

VISUALIZING GEO-TEMPORAL
DOCUMENTS: AN APPLICATION TO
DATA FROM CRISIS MAPS

HINA AMAN

*A thesis submitted to the Faculty of Graduate Studies of
the University of Manitoba
in partial fulfilment of the requirements of the degree of*

MASTER OF SCIENCE

Department of Computer Science
University of Manitoba
Winnipeg, Manitoba, Canada

Copyright © Hina Aman

ABSTRACT

Crowd-sourced crisis mapping is a relatively new phenomenon and platform that enables the collection and visualization of real-time crisis data submitted by users through social media tools and cellular technologies. Crisis maps are generally used by both state and non-state actors for sense-making and as a reference point for action. The current crisis map visualizations only show the location documents such as reports or short messages have been generated from. Such a limited representation fails to immediately show important content, such as themes from a document and their changes over time. As a result, sense-making becomes time-consuming and cognitively demanding. I present a set of visualization tools: Geo-Temporal Tag Visualization (GTViz), Geo-Temporal Pies and Geo-SparkClouds that treat the tags on the crowd-sourced reports as spatio-temporal textual datasets and provide interaction tools to explore the content of the reports. I also demonstrate the value of such tools with case studies and a controlled user study.

PUBLICATIONS

Content, ideas and figures in this thesis have appeared previously in the following publications by the author:

Hina Aman, Pourang Irani and Hai-Ning Liang. A Review of Common Tasks Supported by Information Communication Technology for Times of Emergency. In *Proceedings of the 9th international conference on Information Systems for Crisis Response and Management (ISCRAM)*, 2012.

Hina Aman, Pourang Irani and Fereshteh Amini. Revisiting Crisis Maps with Geo-Temporal Tag Visualization. In *Pacific Visualization Symposium (PacificVis)*, 2014.

ACKNOWLEDGMENTS

It has been a wonderful and a very important experience being a part of the Human Computer Interaction (HCI) Lab. Therefore, I am deeply grateful to my supervisor, Dr. Pourang Irani, for his unwavering support, guidance and for giving me the opportunity to immerse in such a creative environment.

I have received generous support from Fereshteh as well. She has been a lovely mentor, a friend and a kind coffee sharer. I also appreciate the feedback offered by my committee members: Dr. Carson Leung and Dr. Saumen Mandal. Plus, I owe a special thank you to Srikanth, Roi, Khalad, Juan, Barrett and Amir.

The staff at International Students Centre helped me find a home away from home. Gitan and Julia (Ossome), I do not know how I will ever be able to thank you enough for your kindness and friendship and for always being there. If I could, I would gift you a private jet so that you'd visit me often.

Casual encounters in the engineering building led me to a beautiful friendship with Masoumeh and Ameneh - the best roomies I ever had! Thank you for the persian songs and the persian tea. I will forever hold dear the memories of the time we spent together.

I would also like to thank Asagi, Yue, Miles, Alaleh, Safoura, Marcelle, Choo, Naghmeh, Jose, Bhannu, Eid, Phoebe and everyone else (really!) for being such lovely friends. Always fun hanging out with you and it definitely made the winters more tolerable.

Sara, thank you for your brilliant ideas and sketches. I feel very lucky to have met you as I know I can always bank on your creativity.

MOST IMPORTANTLY, I'm deeply grateful to my parents for helping me at every step and for being so amazing. This research would not have been possible without their help and I really would never be able to express the immense amount of respect and love I have for them and for the rest of my family. Azhar, Asad and Mun-sif: Thank you for all the interesting conversations about quantum physics, pulsars, dark matter and what-nots. You guys rock! Nabiha, Lubna, Ubed and Haniya: I miss your company and think about you every day. Thank you for all your help and for being such a fantastic family.

CONTENTS

1	Introduction	1
2	Communication Technologies in Times of Emergency	7
2.1	Methodology	9
2.2	Communication	11
2.3	Crisis Maps	13
2.4	Event Detection and Assessment	14
2.4.1	Floods	15
2.4.2	Earthquakes	16
2.4.3	Terrorism	16
2.5	Warning	17
2.6	Training	18
2.7	Navigation	18
2.8	Evacuation	19
2.9	Conclusion	19
3	Related Work	22
3.1	Tag Clouds on Maps	22
3.2	Temporal Word Clouds for Text and Trend Analysis	24
3.3	Tag Clouds for Document Comparison	25
3.4	IEEE VAST 2011 Mini Challenge	26
4	Design Rationale	29
4.1	Design Goals: Making Sense of it All	29
4.2	The Story in The Data: Choosing What to Present	31
4.2.1	KWIC (Key-Word in Context) Method	32
4.2.2	Metadata Method	34
4.3	Designing the Tool: Data Visualization	35
4.3.1	First Iteration: Floating Windows	35
4.3.2	Second Iteration: Unfolded Dots	36
4.3.3	Final Iteration: Geo-Temporal Tag Visualization, Geo-Temporal Pies and Geo-SparkClouds	37
5	Case Studies	44
5.1	Exploring the Pakistan Flood Map	44
5.2	Exploring the Atlanta Crime Map	46

6	User Study	48
6.1	Experiment Design	48
6.1.1	Tasks and Datasets	48
6.1.2	Study Design	50
6.1.3	Apparatus and Procedure	51
6.1.4	Participants	52
6.2	Results	52
6.2.1	Task Time	52
6.2.2	Error	55
6.2.3	Subjective Preferences	56
7	Future Work and Conclusion	57
7.1	Limitations and Future Work	57
7.2	Conclusion	58
	Bibliography	68

LIST OF FIGURES

- Figure 1 Crisis Maps in action: A) Sahana Eden Pakistan Flood Situation Map B) Japan Quake Ushahidi Crisis Map C) Google Crisis Map for Superstorm Sandy. 3
- Figure 2 Crisis Reports from Pakistan in the aftermath of the 2010 floods as visualized (a) by UShahidi, a popular Crisis Mapping platform and (b) by Geo-Temporal Tag Visualization. The reports from the surrounding regions are aggregated into one visualization as it displays data over a twelve month period. Users can interact with the compact timeline getting exact values for each month. 5
- Figure 3 Classification of the corpus based on the eight identified tasks in the review. 11
- Figure 4 (a) A screen capture of Topigraphy [23] around the word "Cooking" (b) A screenshot of Trendsmap [6] showing the most talked about topics on the social media. 23
- Figure 5 (a) SparkClouds showing top words [39] and their trends over time. (b) LeadLine [6] visualizing events related to Steve Jobs, an American entrepreneur and former CEO of Apple. 25
- Figure 6 A screen capture of Scatterblog [9] showing key terms on the map. 28
- Figure 7 A screen capture of epSpread [4] showing origins of reports on the map. A tag cloud summary is presented separately. 28
- Figure 8 The data analysis process starts by first aggregating the reports according to their region. The meaningless words in the reports (highlighted in red in the example) are filtered out with a stoplist. A human expert then selects important words in context for visualization. 34

- Figure 9 (A) First iteration of visualization showing all the information on the map. (B) First iteration of visualization showing information distributed on the map (left) and activated windows showing the details (right). 41
- Figure 10 An early version of the main component of Geo-Temporal Tag Visualization showing the frequency of words listed on the left. The number in the center shows the number of reports received in that time period. 42
- Figure 11 Geo-Temporal Tag Visualization providing information on reports received from Atlanta's downtown area. The x-axis encodes time (Jan 2010 to Dec 2010) and the y-axis shows the number of reports filed under each tag/category. From the figure, we note that most of the reports were submitted in October tagged under "Larceny". (Best seen in color) 42
- Figure 12 The toolbar for filtering data. Users can either select a month or define a time period to view the top tags from the specified time. The zoom slider controls the line charts. If there are more than 10 words displayed, the toolbar activates more buttons to allow the user to navigate through the list by using First, Prev, Next and Last buttons. The Clear Charts button will clear the line charts. The Default button restores the visualization to its original state. 43
- Figure 13 Geo-Temporal Pies visualizing a year's data. Each month has its own pie. Users can click words to add or remove them from the pie charts. 43
- Figure 14 Based on SparkClouds [39], Geo-Sparkcloud visualizing a year's data. Users can interact with the line charts to get the month and the exact values. 43
- Figure 15 Geo-Temporal Tag Cloud visualizing the data collected in the aftermath of Pakistan Floods in the year 2010. The floods hit the country in late July. (Best seen in color) 45

Figure 16	Geo-Temporal Tag Visualization showing the crime data collected in the city of Atlanta between the years 2009 and 2013. (Best seen in color)	46
Figure 17	A screenshot of the experiment system showing the visualization and the task to perform.	51
Figure 18	Mean task completion time for each visualization.	53
Figure 19	Mean task completion time for each task.	53
Figure 20	Mean task completion time for each task type.	54
Figure 21	Mean task completion time for each group.	55

LIST OF TABLES

Table 1	The nine tasks for the experiment.	49
---------	------------------------------------	----

ACRONYMS

ANOVA	Analysis of Variance
GIS	Geographic Information System
GPS	Global Positioning System
GTP	Geo-Temporal Pies
GTS	Geo-SparkClouds
GTViz	Geo-Temporal Tag Visualization
HCI	Human Computer Interaction
ICT	Information Communication Technology
IEEE VAST	The IEEE Conference on Visual Analytics Science and Technology
ISCRAM	Information Systems for Crisis Response and Management
KWIC	Key-Word in Context
MEA	Mobile Emergency Announcements

INTRODUCTION

Geo-temporal documents consist of textual data generated from a geographical or spatial location and change over time. Any visual overview of geo-temporal textual documents is not complete if the geographical attributes of the documents are overlooked. For instance, a visualization of textual documents coming in from crisis zones will give an incomplete overview of the situation if the geographical attributes are ignored. Similarly, visualizing trends of content change over time is crucial to the understanding of these vast data repositories. Researchers have developed several visualizations to provide an overview of one or more documents, changes in documents over time and differences in documents. However, very little research exists on visualizations of geo-temporal documents.

