helens*. New York, NY: Irvington Publishers.
- Petrescu-Prahova, M., & Butts, C. T. (2008). Emergent Coordinators in the World Trade Center Disaster. *International Journal of Mass Emergencies and Disasters*, 26(3), 133–168.
- Petrovic, S., Osborne, M., & Lavrenko, V. (2010). Streaming first story detection with application to twitter. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics* (p. 181-189). Association for Computational Linguistics.
- Pirolli, P. (1999). Cognitive Engineering Models and Cognitive Architectures in Human-Computer Interaction. In F. Durso, R. Nickerson, R. Schvaneveldt, S. Dumais, D. Lindsay, & M. Chi (Eds.), *Handbook of Applied Cognition* (pp. 443–477). John Wiley & Sons Ltd.
- Polgreen, P. M., Chen, Y., Pennock, D. M., & Nelson, F. D. (2008). Using Internet Searches for Influenza Surveillance. *Clinical Infectious Diseases*, 47(11), 1443–8.
- Pontius, G. R., & Malanson, J. (2005). Comparison of the Structure and Accuracy of Two Land Change Models. *International Journal of Geographical Information Science*, 19(2), 243–265.
- Popescu, A.-M., & Etzioni, O. (2005). Extracting Product Features and Opinions from Reviews. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing* (pp. 339–346). Morristown, NJ, USA: Association for Computational Linguistics.
- Proulx, G. (2002). Movement of people: the evacuation timing.
- Proulx, G., & Sime, J. (1991). To Prevent 'Panic' in an Underground Emergency: Why Not Tell People the Truth? In *Fire Safety Science - Proceedings of the Third International Symposium* (pp. 843–852). University of Edinburgh, Scotland, UK.
- Quarantelli, E. L. (1954). The nature and conditions of panic. *American Journal of*

- Sociology*, 267-275.
- Rahm, E., & Do, H. H. (2000). Data Cleaning: Problems and Current Approaches. *IEEE Data Engineering Bulletin*, 23(4), 3–13.
- Ramezankhani, A., & Najafiyazdi, M. (2003). A System Dynamics Approach on Post-Disaster Management: A Case Study of Bam Earthquake, December 2003. In *26th International System Dynamics Conference*. Athens, Greece.
- Rao, A. S., & Georgeff, M. P. (1995). BDI agents: From theory to practice. In *ICMAS* (Vol. 95, p. 312-319).
- Ratti, C., Pulselli, R. M., Williams, S., & Frenchman, D. (2006). Mobile Landscapes: Using Location Data from Cell Phones for Urban Analysis. *Environment and Planning B: Planning and Design*, 33(5), 727–748.
- Regnier, E. (2008). Public Evacuation Decisions and Hurricane Track Uncertainty. *Journal of Management Science*, 54, 16–28.
- Reinhardt, E., & Crookston, N. (2003). *The Fire and Fuels Extension to the Forst Vegetation Simulator* (Tech. Rep.). Fort Collins, CO: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station.
- Rich, J., & Mulalic, I. (2012). Generating Synthetic Baseline Populations from Register Data. *Transportation Research Part A: Policy and Practice*, 46(3), 467–479.
- Riloff, E., & Wiebe, J. (2003). Learning Extraction Patterns for Subjective Expressions. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing* (Vol. 10, pp. 105–112). Morristown, NJ, USA: Association for Computational Linguistics.
- Ripley, A. (2009). *The unthinkable: Who survives when disaster strikes - and why*. New York: Harmony.
- Ritchie-Dunham, J. L., & Méndez Galván, J. F. (1999). Evaluating Epidemic Intervention Policies with Systems Thinking: a Case Study of Dengue Fever in Mexico. *System Dynamics Review*, 15(2), 119–138.
- Robinson, R. M., & Khattak, A. (2012). Evacuee Route Choice Decisions in a Dynamic