We take the domain of Crisis Mapping as a real life example where visualization of geo-temporal documents is applied but with shortcomings in visualizations that are perceptible and can be addressed. During times of mass emergency, there are numerous reports coming in from the crisis hit areas. These reports are contributed by affected people, volunteers, or officials responsible for providing relief. Putting this geo-tagged data on maps is referred to as Crisis Mapping, which provides a platform for sense making and resource

management. Crisis Mapping is often provided as a part of emergency management and decision making tools.

In recent years, Crisis Mapping platforms like UShahidi¹ and Sahana Eden² have proven remarkably popular as tools to collect information from the public using multiple channels, including SMS, email, Twitter [32], and the web [44]. One popular and common use of such platforms is to aid in emergency management by collecting, classifying, and possibly verifying information contributed by affected people in a crisis zone or by volunteers from different parts of the world leading to a truly globalized relief process. The UShahidi Haiti Project set up after the devastating 2010 earthquake, for example, had several hundred volunteers contributing information, which, among others, was used by the US Coast Guard, Joint Task Force Command Centre and the US Marine Corps to bring aid to the Haitians and "save hundreds of lives" [46].

An independent evaluation of the UShahidi Haiti Project suggests that the most common use of this platform was for situational awareness [49], defined as a collective or individual understanding of the current situation at a high level. Review of selected Crisis Mapping initiatives by Cavelti and Giroux [15] also shows that the maps are generally used to inform and increase situational awareness. Ortman et al. [52] surveyed disaster management experts about their familiarity and experience, among other things, with crowd-sourced data, mapping services and social media during times of emergency. The results show that most of the experts are familiar with the new technologies and use them as well. They recognize the benefits of

¹ <http://ushahidi.com>

² <http://eden.sahanafoundation.org>

crowd-sourced data despite the verification and validation issues and mention a need for ranking, filtering and efficiently having access to the information hidden in the sheer amount of data.

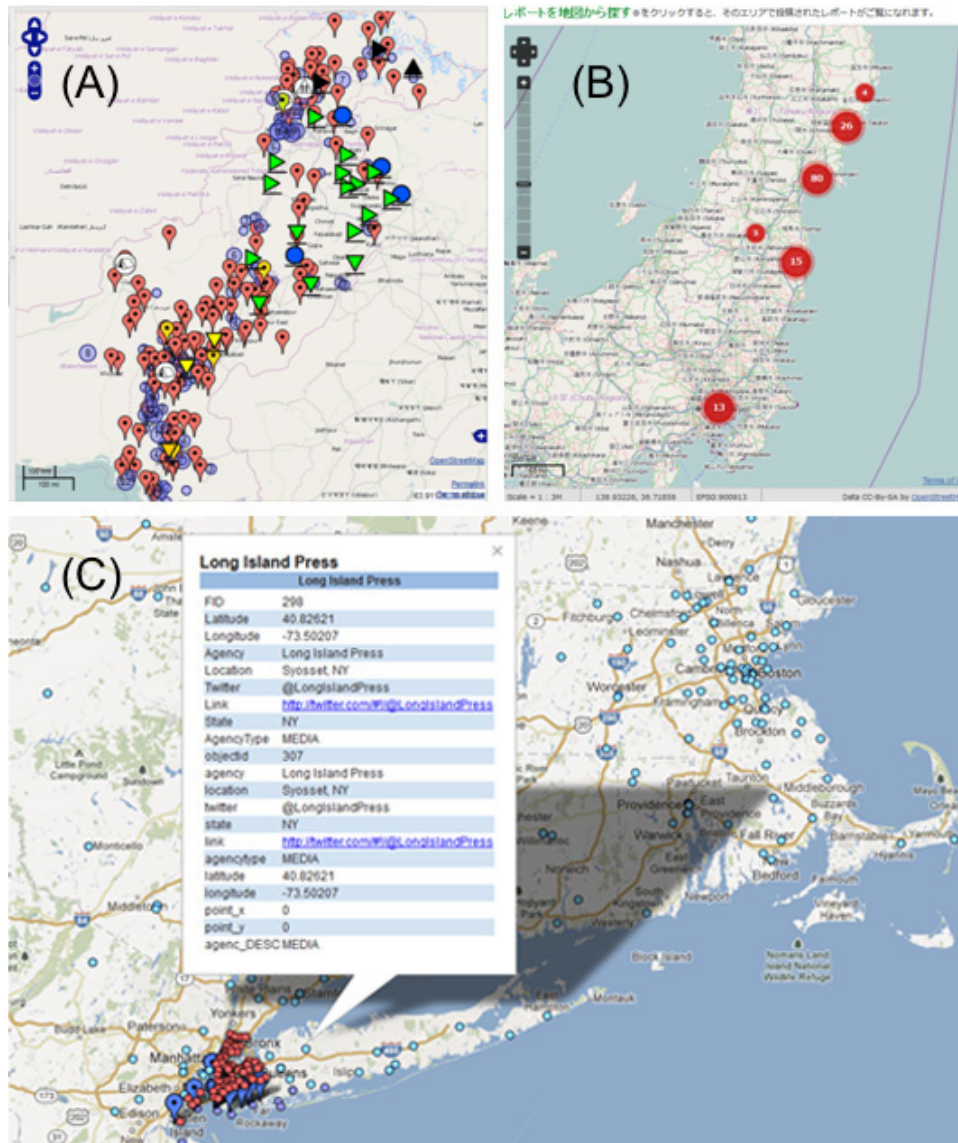


Figure 1: Crisis Maps in action: A) Sahana Eden Pakistan Flood Situation Map B) Japan Quake Ushahidi Crisis Map C) Google Crisis Map for Superstorm Sandy.

Studies similar to the one above also suggest that crisis maps popularly serve as a platform to support sense-making activities

undertaken to understand crises better [42]. However, the current visualizations utilized by common Crisis Mapping platforms are not optimized to support this task. Please see Figure 1 on page 3 for snapshots of visualizations offered by these platforms. Such visualizations are effective in showing where reports emanate from and/or how many, but do not provide an immediate sense of the overall picture, i.e. a glance-view, of the content and how it evolves over the length of the crisis.

In my thesis, I tackle the challenge of providing better visualization tools for geo-temporal data, which the field of Crisis Mapping can benefit from. I present Geo-Temporal Tag Visualization (GTViz) and Geo-Temporal Pies (GTP), two new visualization techniques that use the metadata of the reports to provide a summary of the reports' contents and the changes over time while preserving the geographic information. I also present Geo-SparkClouds building upon SparkClouds [39]. GTViz, GTP and Geo-SparkClouds extend the concept of Word or Tag Clouds [73], a weighted list of words widely used to visualize tags on the web by news websites for highlighting the most popular news or story themes, social bookmarking sites for popular topics, and for easy information retrieval and categorization [28]. Existing visualization of tags do not incorporate the spatial and temporal elements of geo-temporal tags. Very few visualizations exist which consider the temporal elements of tags, LeadLine [21] and SparkCloud [39], and even fewer application specific visualizations have the spatial element included, Trend Maps [6] and TwitterVision [69].

My main contributions are:

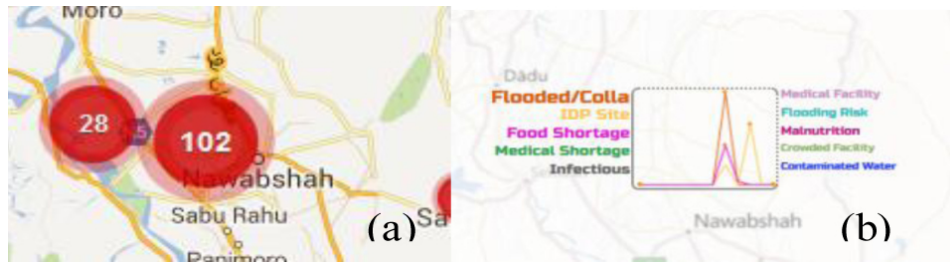


Figure 2: Crisis Reports from Pakistan in the aftermath of the 2010 floods as visualized (a) by UShahidi, a popular Crisis Mapping platform and (b) by Geo-Temporal Tag Visualization. The reports from the surrounding regions are aggregated into one visualization as it displays data over a twelve month period. Users can interact with the compact timeline getting exact values for each month.

1. Presenting a set of compact visualization tools: Geo-Temporal Tag Visualization, Geo-Temporal Pies and Geo-SparkClouds for summarizing texts using tag clouds. These visualizations represent all the three properties of a geo-temporal textual document dataset: content, area of origin and time of origin. This compact integration of a word cloud with a line chart is placed on the map at the location of its origin as shown in Figure 2. A user is also provided several interaction techniques to filter and analyze the data. These details are discussed in Chapter 4, the design section.
2. Conducting two case studies to demonstrate the value of my proposed visualization techniques. The case studies use the data collected from 2010 Pakistan Floods and Atlanta Crime Reports and highlight the interesting trends.
3. Performing a user study to evaluate all the visualizations. My findings show that users preferred to use Geo-Temporal Tag

Visualization, which also performed better in trend analysis and topic comparison tasks.

The chapters of this thesis are structured as follows: In chapter 2, I review how technology is used during times of emergency. I highlight some shortcomings and discuss communication technologies during times of crises in detail. In chapter 3, I discuss existing research work related to visualizing geo-temporal data. In chapter 4, I discuss the design rationale and details of my proposed visualization techniques. In chapter 5, I present two case studies applying Geo-Temporal Tag Visualization on real crisis maps. In chapter 6, I explain user evaluation of the visualizations and present the results. In chapter 7, I provide a conclusion and future work.

COMMUNICATION TECHNOLOGIES IN TIMES OF EMERGENCY

Content in this chapter also appears in [3], to which I contributed as the first author with professor Pourang Irani and Hai-Ning Liang.

As the first phase of my thesis, I conducted a research survey to:

1. Understand the current trends of technology use in times of emergency,
2. Get a sense of the Information Communication Technology (ICT) researchers' contributions in this field and shed light on how ICT research is being applied to real world problems, and
3. Identify emerging issues and overlooked problem areas.

The public and organizations' response to crises is changing due to the new developments in social media outlets, private communication devices, sensor and broadcasting technologies, which are influencing the way people seek, gather, process, and communicate information. However, we are yet many years away from utilizing the full potential of Information Communication Technology (ICT) in times of crises. To acknowledge and understand ICT's increasing role in crisis management, we need to know what types of ICT is available, how it is being used, and by what groups of users. This knowledge will inform the design and development of future ICT.