- Hurricane Evacuation Context. *Transportation Research Record: Journal of the Transportation Research Board*, 2312(-1), 141–149.
- Roche, S., Propeck-Zimmermann, E., & Mericskay, B. (2011). GeoWeb and Crisis Management: Issues and Perspectives of Volunteered Geographic Information. *GeoJournal*, 78(1), 21–40.
- Rochoux, M., Delmotte, B., Cuenot, B., Ricci, S., & Trouvé, A. (2013). Regional-scale Simulations of Wildland Fire Spread Informed by Real-time Flame Front Observations. *Proceedings of the Combustion Institute*, 34(2), 2641–2647.
- Rodríguez, R., Cortés, A., & Margalef, T. (2010). Towards Policies for Data Insertion in Dynamic Data Driven Application Systems: a Case Study Sudden Changes in Wildland Fire. *Procedia Computer Science*, 1(1), 1267–1276.
- Rosenthal, U., Charles, M. T., & 't Hart, P. (Eds.). (1989). *Coping with Crises: The Management of Disasters, Riots and Terrorism*. Springfield, MD.
- Ryan, J., Maoh, H., & Kanaroglou, P. (2007). *Population Synthesis: Comparing the Major Techniques Using a Small, Complete Population of Firms* (No. October).
- Saadatseresht, M., Mansourian, A., & Taleai, M. (2009). Evacuation Planning Using Multiobjective Evolutionary Optimization Approach. *European Journal of Operational Research*, 198(1), 305–314.
- Sabino, A., Nóbrega, R., Rodrigues, A., & Correia, N. (2008). Life-Saver: Flood Emergency Simulator. In *Proceedings of the 5th International Conference on Information Systems for Crisis Response and Management* (pp. 724–733). Washington, DC.
- Sahito, F., Latif, A., & Slany, W. (2011). Weaving Twitter Stream into Linked Data: A Proof of Concept Framework. In *Proceedings of the 7th International Conference on Emerging Technologies*. Islamabad, Pakistan.
- Sakaki, T., Okazaki, M., & Matsuo, Y. (2010). Earthquake Shakes Twitter Users. In *Proceedings of the 19th International Conference on World Wide Web* (p. 851). New York, New York, USA: ACM Press.
- Salgado, M., Marchione, E., & Gill, A. (2010). The Calm After the Storm? Looting in

- the Context of Disasters. In *3rd World Congress on Social Simulation: Scientific Advances in Understanding Societal Processes and Dynamics* (pp. 1–8). University Of Kassel, Kassel, Germany.
- Samuelson, D. A., Zimmerman, A., Guerin, S., Densmore, O., Miller, L., & Thorp, J. (2008). Agent-Based Simulations of Mass Egress After Improvised Explosive Device Attacks. In *Proceedings of the 5th International Conference on Information Systems for Crisis Response and Management* (pp. 59–70). Washington, DC.
- Santos, G., & Aguirre, B. E. (2004). A Critical Review of Emergency Evacuation Simulation Models. In *NIST Workshop on Building Occupant Movement during Fire Emergencies* (pp. 25–50). Gaithersburg, Maryland.
- 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 *ACM Conference on Computer Supported Cooperative Work*. Bellevue, WA.
- Scellato, S., Noulas, A., Lambiotte, R., & Mascolo, C. (2011). Socio-Spatial Properties of Online Location-Based Social Networks. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media Socio-Spatial* (pp. 329–336).
- Scharl, A., & Tochtermann, K. (Eds.). (2007). *The Geospatial Web: How Geobrowsers, Social Software and the Web 2.0 are Shaping the Network Society*. London, UK: Springer.
- Scherer, K. R. (1999). Appraisal theory. *Handbook of Cognition and Emotion*, 637-663.
- Schmidt, B. (2000). *The modelling of human behaviour: The PECS reference models*. SCS-Europe BVBA.
- Scott, J. (1999). Nexus: A System For Assessing Crown Fire Hazard. *Fire Management Notes*, 59(2), 20–24.
- Séro-Guillaume, O., Ramezani, S., Margerit, J., & Calogine, D. (2008). On Large Scale Forest Fires Propagation Models. *International Journal of Thermal Sciences*, 47(6), 680–694.