Recent years have seen an increase in the number of tools and systems available to emergency responders and to those affected by a crisis. This has been matched with a significant amount of research that discusses various aspects of emergency prevention, response and preparedness. However, because research on emergency and crisis management using ICT is interdisciplinary, the literature is spread across research areas making it hard to capture the full breadth of the research landscape and discern patterns from it.

A few studies have provided some review of specific research in emergency response. The 59th issue of *Participatory Learning and Action*, published by the International Institute for Environment and Development, discusses how web 2.0 technologies are crucial in development work and presents a few case studies of Web 2.0 tools being used for development work [5].

Addams-Moring et al. [1] survey early warning technologies available and provide a simple taxonomy for mobile emergency announcement systems. Palen and Liu provide an overview of how the public makes use of technology-mediated communication in the event of a crisis [55]. They mention the different characteristics of citizen-to-citizen communications, deriving observations from case studies of different types of crisis. My survey, on the other hand, takes a wider approach and aims to provide an overview of ICT used in common tasks associated with crisis events.

In this chapter, I first classify tasks common to all forms of crises. I employ this classification to organize relevant papers based on whether and how ICT presented in each paper supports the tasks.

I first report on our methodology used for this survey and then present the results.

I use the terms crisis, emergency, natural hazards, and disaster inter-changeably despite the subtle differences in their definitions. I refer to terrorism as disaster as well because terrorists, in the last two decades, have increased the scale of the attacks to maximize casualties and economical loss [77]. The damages incurred by recent terrorist attacks are very similar to those incurred by natural hazards.

2.1 METHODOLOGY

The first step in my approach was to collect papers touching on the breadth of technology in times for emergency. I distilled literature from proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI) and the Information Systems for Crisis Response and Management (ISCRAM) over the years. I also retrieved papers from other sources that were available through the ACM and SpringerLink Digital Libraries. I used several keywords, including: "emergency", "information systems", "disaster relief", "floods", "earthquake", "tsunami", "hurricane", "cyclone" and "fire-fighting". I went through the references of the papers retrieved to pull out more relevant works.

I collected over 100 papers which mentioned a certain aspect of technology use. Most of the papers touched upon the subject of communication, as shown in Figure 3. The reason for this is based on the fact that communication exists throughout all the stages of a disaster and the activities undertaken.

I analyzed each paper with an answer to each of the following questions:

1. What is the main motivation?
2. What methodology does the research employ (e.g., case study, ethnography, etc.)?
3. What types of emergency situations are addressed?
4. What are the tasks?
5. What technological tools and/or methods are used or presented for the tasks?
6. How are these tools and/or methods evaluated?
7. What are the results and conclusions?
8. What are the different terms used in the paper?

My review of the papers led to the identification of patterns of tasks and actions that would normally take place independently of the type or magnitude of the emergency situations. I grouped the papers according to the tasks they supported. See Figure 3 on page 11 for the distribution of works in each of these categories.

Further to the above categories, I consider field/ethnographic studies as a very important area of work as well. Technologists need to understand the behavior and the actions of first responders, volunteers or the affected to design effective and sustainable technologies. Despite a dearth of literature in this area, I came across some very valuable studies [38, 18, 82, 38]. However, I have excluded from review the results from field/ethnographic studies.

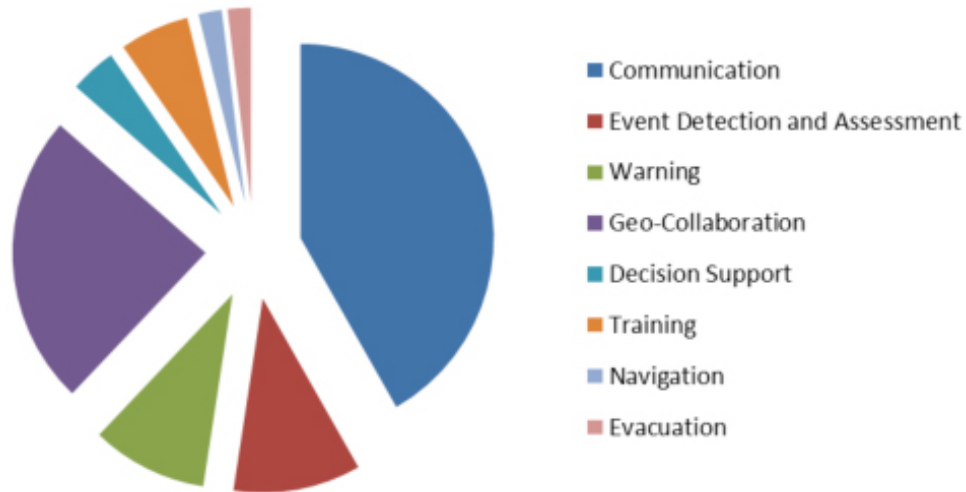


Figure 3: Classification of the corpus based on the eight identified tasks in the review.

2.2 COMMUNICATION

Release of the major micro-blogging platform Twitter, in 2006, brought about a key change in online communication by expanding the bounds of social networking. It was followed by an explosion of twitter web, desktop and cell phone clients like Twitterific, Seismic and Twihirl. The possibility of making twitter updates through cellphones particularly enabled people to broadcast information from crisis hit areas. Case studies of twitter usage during and in the aftermath of Oklahoma Grassfires - April 2009 [75], Red River Floods - March and April 2009 [75], Haiti Earthquake - January 2010 [82], Chile Earthquake - February 2010 [12] and Yushu Earthquake [56] suggest that micro-blogging can serve as an important platform for situational awareness and disseminating information about relief activities.

Following the mass use of twitter for relief activities, Twitter launched Hope140, a portal to highlight social movements and pro-

vide a starting point to interested volunteers. Hughes and Palen [33] provide preliminary evidence that users who adopt the micro-blogging technology in irregular emergency situations are more likely to become long-term users of the technology after finding a usefulness factor in it.

Such open communication raises several ethical issues involving the privacy and manipulation of those affected, which have not received much attention [76] [30]. Rumors are just as easily propagated. However, as Mendoza et al. note, the online community tends to question rumors more compared to verified news, a pattern which can lead to the identification of rumors through an aggregate analysis of collected data [47].

The introduction of micro-blogging and social networking technologies has led to the emergence of citizen journalism [48]. Leading news avenues are no longer the only popular source of updates on the world events. The internet has brought broadcasting to the fingertips of the general public. During times of emergency, individuals present at the location can provide updates. Apart from micro-blogging, photo sharing community websites like Flickr are used for photographic documentaries of disaster events [43]. Palen et al. provide a good overview of how citizens are using technology to participate in relief activities [55]. Many news agencies, like CNN, encourage the public to send in their stories as text, photographs or videos [63]. The trust factor is, of course, a concern and something requiring an innovative technological solution.

2.3 CRISIS MAPS

The current decade has seen technology opening new channels for information sharing, making mass communication accessible to members of the public and enabling them to contribute important information. These contributions not only aid emergency response teams and law-enforcing personnel but also allow the public to make informed decision. Such active public participation has led to a new vision of reframing emergency management as a socially distributed information system acknowledging citizenry as a significant force [54].

The most popular type of crisis communication today is through web-based Geographic Information System Collaboration. There are ongoing discussions on what to name the phenomena and activities which the new interactive online mapping tools enable. Some inter-related terms include Neogeography [51], Ubiquitous Cartography [24], Web Mapping [24], Collaboratively Contributed Geographic Information [8], and Volunteered Geographic Information [25]. The first three terms in the list stress on the cartographic representation being open to a new set of users (that is, the members of the public). The later emphasizes on shifts in the availability, and the process of creating, sharing and the utilization of the data. There are still debates surrounding the usage of the terms [14, 53].

The International Network of Crisis Mappers is a very large international online community of individuals with different professional backgrounds who engage in mapping during humanitarian crisis. An international conference for Crisis Mappers is held annually and

brings together professionals from different areas to assess the role of Crisis Mapping and technology in times of emergency.

No matter what the type of crisis is, GIS tools play an important role in the planning of relief activities. The origin of two popular open-source GIS collaboration tools, UShahidi and Sahana, are evidence to this fact. UShahidi was initially developed to map reports of violence due to poll rigging suspicions in Kenya following elections in 2008. Sahana, on the other hand, was developed in the aftermath of the 2004 Indian Ocean Tsunami [13]. The system has been deployed to track requests for assistance and information through several natural disasters, the most recent one being Japan's tsunami. A case study of active deployments is available on Sahana Foundation website .

One of the problems with GIS collaboration is the difficulty of keeping track of current situations. Crisis mappers post requests for helps but these requests are usually not followed up, and therefore one does not know if those requirements have been fulfilled [26].

Another difficulty of working with GIS, is due to the use of symbols. While working with maps, symbology becomes important. Because of the complicated nature of the maps and the types of purpose they serve, it is impractical to have common standardized symbols. Different maps use different symbols, which sometimes become hard to interpret [61].

2.4 EVENT DETECTION AND ASSESSMENT

There is an element of unpredictability in most types of disasters. Ideally, an event should be predicted accurately before it happens

but, in each case, an early detection can help rapid humanitarian responses which can result in saving of lives and infrastructure. Disastrous earthquakes can be distinguished from inconsequential ones within less than an hour after the event. This is important for rescue operations because for long periods (i.e., days) information usually does not flow from the hardest hit areas to the outside [80, 81].

Systems for prediction and detection of natural disasters need to be modeled in a way that they withstand the severe shocks. There is not much literature available that highlights these challenges. However, there are several websites which provide alerts and early warning regarding natural disasters.

2.4.1 *Floods*

A seamless way of monitoring floods is through remote sensing. Powerful data analyzing algorithms are employed to detect floods in real time. Global Disaster Alert and Coordination System (GDACS), a web-based public tool for monitoring floods, utilizes image processing algorithms to monitor floods at a global scale [19]. Light Detection And Ranging (LIDAR), among others, is popularly being used to generate Digital Elevation Models (DEMs), which can be used to simulate flood maps and provide a simulation for flood planning [11].

2.4.2 *Earthquakes*

There are several organizations that provide information about earthquakes as soon as one is detected. PAGER (Prompt Assessment of Global Earthquakes for Response [80]), World Agency of Planetary Monitoring and Earthquake Risk Reduction [81] and The Global Disaster Alert and Coordination System [19] are some examples of such organizations.

A novel approach in detecting earthquake uses humans as sensors [62]. When an earthquake occurs, people tend to update their status messages on social networks or write micro-blogs about it, which enables detection of an occurrence and propagation of an earthquake.

2.4.3 *Terrorism*

Detecting certain types of events (such as the presence of radioactive material) can pose a significant threat to security officials. Testing strategies to intercept such acts of terrorism can impose dangers to the protagonists involved. As a solution, some researchers are exploring publicly-available software like Second Life (a virtual world environment) to test strategies before introducing them in real life [79].

Researchers at UC Berkeley are currently developing small sensor packages that could be integrated into cell phone handsets to turn cell phone networks into sophisticated mobile sensor networks that can detect abnormalities like unusual traffic or bio-weapons. As an

interface to such system, WIPER is one example that can provide emergency planners and responders with a system that will help to detect possible emergencies as well as to suggest and evaluate possible courses of action to deal with the emergency [64].

2.5 WARNING

An efficient public warning system is essential to alert and instruct the general public about crisis and diverting them from danger. Most countries still use a siren as a basic warning system, which automatically rules out a population with hearing problems. Research conducted on a population showed that on average 37 percent of the population did not hear the siren and about 61 percent did not know what action to take when the siren was sounded [66]. Accessibility is an important aspect of an effective public warning system [10]. CAP-ONES is one example of an alert system which focuses on people with disabilities [45]. Nowadays, mobile phones are largely being used now to alert and instruct the public. These systems are referred to as Mobile Emergency Announcements (MEA). Addams-Moring et al. provide a simple taxonomy of existing MEAs [1], which distinguishes three types of such systems: pre-planned, ad-hoc, and semi ad-hoc. The key difference between these categories is whether or not they are known to the public and/or authorities.

2.6 TRAINING

The most important measure that can be taken to face any kind of emergency is to be prepared for it. A great deal of work is being done to train fire-fighters at the front-line through simulations as replicating a real fire-fighting scenario can be expensive and dangerous [7]. Apart from front-line action, trainings to facilitate communication and coordination have also received considerable attention. There is valuable literature available where researchers have worked alongside firefighters to better understand their nature of work, which would allow technologists to come up with tools to support front-line work and to design high-fidelity training simulations [20]. However, low fidelity training tools should not be ignored altogether given that they have been found to work just as well for such systems [68].

It is also very difficult to arrange training to intercept terrorist activities considering the risks posed especially if there are biochemical weapons. I did not come across any literature addressing this particular area.

2.7 NAVIGATION

Disasters like fires and floods often damage a great amount of infrastructure. It makes navigating to and from as well as within disaster areas very challenging, and also puts the life of the first responders at risk. In one of the aerial search and rescue missions during the Australian floods, a pilot had to improvise using an

iPhone app to navigate when the maps in the chopper were not of much use due to flooded roads [22].

Indoor reconnaissance missions taken by firefighters are very risky activities, which involve navigating through an unknown burning building with a limited supply of oxygen. Firefighters rely on their skills for way finding. Technology for this kind of job has to be very reliable. Ramirez et al. [57] show how they choose ambiguity as a tool to guide their design process of a simple way finding tool in an environment as rigid as firefighting.

2.8 EVACUATION

Mass evacuation planning requires an understanding of human behavior, which is likely to vary from place to place given different cultural norms. Many research groups are working on developing simulations that exhibit human behavior patterns in order to help practitioners plan effective evacuation techniques [35]. Moreover, evacuating a building is different from evacuating a flooded area. Documenting the different types of evacuations and field studies remain an open research area.

2.9 CONCLUSION

While a considerable amount of work has been done in crisis prevention and damage mitigation through the use of ICT, there is still further work ahead. Considering the damages that still incur due to

natural or man-made disasters, the objective of using technology to mitigate loss is far from having been achieved. Each year, we hear of floods and earthquakes causing a great deal of damage in terms of human lives and infrastructures. The geographic and economic condition of countries presents a greater challenge to technology. A majority of countries in the world are poor and do not have the resources to invest in advanced technologies to manage disaster risks.

I conclude my literature review by mentioning some open areas for research:

- There have been several types of terrorist activities noted, which could involve explosives, chemical or biological weapons. Training professionals to intercept terrorists involves a lot of risks that technology can help overcome by providing a virtual environment for training.
- Several online tools are available which help organizations and volunteers raise funds for disaster victims. These tools, however, are only limited to collection of funds. People who contribute may want (and have the right) to know how and where the funds are spent. Providing a platform to support such transparency can engage more contributors.
- Though online communication has accelerated the propagation of information, it has accelerated the propagation of rumors as well. Information validation and verification are two areas that deserve more attention especially if tools are being misused and draining tight resources allocated to the relief work.

- Maps are important visualization tool. Crisis Mapping by professionals and volunteers has become a mass and diverse but very important collaborative activity. User studies in this area can further improve the visualization and interaction tools offered by Crisis Mapping platforms. Moreover, we also see a need to improve the use of map symbology and provide a consistent vocabulary across different sites. Perhaps an international effort can be raised to ensure some consistent methods to provide user input to indicate specific needs during a crisis.

RELATED WORK

As part of the survey embodied in the previous chapter, I identified several open areas for research including the need for new and improved visualization of a large amount of geo-temporal documents or reports coming in from several sources. My thesis work attempts to address this need.

I draw inspiration from the various interfaces for visualizing documents, very few of which consider the visualization of documents with temporal and spatial properties. Tag clouds, which show the frequencies of the words in textual documents, have been applied in numerous contexts [60]. Our base representation builds on Tag Clouds, which we categorize based on various applications.

3.1 TAG CLOUDS ON MAPS

In a few instances Word Clouds have been applied to show content with a spatio-temporal dimension. Tag Maps [65], for example, puts representative textual tags (of photographs) on relevant map locations to highlight important concepts in the regional corpus. The visualization offered by Tag Maps is very simple. The textual tags are placed on the map but a higher density of text affects readability.

Topigraphy (topic + topography) [23] is a similar visualization on which words are displayed on a topographic image. Topigraphy is meant to help users intuitively understand the relationship between tags in a large-scale tag cloud and thus, making it easier for users to find the desired tag.

Trendsmap [6] is another tool that places words on maps. It is available online for users to analyze trends of conversations taking place on the internet. Trendsmap shows only the popular words, or topics, used in messages posted on Twitter. One can see the sizes of words or tags change in real-time as their popularity increases or decreases. The user has the option to focus on a single geographic area by zooming in. He can also specify topics or words of interest but he does not have the option to access the data of a previous time.

The above tag cloud based visualizations incorporate the geographic attributes of the data but they do not show the changes and trends of the tag cloud over time.

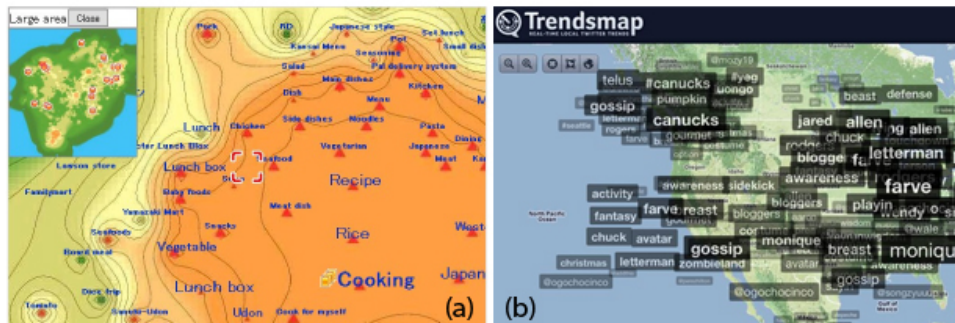


Figure 4: (a) A screen capture of Topigraphy [23] around the word "Cooking" (b) A screenshot of Trendsmap [6] showing the most talked about topics on the social media.

3.2 TEMPORAL WORD CLOUDS FOR TEXT AND TREND ANALYSIS

Word clouds with a temporal dimension have been explored to visualize topic trends and evolution. Themail [72] uses a column-based visualization to typographically present the progress of a conversation and relationship over time using the content of e-mails exchanged between individuals. Users can compare, with interaction, the contents of messages exchanged during different time periods. Individual messages are represented as colored circles in the background. The color is used to distinguish between sent and received messages.

SparkClouds [39] takes a more general approach towards text comparison. SparkClouds integrates Sparklines [70] into a word cloud to show trends between multiple word clouds. Each word is set on a compact spark chart that shows its frequency over time. With the resulting visualization, one can identify the most frequent words and their trend over a period of time.

In contrast to the above discussed methods, LeadLine [21] represents text corpora in the form of meaningful events. The visualization, which integrates word clouds for summarization, highlights events that trigger major changes in temporal trends.

CloudLines [37] is another visualization that shows events on a timeline with circles, the opacity and size of which shows the importance of the event it represents.

TIARA (Text Insight via Automated Responsive Analytics) [78] offers a visualization to aid users in exploratory text analytic tasks. TIARA presents content evolution over time by showing the derived

keywords as a word cloud on a chart, which represents a topic. The topics are presented as a layer on a chart with the height of each layer representing its strength over time. The visualization itself is very similar to ThemeRiver [29], which presents strengths of topics as currents of varying width flowing across time like a river.

TextFlow [17] is another visualization that uses the river analogy to present evolving topics in text. TextFlow focuses particularly on the evolution and relationships between multiple topics. It uses glyphs on the flow graphs (analogous to currents) to encode critical events. Visual exploration is enabled with a timeline and a word cloud view.