- Sharma, S., Singh, H., & Prakash, A. (2008). Multi-agent Modeling and Simulation of Human Behavior in Aircraft Evacuations. *IEEE Transactions on Aerospace and Electronic Systems*, 44(4), 1477–1488.
- Sherman, M., Peyrot, M., Magda, L., & Gershon, R. (2011). Modeling Pre-evacuation Delay by Evacuees in World Trade Center Towers 1 and 2 on September 11, 2001: A Revisit using Regression Analysis. *Fire Safety Journal*, 46(7), 414–424.
- Shi, X., Tseng, B., & Adamic, L. A. (2007). Looking at the Blogosphere Topology through Different Lenses. In *Proceedings of the International Conference on Weblogs and Social Media*.
- Shim, K.-C., Fontane, D. G., & Labadie, J. W. (2002). Spatial Decision Support System for Integrated River Basin Flood Control. *Journal of Water Resources Planning and Management*, 128(3), 190–201.
- Shimojo, A., Kamada, S., Matsumoto, S., & Nakamura, M. (2010). On Integrating Heterogeneous Lifelog Services. In *Proceedings of the 12th International Conference on Information Integration and Web-based Applications & Services* (p. 263). New York, New York, USA: ACM Press.
- Shneiderman, B. (2008). Catalyzing Community Efforts for Disaster Prevention, Response and Recovery. In F. Fiedrich & B. V. de Walle (Eds.), *Proceedings of the 5th International International Conference on Information Systems for Crisis Response and Management Conference*. Washington, DC.
- Shore, J. (1986). *Disaster Stress Studies: New Methods and Findings*. Washington, DC: American Psychiatric Press.
- Sime, J. D. (1985). Movement toward the Familiar: Person and Place Affiliation in a Fire Entrapment Setting. *Environment and Behavior*, 17(6), 697–724.
- Simon, H. (1957). A Behavioral Model of Rational Choice. In *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting* (pp. 271–280). New York, NY: Wiley.
- Simon, H. (1996). *The Sciences of the Artificial*. MIT Press.

- Smith, A. (2011). *Why Americans use Social Media* (Tech. Rep.). Washington, DC: Pew Research Center. Retrieved from <http://www.pewinternet.org/files/old-media//Files/Reports/2011/Why-Americans-Use-Social-Media.pdf>
- Smith, B. W., Kay, V. S., Hoyt, T. V., & Bernard, M. L. (2009). Predicting the Anticipated Emotional and Behavioral Responses to an Avian Flu Outbreak. *American Journal of Infection Control*, 37(5), 371–80.
- Smith, D. M., Clarke, G. P., & Harland, K. (2009). Improving the Synthetic Data Generation Process in Spatial Microsimulation Models. *Environment and Planning A*, 41(5), 1251–1268.
- Smith, D. M., Clarke, G. P., Ransley, J., & Cade, J. (2006). Food Access and Health : A Microsimulation Framework for Analysis. *Chiikigaku Kenkyu (Studies in Regional Science)*, 35(4), 909–927.
- Smith, N. A. (2011). *Linguistic structure prediction*. Morgan and Claypool.
- Smith, W., & Dowell, J. (2000). A Case Study of Co-ordinative Decision-making in Disaster Management. *Ergonomics*, 43(8), 1153–66.
- Starbird, K., & Palen, L. (2010). Pass It On?: Retweeting in Mass Emergency. In *Proceedings of the 7th International International Conference on Information Systems for Crisis Response and Management Conference* (pp. 1–10). Seattle, Washington.
- Starbird, K., & Palen, L. (2012). (How) Will the Revolution be Retweeted?: Information Propagation in the 2011 Egyptian Uprising. In *ACM Conference on Computer Supported Cooperative Work*. Bellevue, WA.
- 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 *ACM 2010 Conference on Computer Supported Cooperative Work*. Savannah, GA.
- Stefanidis, A., Cotnoir, A., Croitoru, A., Crooks, A., Radzikowski, J., & Rice, M. (2013). Demarcating New Boundaries: Mapping Virtual Polycentric Communities through