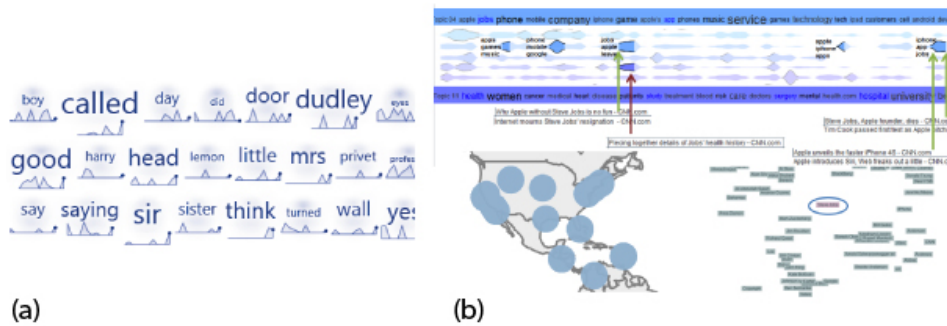


Figure 5: (a) SparkClouds showing top words [39] and their trends over time. (b) LeadLine [6] visualizing events related to Steve Jobs, an American entrepreneur and former CEO of Apple.

3.3 TAG CLOUDS FOR DOCUMENT COMPARISON

Studying trends in reports over time certainly adds to the understanding of a crisis situation. However, comparing the trends between different regions over time provides an additional, helpful insight. Tag clouds have been explicitly or implicitly employed in visualizations to demonstrate similarities between different texts. Rembold

and Spath [59] present an aesthetically pleasing radial space-filling visualization to allow comparison between essays. The visualization has different levels to represent different information about the essays including the authors, length of the essays, structure of the essays, the most frequently used words and the number of coinciding words. The final design embodies word clouds in the third level with each word wrapped in a circle proportional to the font-size to represent the word's frequency.

Parallel Tag Clouds (PTCs) [16] presents differences among facets of very large text corpora by combining parallel coordinates with word clouds. The words are arranged into parallel columns, each column representing a distinct subset of the data across a facet of interest. Common words are connected by nearest-neighbour edges, which make it easier to find coinciding words within the corpus.

IBM's Many Eyes [74] also allows comparison of two bodies of text in a single word cloud but like the other visualizations mentioned above, it does not provide a temporal advantage.

3.4 IEEE VAST 2011 MINI CHALLENGE

In 2011, VAST presented a mini challenge problem titled, 'Characterization of an Epidemic Spread' [71]. The challenge provided participants with two datasets. The first dataset contained microblog messages collected from various GPS capable devices, which meant that the exact location of the origin of the messages was known. The second dataset contained a detailed satellite image of a fictional metropolis, Vastopolis, where the first dataset was coming from.

The challenge was to identify the origin and spread of an epidemic outbreak using the datasets provided. The scenario presented is very similar to a real-life crisis situation where individuals from a crisis-hit area contribute micro messages or reports. We review two of the winning visualizations, the most relevant to our proposed visualization, developed and contributed as a solution to the challenge problem.

ScatterBlog [9] used Tag Clouds along with density hotspots on the map to support geo-spatial document analysis tasks. A 3d scatter plot was used to show messages over time, where color was used to highlight queried keywords in messages, which are plotted on the map as dots. epSpread [4] combined stream graphs, word clouds and map visualizations to provide a querying interface for document analysis. All these visualizations are presented as separate but linked components of the tool.

The visualizations proposed as a solution to this challenge were mostly aimed at analytical tasks often carried out by experts seeking a solution to a particular problem. With our proposed visualization, we focus on the general task of sense-making carried out by experts and non-experts alike. We provide an overview of all the textual documents along with information on their regions of origins and differences in content over time.

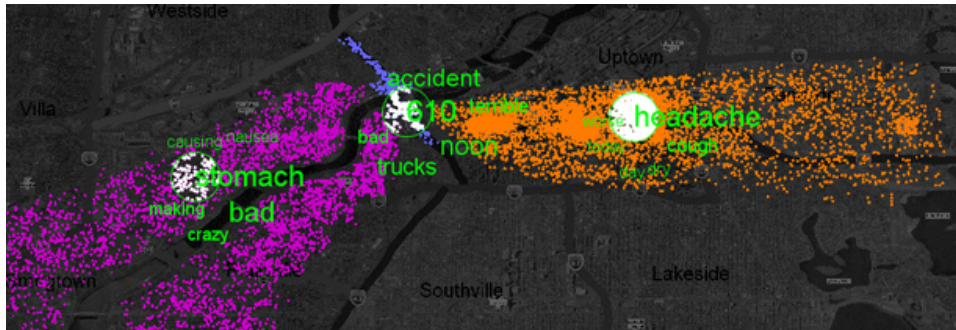


Figure 6: A screen capture of Scatterblog [9] showing key terms on the map.

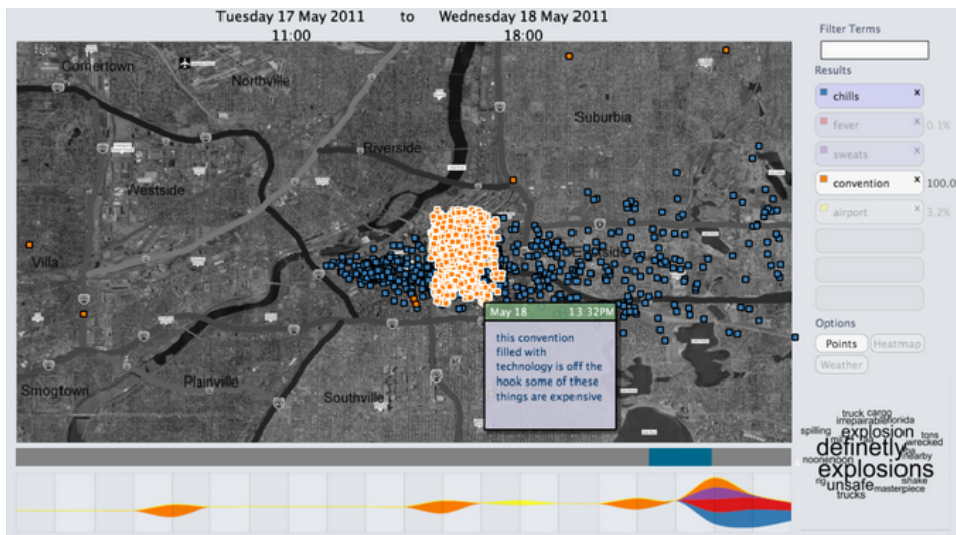


Figure 7: A screen capture of epSpread [4] showing origins of reports on the map. A tag cloud summary is presented separately.

DESIGN RATIONALE

Content in the following chapters also appears in [2], to which I contributed as the first author with professor Pourang Irani and Fereshteh Amini.

I designed the new set of visualization tools to provide a simple and understandable summary of the geo-temporal data at hand. I focused on presenting minimalistic designs that lowers the cognitive load and reduces the time consumed to understand the crisis situation being mapped. Therefore, I take inspiration from tag clouds considering their compact nature. The readability of tag clouds can easily be controlled with proper spacing and font-size choices [27]. I used line charts, which are very widely used and understood, to show the trends. Before describing the design, I present the usage scenario of sense-making in detail to provide context and motivate our design choices.

4.1 DESIGN GOALS: MAKING SENSE OF IT ALL

Several studies mention sense-making and situational awareness as one of the main uses of Crisis Maps [67, 50, 14]. A widely accepted definition of situational awareness is given as having knowledge about current data elements, or inferences drawn from these data elements. On the other hand, sense-making is described as the pro-

cess undertaken to achieve situational awareness [36]. I use these two terms interchangeably to refer to the process of inferring the conditions in a crisis-hit area by analyzing the crowd-sourced data (reports or messages) from that region.

Figure 1 on page 3 shows a screenshot of Crisis Maps, which were active following different natural disasters. The visualizations employed by these tools are analogous to the traditional non-technological method of placing push-pins on paper maps to mark a position. Icons, dots of varying sizes and numbers are used to encode information about the reports. A user is required to interact further with the visualization and view the contents of all or at least most of the reports to obtain increased situational awareness. Designers choose a simple representation to promote general legibility. However, this is achieved at the cost of not providing a general and glance-able overview of the reports' content.

To guide the design process, I came up with a list of questions about the data:

1. What are the reports in this region or Crisis Map about?
2. What is the most prominent topic/category of reports?
3. What was the most prominent topic/category in a particular month or time period?
4. How do the content of the reports change over time?
5. How do the reports from one area differ from another? What are the similarities?

6. How has one topic evolved compared to another in one or multiple areas?

The answers to the above questions provide a basic sense of the crisis data, which I aim to provide through the new visualization tools. To help answer these questions, I identified the following dataset elements to incorporate in the visualization design:

1. A list of areas (containing aggregated reports)
2. Time periods (of the reports received)
3. Number of reports received (in each time period from each area)
4. A list of words (tags) and their frequencies (for each time period from each area).

Apart from providing users with an overview of the above data, it was necessary to enable them to compare the data between two or more regions over different time periods by making the visualization interactive.

4.2 THE STORY IN THE DATA: CHOOSING WHAT TO PRESENT

The language available in the crisis reports depends a lot on their origin. For example, the reports from Japan following the tsunami were mostly in Japanese and the reports from Jakarta following torrential rains were in Bahasa Indonesian . This limited the number of datasets we could use due to unfamiliarity with the languages.

However, the visualization technique is applicable to text of any language, as long as key terms can be extracted from the content.

To design the prototypes, I obtained the dataset of 2010 Pakistan Flood Reports with approximately two thousand reports which I further filtered and considered only the ones verified by volunteers. This reduced the number of reports by approximately half, which I aggregated by area of origin. The details of the datasets and their implementations are explained in the next section.

I identified two possible ways of representing the textual content of the reports:

1. KWIC (Key-Word in Context) Method
2. Metadata Method

4.2.1 *KWIC (Key-Word in Context) Method*

The KWIC Method is based on doing a frequency count of selected words in the reports. Such a method seems appropriate in the context of crisis hit areas where most of the reports, as observed, carry request for resources. Many of these text messages received contain similar text, and a frequency count can shed some light on the most mentioned resources or concerns. Here are two example text messages¹ received from individuals seeking help during a crisis:

Lost everything in flood and looking for food items and issuance of watan card.

¹ <http://pakreport.org/ushahidi/reports>

We want some donation and food items (for unsettled areas) and some books from class 1-12 and bags (for settled areas).

The frequency count is preceded by aggregation of the reports according to their regions of origins. I prepared a list of stop-words comprising of articles, pronouns, conjunctions, prepositions, adverbs and interjections. However, a mere list of stop-words is not enough to filter out words of importance from a large corpus. Therefore, the method requires analyst intervention, as explicitly recommended by Rayson and Garside [58]. Given the nature and context-dependency of the crisis data, human expertise is better at identifying meaningful terms. The method adopts a standard corpus technique of KWIC (key-word in context). After doing a frequency count, the system lists all the words that made it through the stop list for the human analyst to select words deemed meaningful in the context. These words were added to the list of go-words. The analyst also has the option to merge words. This process is illustrated in Figure 8.

At the end of the process, we have a list of meaningful words from the corpus and their frequency from each area for each time period.

This is a very simple approach of extracting useful information from datasets but it is sufficient for our purpose. Frequency pattern mining algorithms and visualizations similar to FpViz [40] and FpMapViz [41] provide a direction for getting a richer insight about the data.

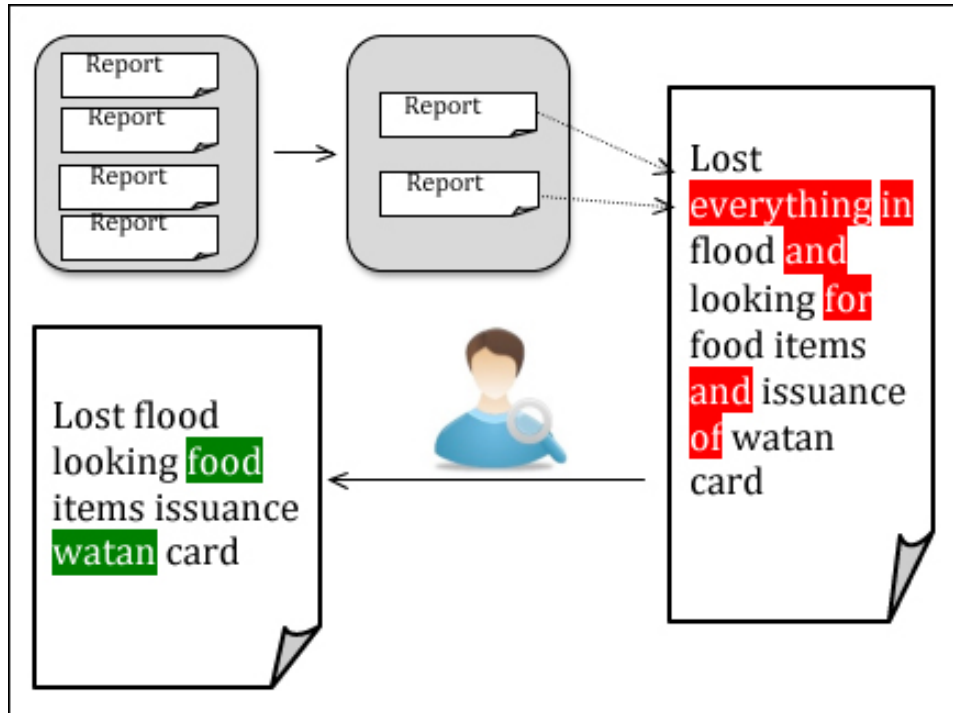


Figure 8: The data analysis process starts by first aggregating the reports according to their region. The meaningless words in the reports (highlighted in red in the example) are filtered out with a stoplist. A human expert then selects important words in context for visualization.

4.2.2 Metadata Method

The reports in the datasets we obtained had a time stamp, latitude, longitude and a tag for categorization. From the reports under the tag "Food", here is an example text message² received from individuals seeking help during a crisis:

We want some donation and food items (for unsettled areas)...

The tag appropriately sums up the content of the reports. The same is true for other reports. Therefore, I found it appropriate to use this

² <http://pakreport.org/ushahidi/reports>

metadata to represent the content of the reports in the visualization. The reports are grouped according to these tags/categories and aggregated according to the regions of their origin.

Since this method does not require any human intervention, I apply this on the case studies presented in the next section.

4.3 DESIGNING THE TOOL: DATA VISUALIZATION

It seemed (and it is) quite a challenge to visualize all the above components of data without overwhelming the user with details and without causing clutter. I opted to use a 2D space and keep the layout similar to the current visualizations. The final design is the result of an iterative process, the steps of which I will discuss in detail.

4.3.1 *First Iteration: Floating Windows*

In the first design iteration, I replaced the circle or icons on the map with a pie chart to represent the frequency of top words, distinguishable by their colors. Apart from a legend on the map showing the color-word association, the user could also hover over pie slices to know its representative word. To keep the layout simple, each pie chart represented only the top five words. User could access more information by clicking on the pie chart, which popped up an infobox with several tabs. Each tab offered a different visualization: a word cloud and a compound bar chart to show frequency of the

words and number of reports over time. I also experimented with the placement and portability of the infoboxes as shown in Figure 9.

In one version, the map occupied the entire screen area. Information in the documents was represented as pie-charts with associated infoboxes placed on the map when triggered. In second variation, we divided the screen into two areas. The map with pie charts on it occupied half of the area. The other half of the screen was meant for the infoboxes. This version enabled a more organized workspace for the user. Despite the ease of having part of the data in draggable infoboxes activated or made visible by a click, we felt the need for a more detailed visualization that allowed a glimpse of all the data without additional interaction.

4.3.2 *Second Iteration: Unfolded Dots*

In the second design iteration, I replaced the pie-chart with a line-chart enclosed in a circle, which shows the frequency of the top words (y-axis) over time (x-axis). The words being represented are placed on the side of the circle in an ascending order of their frequency as shown in Figure 10 . The frequency of the word is also encoded in the font-size of the words to enable comparison between words.

4.3.3 *Final Iteration: Geo-Temporal Tag Visualization, Geo-Temporal Pies and Geo-SparkClouds*

The final iteration saw many several changes in the design, interactivity support, and new visualizations to support more tasks. I will first discuss the properties of the final design, which will be followed by an introduction to the two new visualizations (Geo-Temporal Pies and Geo-SparkClouds), which share the same interaction tools.

Representing Content:

The content of the reports is represented by using a modified version of a tag cloud where the words are aligned in descending order of their frequency. This alignment along with the font-size makes it easier to identify the most frequently occurring themes in the reports. The font-size of the words aids in comparison of frequencies between words as well. Most tag clouds use a linear transformation to translate frequency or word weight into the font-size. I opted for the logarithmic algorithm because it distributes the font-size more evenly giving visibly improved results [31]. To make efficient use of the space and to maintain its even distribution among words, a part of the longer word is hidden and made visible only when the user hovers over it.

Each word is assigned a unique color, which links the word to its line chart. To select usable colors for the visualization, we referred to the work by Ichihara, et al. [34]. I selected colors from their recommended color palette to ensure that the chosen colors are easily distinguishable by people with and without color-vision deficiencies.

Representing Temporal Information:

In Geo-Temporal Tag Visualization, the temporal information is presented through a composite line chart. The scale of the chart can be adjusted. In this example, a point represents the aggregate result over a period of a month. Each word's line on the chart is distinguishable by its unique color. To reduce clutter, the line chart shows information about the top three words only. This default state can be changed by the user.

Geo-SparkCloud uses line charts as well for the temporal information. However, in Geo-SparkCloud, each word has its own line chart embedded above it. Geo-Temporal Pies represent each unit of time (day, week, month or year) with a pie chart enabling easy comparison of tags.

Representing Geographic Information:

The geographic information is entailed in the map on which the words along with their temporal information are placed. To improve readability of the words and the charts, I applied opacity on the map. This chosen opacity level chosen is enough to highlight the foreground while still maintaining the readability of the graphical background (i.e. the map).

Interactions on the Visualization:

The default view of the visualization shows a summary of the data collected over all periods of time. However, a user can interact with the toolbar, shown in Figure 12, to explore the datasets for a specific

month or a time period. The user can also click on a word to show or hide its line chart. Users can also control the bounds (x-axis: time) of the lines in the chart of the active words by using the zoom slider. The clear button will clear the line charts for all areas. The user can then click on words of interest to activate their charts. The selected word can be in any area but it will be included in the charts of all the areas where this word occurs. It provides the user with a convenient way of identifying the areas with similar content.

To get the values of the number of reports for each word or tag, the user can hover on the line chart or pie slices. It activates a tool tip which shows the exact values. The month is shown separately on the top of the chart. The Default button restores the default view of the visualization.

Scalability of the Visualization:

The visualization can be scaled to include a number of words. The user can browse between the words using navigation controls placed below each chart. The control under each chart is area specific. The control in the toolbox is global. Increasing the number of words, however, comes with a limitation. The greater the number of words, the more difficult it becomes to select a distinguishable color for each word. To include reports that span over the course of several years, we can simply add an additional year button bar along with a slider.

Geo-Temporal Pies:

I revisit the first iteration of the design with Geo-Temporal Pies. Each month (or any unit of time) has a small pie chart. The top five active

words are represented by their colors in each pie. Clicking a tag in the list will either add or remove its values from the pie depending whether or not the tag is already represented there. Hovering on the pie slices activates a tooltip. The tooltip contains the name of the tag the slice is representing and the exact frequency of the tag. The name of each month is printed above each pie.

Geo-SparkClouds:

SparkClouds [39] integrate sparklines (small line charts) into a tag cloud to present trends between multiple tag clouds. Given its ability to incorporate temporal data with words while still being compact, we put it on a map to evaluate its performance along with the other visualizations. I modified Geo-SparkCloud to maintain consistency in the visualizations.

Words are arranged in the order of their frequency. Users can get the exact values for each month by hovering on the small charts. The month is displayed on the top of the cloud. As with the other two visualizations, each word is assigned a space. The overflow of longer words is hidden but the entire word becomes visible on mouse-over.