- Social Media Content. *Cartography and Geographic Information Science*, 40(2), 116–129.
- Stefanidis, A., Crooks, A. T., & Radzikowski, J. (2013). Harvesting Ambient Geospatial Information from Social Media Feeds. *GeoJournal*, 78(2), 319–338.
- Stickgold, E., Lofdahl, C., & Farry, M. (2013). Trust Definitions and Metrics for Social Media Analysis. In *Proceedings of the 22nd Annual Conference on Behavior Representation in Modeling and Simulation*. Ottawa, Ontario, Canada.
- Stonedahl, F., Rand, W., & Wilensky, U. (2010). Evolving viral marketing strategies. In *Proceedings Of The 12Th Annual Conference On Genetic And Evolutionary Computation* (p. 1195-1202). ACM.
- Stratton, R. D. (2006). *Guidance on Spatial Wildland Fire Analysis: Models, Tools, and Techniques* (Tech. Rep. No. December). Fort Collins, CO: U.S. Department of Agriculture, Forest Service, Rocky Mountain Research Station.
- Su, Y., Huang, X., & Jin, Z. (2009). Modeling Emergency Service Centers with Simulation and System Dynamics. In *11th International Conference on Computer Modelling and Simulation* (pp. 305–310). Emmanuel College, Cambridge, UK: IEEE.
- Sui, D., & Goodchild, M. (2011). The Convergence of GIS and Social Media: Challenges for GIScience. *International Journal of Geographical Information Science*, 25(11), 1737–1748.
- Sullivan, A. L., & Knight, I. K. (2008). A Hybrid Cellular Automata/Semi-physical Model of Fire Growth. *Complexity International*, 12, 10.
- Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies and nations*. New York: Doubleday.
- Sutton, J., Palen, L., & Shklovski, I. (2008). Backchannels on the Front Lines: Emergent Uses of Social Media in the 2007 Southern California Wildfires. In *Proceedings of the 5th International Conference on Information Systems for Crisis Response and Management* (pp. 624–632). Washington, DC.

- Szabò, G., Alava, M., & Kertész, J. (2003). Structural transitions in scale-free networks. *Physical Review E*, 67(5).
- Szabo, G., & Huberman, B. a. (2010). Predicting the Popularity of Online Content. *Communications of the ACM*, 53(8), 80.
- Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267-307.
- Tang, T., & Wannemacher, M. (2005). GIS Simulation and Visualization of Community Evacuation. *Middle States Geographer*, 38, 22-30.
- Tapscott, D., & Williams, A. D. (2006). *Wikinomics: How Mass Collaboration Changes Everything*. London: Atlantic Books.
- TechAmerica Foundation. (2012). *Demystifying Big Data: A Practical Guide To Transforming The Business of Government* (Tech. Rep.). TechAmerica Foundation.
- Thompson, J. (1986). The Impact of Crisis Induced Stress on Decision-making. In F. Solomon & R. Q. Marston (Eds.), *The Medical Implications of Nuclear War* (pp. 290-316). Washington, DC: National Academy of Sciences, National Academy Press.
- Thorp, J., Guerin, S., Wimberly, F., Rossbach, M., Densmore, O., Agar, M., & Roberts, D. (2006). Santa Fe on Fire: Agent-Based Modelling of Wildfire Evacuation. In *Proceedings of the Agent 2006 Conference on Social Agents: Results and Prospects*. University of Chicago and Argonne National Laboratory, Chicago, Illinois.
- Tobler, W. (1970). A Computer Movie Simulating Urban Growth in the Detroit Region. *Economic Geography*, 46(2), 234-240.
- Tomintz, M. N., Clarke, G. P., & Rigby, J. E. (2008). The Geography of Smoking in Leeds: Estimating Individual Smoking Rates and the Implications for the Location of Stop Smoking Services. *Area*, 40(3), 341-353.
- Tonkin, E., Pfeiffer, H. D., & Tourte, G. (2012). Twitter, information sharing and the London riots? *Bulletin of the American Society for Information Science and Technology*, 38(2), 49-57.

- Torrens, P. M., & McDaniel, A. W. (2012). Modeling Geographic Behavior in Riotous Crowds. *Annals of the Association of American Geographers*, 120619095844009.
- Tsur, O., Davidov, D., & Rappoport, A. (2010). A Great Catchy Name: Semi-supervised Recognition of Sarcastic Sentences in Online Product Reviews. In *Fourth International AAAI Conference on Weblogs and Social Media In Proceedings of the Fourth International Conference on Weblogs and Social Media*. George Washington University, Washington, DC.
- Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2011). Election forecasts with Twitter: How 140 characters reflect the political landscape. *Social Science Computer Review*, 29(4), 402-418.
- Tung, T. (2012). *Evaluation of Natural Language Processing Techniques for Sentiment Analysis on Tweets*. Unpublished doctoral dissertation, Technische Universitat Darmstadt.
- Tunkelang, D. (2009). *A Twitter Analog to PageRank*. Retrieved March 18, 2013, from <http://thenoisychannel.com/2009/01/13/a-twitter-analog-to-pagerank/>
- Turner, A., Forrest, B., Lorica, B., & Magoulas, R. (2008). *Where 2.0: the state of the geospatial web*. O'Reilly Media, Incorporated.
- Turney, P. D. (2002). Thumbs Up or Thumbs Down?: Semantic Orientation Applied to Unsupervised Classification of Reviews. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics* (p. 417). Morristown, NJ, USA: Association for Computational Linguistics.
- Twitter. (2014). *About Twitter*. Retrieved March 1, 2014, from www.twitter.com
- Tymstra, C., Bryce, R. W., Wotton, B. M., Taylor, S. W., & Armitage, O. B. (2010). *Development and Structure of Prometheus: the Canadian Wildland Fire Growth Simulation Model* (Tech. Rep.). Edmonton, Alberta: Natural Resources Canada, Canadian Forest Service.
- Udell, E. (2012). *Colorado Wildfire: 32,000 People Evacuated in Waldo Canyon Fire*.
- United States Census Bureau. (2011). *2010 Census Summary File 1* (Vol. 4; Tech. Rep.

- No. September). United States Census Bureau.
- United States Census Bureau. (2012). *Public Use Microdata Areas (PUMAs)*. Retrieved March 1, 2014, from <https://www.census.gov/geo/reference/puma.html>
- United States Census Bureau. (2013). *Computer and Internet Use in the United States* (Tech. Rep. No. May). United States Census Bureau.
- United States Geological Study. (2012). *United States Geological Study*. Retrieved September 19, 2012, from <http://www.usgs.gov/>
- Ushahidi. (2014). *Ushahidi*. Retrieved March 1, 2014, from <http://ushahidi.com/>
- Vakalis, D., Sarimveis, H., Kiranoudis, C. T., Alexandridis, A., & Bafas, G. (2004). A GIS-based Operational System for Wildland Fire Crisis Management I. Mathematical Modelling and Simulation. *Applied Mathematical Modelling*, 28(4), 389–410.
- Van Wagner, C. (1977). Conditions For The Start And Spread Of A Crown Fire. *Canadian Journal of Forest Research*, 71(3), 23–34.
- Vatsavai, R., & Bhaduri, B. (2013). Geospatial Analytics for Big Spatiotemporal Data: Algorithms, Applications, and Challenges. In *NSF Workshop on Big Data and Extreme-Scale Computing*. Charleston, SC.
- Vaught, C., Brnich, M., Mallett, L., Cole, H., Wiehagen, W., Conti, R., ... Litton, C. (2000). *Behavioral and Organizational Dimensions of Underground Mine Fires* (Tech. Rep. No. 2000-126). Pittsburgh, PA: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, National Institute for Occupational Safety and Health.
- Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010). Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness. In *CHI 2010: Crisis Informatics* (pp. 1079–1088). Atlanta, GA.
- Wang, J., & Gilmer, G. (2012). Disaster Response Modeling through Discrete-Event Simulation. In (pp. 68–73).
- Wang, X., Wei, F., Liu, X., Zhou, M., & Zhang, M. (2011). Topic Sentiment Analysis in Twitter: A Graph-based Hashtag Sentiment Classification Approach. In *Proceeding of*