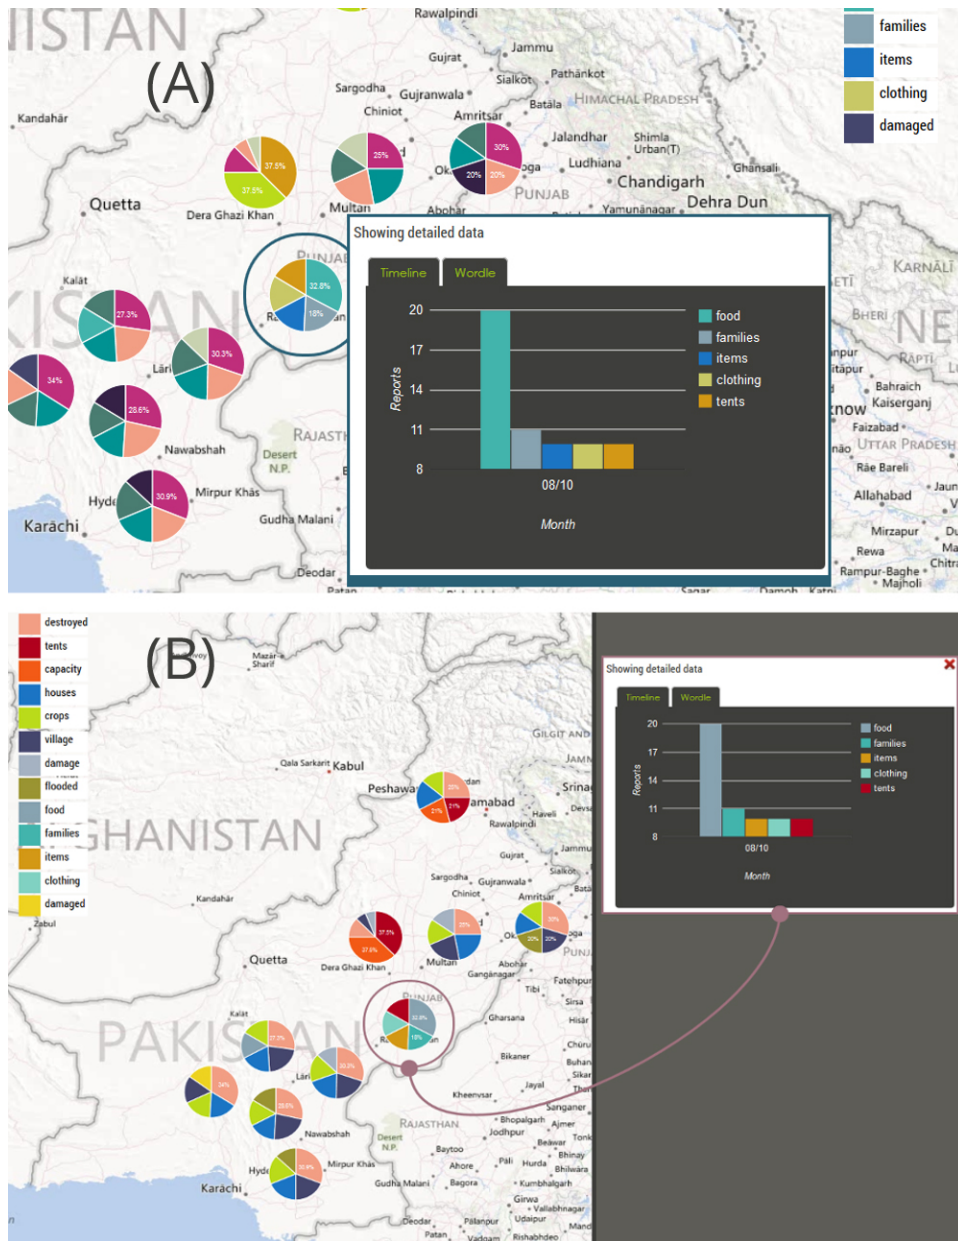


Figure 9: (A) First iteration of visualization showing all the information on the map. (B) First iteration of visualization showing information distributed on the map (left) and activated windows showing the details (right).



Figure 10: An early version of the main component of Geo-Temporal Tag Visualization showing the frequency of words listed on the left. The number in the center shows the number of reports received in that time period.

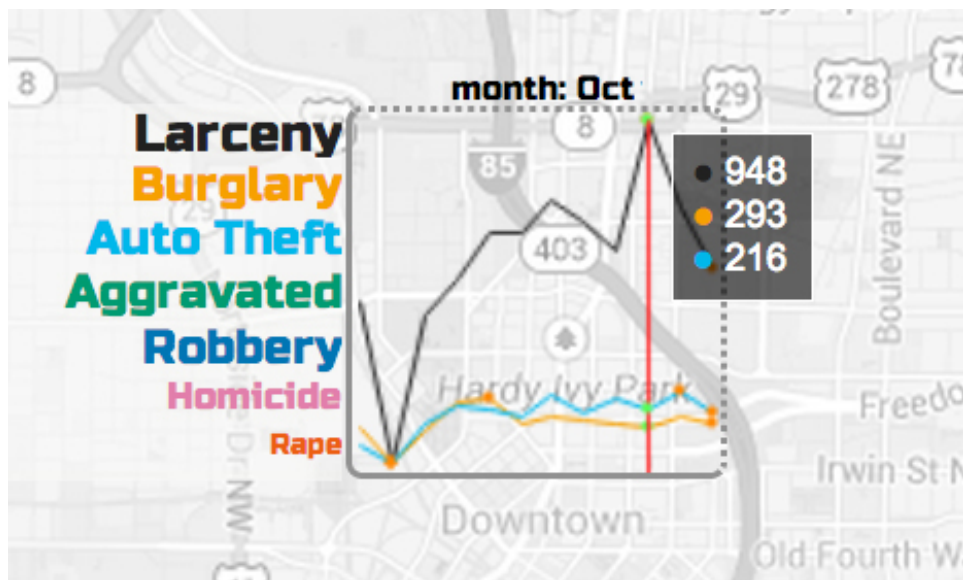


Figure 11: Geo-Temporal Tag Visualization providing information on reports received from Atlanta’s downtown area. The x-axis encodes time (Jan 2010 to Dec 2010) and the y-axis shows the number of reports filed under each tag/category. From the figure, we note that most of the reports were submitted in October tagged under "Larceny". (Best seen in color)

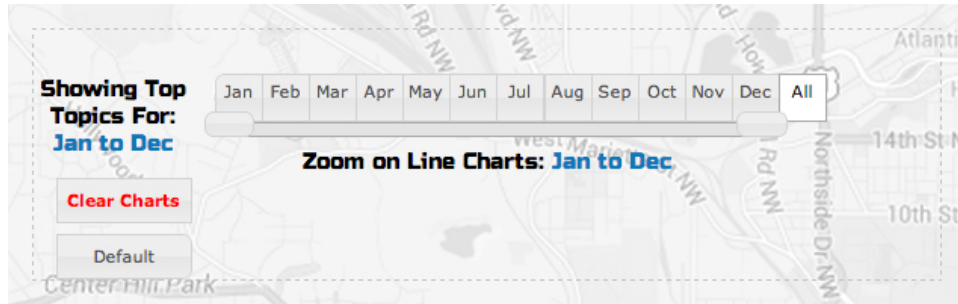


Figure 12: The toolbar for filtering data. Users can either select a month or define a time period to view the top tags from the specified time. The zoom slider controls the line charts. If there are more than 10 words displayed, the toolbar activates more buttons to allow the user to navigate through the list by using First, Prev, Next and Last buttons. The Clear Charts button will clear the line charts. The Default button restores the visualization to its original state.



Figure 13: Geo-Temporal Pies visualizing a year’s data. Each month has its own pie. Users can click words to add or remove them from the pie charts.

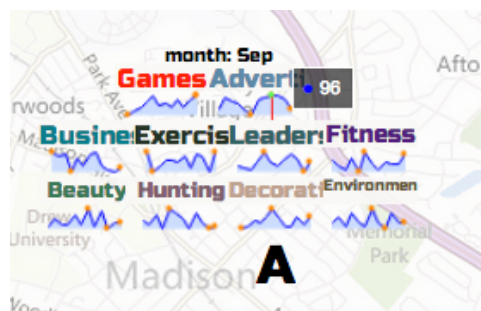


Figure 14: Based on SparkClouds [39], Geo-Sparkcloud visualizing a year’s data. Users can interact with the line charts to get the month and the exact values.

CASE STUDIES

I undertook two case studies to observe how well the visualization integrated with real-world datasets. I give a background on the datasets and the visualization processes before describing some of the patterns and phenomenon that the visualization highlights in each case.

5.1 EXPLORING THE PAKISTAN FLOOD MAP

Severe floods affected Pakistan in the month of July 2010 after torrential seasonal Monsoon rains. It killed an estimated 1,600 people and affected more than fourteen million people. Pakreport, a customization of UShahidi, was deployed in the beginning of August 2010 collecting data through several channels including SMS. The dataset I downloaded from Pakreport website had 2262 report entries. I used a subset of 753 verified entries.

Figure 15 provides an overview of the crowd-sourced crisis data of Pakistan Flood 2010 with Geo-Temporal Tag Cloud. In this section, I discuss a few things notable about the data by looking at this visualization.