- the ACM Conference on Information and Knowledge Management* (pp. 1031–1040). Glasgow, Scotland, UK.
- Wasserman, S., & Faust, K. (1994). *Social Network Analysis: Methods and Applications*. Cambridge University Press.
- Watts, D. (1999). *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, New Jersey: Princeton University Press.
- Watts, D. (2003). *Six Degrees: The Science Of A Connected Age*. New York, New York: W.W. Norton and Co.
- Watts, D. J. (2013). Computational Social Science: Exciting Progress and Future Directions. *The Bridge on Frontiers of Engineering*, 43(4), 5–10.
- Watts, D. J., & Strogatz, S. H. (1998). Collective Dynamics of ‘Small-world’ Networks. *Nature*, 393(6684), 440–2.
- Wayant, N., Crooks, A. T., Stefanidis, A., Croitoru, A., Radzikowski, J., Stahl, J., & Shine, J. (2012). Spatiotemporal Clustering of Twitter Feeds for Activity Summarization. In *International Conference on Geographic Information Science* (pp. 1–6). Columbus, Ohio.
- Weber, R. (1991). Modelling Fire Spread Through Fuel Beds. *Progress in Energy and Combustion Science*, 17(1), 67–82.
- Wein, L. M., Craft, D. L., & Kaplan, E. H. (2003). Emergency Response to an Anthrax Attack. *Proceedings of the National Academy of Sciences of the United States of America*, 100(7), 4346–51.
- Weinberger, S. (2011). Web of War: Can Computational Social Science Help to Prevent or Win Wars? *Nature*, 471(7340), 566–8.
- Wellman, B. (1992). Which Types of Ties and Networks Provide What Kinds of Social Support? In *Advances in Group Processes* (pp. 207–235). JAI Press Inc.
- Wellman, B., & Wortley, S. (1990). Different Strokes from Different Folks: Community Ties and Social Support. *American Journal of Sociology*, 96(3), 558–588.
- Weng, J., Lim, E.-p., Jiang, J., & He, Q. (2010). TwitterRank. In *Proceedings of the third*

- ACM international conference on Web search and data mining* (p. 261). New York, New York, USA: ACM Press.
- West, G. (2013). Big Data Needs a Big Theory to Go With It. *Scientific American*.
- Wheaton, W. D., Cajka, J. C., Chasteen, B. M., Wagener, D. K., Cooley, P. C., Ganapathi, L., & Roberts, D. J. (2009). Synthesized Population Databases: A US Geospatial Database for Agent-Based Models. *RTI Press, 2009*(10), 905.
- Wiebe, J., & Riloff, E. (2005). Creating Subjective and Objective Sentence Classifiers from Unannotated Texts. In A. Gelbukh (Ed.), *Computational Linguistics and Intelligent Text Processing* (pp. 486–497). Springer Berlin Heidelberg.
- Wiebe, J. M. (1990). Identifying Subjective Characters in Narrative. In *Proceedings of the 13th Conference on Computational Linguistics* (Vol. 2, pp. 401–406). Morristown, NJ, USA: Association for Computational Linguistics.
- Wiebe, J. M., Bruce, R. F., & O’Hara, T. P. (1999). Development and Use of a Gold-standard Data Set for Subjectivity Classifications. In *Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics on Computational Linguistics* (pp. 246–253). Morristown, NJ, USA: Association for Computational Linguistics.
- Wilensky, U. (1997). *Traffic Grid*. Retrieved March 1, 2014, from <http://ccl.northwestern.edu/netlogo/models/TrafficGrid>
- Wilson, R. D., & Cales, B. (2008). Geographic Information Systems, Evacuation Planning and Execution. *Communications of the IIMA, 8*(4), 13–30.
- Wilson, T., Wiebe, J., & Hoffmann, P. (2005). Recognizing contextual polarity in phrase-level sentiment analysis. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing* (pp. 347–354). Morristown, NJ, USA: Association for Computational Linguistics.
- Wilson, T., Wiebe, J., & Hwa, R. (2004). Just How Mad are You? Finding Strong and Weak Opinion Clauses. In *Proceedings of National Conference on Artificial Intelligence*. San Jose, California.
- World Economic Forum, & Vital Wave Consulting. (2012). *Big Data, Big Impact: New*

- Possibilities for International Development* (Tech. Rep.). World Economic Forum.
- Wu, F., Huberman, B. a., Adamic, L. a., & Tyler, J. R. (2004). Information Flow in Social Groups. *Physica A: Statistical Mechanics and its Applications*, 337(1-2), 327–335.
- Wu, S., Hofman, J. M., Mason, W. a., & Watts, D. J. (2011). Who Says What to Whom on Twitter. In *Proceedings of the 20th International Conference on World Wide Web* (p. 705). New York, New York, USA: ACM Press.
- Wu, S., Shuman, L., Bidanda, B., Kelley, M., Sochats, K., & Balaban, C. (2008). Agent-based Discrete Event Simulation Modeling for Disaster Responses. In J. Fowler & S. Mason (Eds.), *Proceedings of the 2008 Industrial Engineering Research Conference* (pp. 1908–1913). Vancouver, British Columbia, Canada.
- Wu, Y., Zhang, Q., Huang, X., & Wu, L. (2009). Phrase Dependency Parsing for Opinion Mining. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing* (pp. 1533–1541). Singapore.
- Xiong, F., Liu, Y., Zhang, Z.-j., Zhu, J., & Zhang, Y. (2012). An Information Diffusion Model Based on Retweeting Mechanism for Online Social Media. *Physics Letters A*, 376(30-31), 2103–2108.
- Yang, S. H., Long, B., Smola, A., Sadagopan, N., Zheng, Z., & Zha, H. (2011). Like like alike Joint Friendship and Interest Propagation in Social Networks. In *Proceedings of WWW 2011* (pp. 537–546). Hyderabad, India.
- Yardi, S., & Boyd, D. (2010). Tweeting for the Town Square: Measuring Geographic Local Networks. In *Proceedings of fourth international AAAI conference on weblogs and social media* (pp. 194–201). Washington, DC.
- Yassemi, S., Dragičević, S., & Schmidt, M. (2008). Design and Implementation of an Integrated GIS-based Cellular Automata Model to Characterize Forest Fire Behaviour. *Ecological Modelling*, 210(1-2), 71–84.
- Yazici, M. A., & Ozbay, K. (2008). Evacuation Modelling in the United States: Does the Demand Model Choice Matter? *Transport Reviews*, 28(6), 757–779.

- Yu, H., & Hatzivassiloglou, V. (2003). Towards Answering Opinion Questions. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing* (Vol. 10, pp. 129–136). Morristown, NJ, USA: Association for Computational Linguistics.
- Zepeda, C., Osorio, M., & Sol, D. (2005). Modeling Evacuation Planning Using A-Prolog. In *15th International Conference on Electronics, Communications and Computers* (pp. 292–297). Puebla, Mexico: IEEE.
- Zepeda, C., & Sol, D. (2007). Evacuation Planning using Answer Set Programming: An initial approach. *Engineering Letters*, 15(2).
- Zerger, A., & Smith, D. I. (2003). Impediments to Using GIS for Real-time Disaster Decision Support. *Computers, Environment and Urban Systems*, 27(2), 123–141.
- Zhang, B., Chan, W. K., & Ukkusuri, S. V. (2009). Agent-Based Modeling for Household Level Hurricane Evacuation. In M. D. Rossetti, R. R. Hill, B. Johansson, A. Dunkin, & R. G. Ingalls (Eds.), *Proceedings of the 2009 Winter Simulation Conference* (pp. 2778–2784). Austin, Texas.
- Zickuhr, K., & Smith, A. (2012). *Digital Differences* (Tech. Rep.). Pew Research Center.
- 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), 6–32.

Curriculum Vitae

Sarah Wise is originally from the Northern Virginia area. She graduated from Thomas Jefferson High School for Science and Technology in 2005. In 2009, she graduated from the University of Chicago with a B.A. in Computer Science (with honors) and East Asian Language and Civilization. She entered a PhD program in Computational Social Science at George Mason University in 2009, and was awarded the George Mason University Presidential Scholar grant, which provided her with full funding for tuition and living costs.

During her undergraduate studies, Sarah worked at Argonne National Laboratory as a technical researcher, developing agent-based models from 2007 to 2009. She also worked for the United States Department of State Office of the POLAD (Political Advisor) Coordinator during 2007. Continuing with her computational research, she worked at SAIC during 2010, exploring social networks and algorithm design. Sarah worked at the Mitre corporation during both 2011 and 2012, again developing agent-based models of pedestrian movement and crisis scenarios.

Over the course of her graduate research, Sarah was awarded funding to attend the Santa Fe Institute Complex Systems Summer School as well as numerous conferences where she presented her research. She has published two journal papers and a book chapter, and has served as a reviewer for multiple journals.

Sarah is a candidate for the PhD degree in Computational Social Science from George Mason University in May 2014.