Most of the reports submitted contained information about flooded areas and collapsed structures as well as sites for Internally Displaced

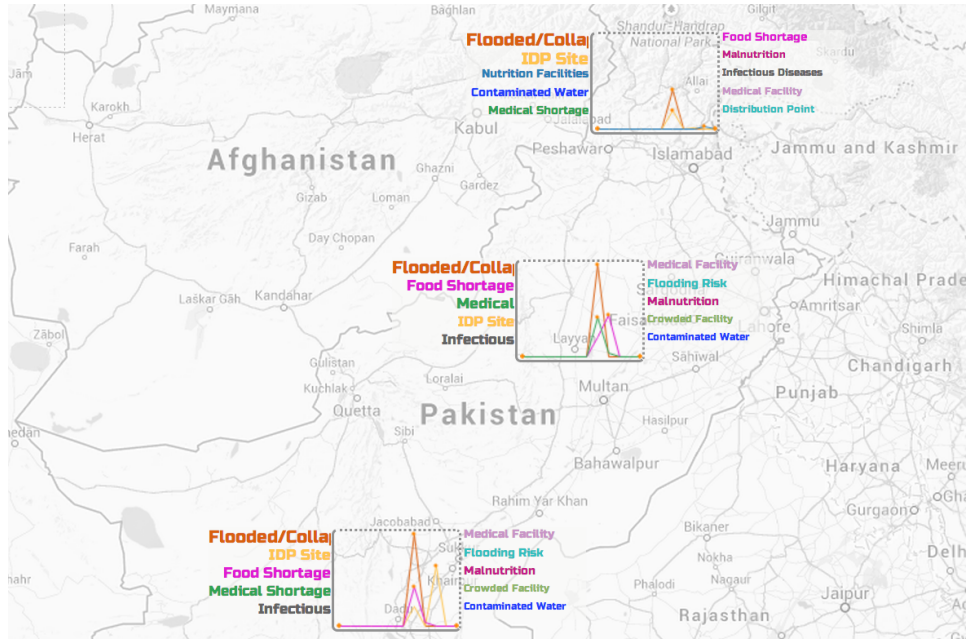


Figure 15: Geo-Temporal Tag Cloud visualizing the data collected in the aftermath of Pakistan Floods in the year 2010. The floods hit the country in late July. (Best seen in color)

Refugees (IDPs). Reports mention Food Shortage, Medical Shortage, Contaminated Water, Malnutrition and Infectious Diseases in almost every area. There is also information about Medical Facilities.

Most of the reports were submitted in the month of August, several weeks after the massive floods hit the country. A few more reports were submitted in September and October. The number of reports from the Southern part of Pakistan is greater compared to the Northern part. This does not necessarily indicate that the Southern part suffered most damage; however, we can safely say that people contributed more reports from or about the Southern region.

5.2 EXPLORING THE ATLANTA CRIME MAP

I obtained the second dataset from Atlanta Open Data Portal filed under Atlanta Crime Map. The data was collected between 2009 and 2013 and is provided by Atlanta's Police Department. This data file contains 146,569 reports. The year 2010 had a set of reports from all twelve months unlike the previous years, so we look at this subset of 35,767 reports in the visualization. Filtering the reports with missing records left us with 35,674 reports.

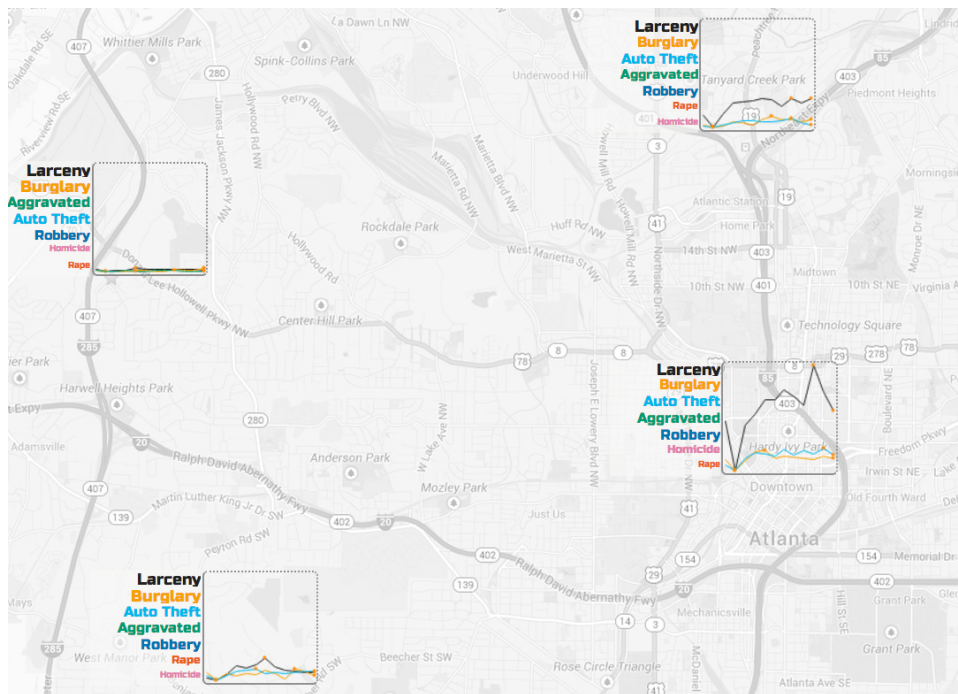


Figure 16: Geo-Temporal Tag Visualization showing the crime data collected in the city of Atlanta between the years 2009 and 2013. (Best seen in color)

The Atlanta Crime Map dataset contained crime reports submitted by the citizens to the Atlanta Police Department. The textual content of the records in the dataset was very compact and comprised of the

following tags describing the type of the incidents: Larceny, Burglary, Auto Theft, Aggravated Assault, Robbery, Rape, and Homicide.

One of the first details that one can observe on the Geo-Temporal Tag Visualization of the Atlanta Crime Map is that the number of incident reports is high in the downtown region with Larceny reaching a high peak in the month of October (bottom right in 16). In the same area, it is interesting to note that the numbers of incident reports are very low for the month of February. Other areas also have relatively fewer reports for the month of February.

Larceny, Burglary, Auto Theft and Assault seem to be the most common crimes committed all over the city. West of Atlanta (top left in Figure 16) saw fewer reports of crimes with the rate being almost constant throughout the year. This low rate is followed by the South of Atlanta (bottom left in Figure 16). Larceny peaks in this region in the month of July. The North of Atlanta (top right in Figure 16) sees the second highest number of reports. Similar to the downtown region, Larceny peaks here in the month of October as well.

USER STUDY

I conducted a controlled experiment to explore how efficiently (quickly and accurately) users could answer three different types of tasks by using Geo-Temporal Tag Visualization, Geo-SparkClouds and Geo-Temporal Pies.

6.1 EXPERIMENT DESIGN

6.1.1 *Tasks and Datasets*

I tested Geo-Temporal Tag Visualization (GTViz), Geo-SparkClouds (abbreviated as GTS for ease) and Geo-Temporal Pies (GTP) across nine different tasks under three categories. These tasks were based on possible questions about geo-temporal crisis data. Each question had the following three variables: location, time, and content with one of them unknown and two known. Table 1 summarizes the nine tasks.

Each of the nine questions had its own dataset including a map. I wrote a program to generate these datasets. The maps were included manually. The program randomly selected 10 words from the list of 72 words for each cluster and assigned them random frequencies for each month (for a total of twelve months). I manually edited the

#	Task	Task Type	Unknown
1	What topic ranked the most/least popular during the month M in the area A?	Specific	Topic
2	In which month was the topic T most/least popular in area A?	Specific	Time
3	In which area was the topic, T, most popular in the month M?	Specific	Area
4	Which topic saw a continuous increase/decrease in rank from Month A to Month B in area A?	Trend	Topic
5	Which area saw a continuous increase/decrease in rank of topic T from Month A to Month B?	Trend	Time
6	What time period did the topic T saw a continuous increase in its rank in area A?	Trend	Area
7	What topic had a rank higher than Topic T in area A during Month M?	Comparison	Topic
8	In which area was topic T with a rank higher than topic T during Month M?	Comparison	Time
9	Which month saw topic T ₁ with a rank higher than topic T ₂ in area A?	Comparison	Area

Table 1: The nine tasks for the experiment.

datasets to ensure that each trial had a unique answer. A synthetic dataset also allowed me to control the characteristics of the data for a very balanced dataset. I found that this was not possible with real datasets from existing Crisis Maps, which didn't contain all the features or events necessary for the target questions, and had greatly varying scales in time or space. It would have made the experimental conditions less well controlled, reducing our statistical power.

To reduce the complexity of the experiment, the questions were given in multiple choice format. The position of each answer was randomly generated as well.

6.1.2 *Study Design*

I ran a 3 (Visualizations) \times 9 (Tasks) within-subject study. Each participant performed all the tasks across the three visualizations. The order of the tasks was constant. The nine tasks were performed twice on each visualization resulting in 54 trials. I measured the accuracy and task completion time for each trial. The 54 trials were divided into nine blocks - one for each task. Before the beginning of each block, the administrator described the question to the participant. The participant then completed a practice session before moving to the real trials. At the end of each block, the user was asked to provide subjective feedback on the preference of visualization for the particular task. At the end of the experiment, the participants were asked to rank the visualizations according to the ease of use and general preference.

6.1.3 Apparatus and Procedure




Figure 17: A screenshot of the experiment system showing the visualization and the task to perform.

The user study was run on a 3.16 GHz Dell OptiPlex 910 computer with 16GB of RAM and a 23" Dell Flat Panel Monitor with a 1920×1080 pixel resolution. The visualization occupied 80% of the screen and the remaining 20% was reserved for the question and the multiple choices as shown in 17. Each trial began by displaying the question first. The users were asked to take their time to understand the question properly so that they could spend the trial time only looking for the answer. Once reading the question is over, the users pressed the start button, which started the timer triggering the display of the data and visualization, as well as the four possible choices. The timer stopped when user selected an answer. The user then had to press the submit button to go to the next question.

6.1.4 *Participants*

Twelve volunteers (4 females) participated in the study with ages ranging from 21 to 47 years of age.

6.2 RESULTS

I used an Analysis of Variance (ANOVA) test at the significance level of $\alpha = 0.05$ using the Bonferroni adjustment to carry out all the statistical analysis. Before this process, I performed a linear regression on the data to remove the outliers with time as the response variable for different predictor variables in each analysis.

6.2.1 *Task Time*

Before the experiment, the participants were asked to do the trials as quickly as possible without sacrificing accuracy. Therefore, I include all the trials in the time analysis interpreting the incorrect answers as characteristics of the visualizations.

The analysis yielded a main effect of visualization ($F_{2,22} = 8.9, p = 0.001$) with Geo-Temporal Tag Visualization (GTViz) having a better task completion time overall (Figure 18 and 19). We also see a main effect of task ($F_{8,88} = 26, p < 0.001$) and a significant interaction of Visualization \times Task ($F_{16,176} = 11, p < 0.001$).

The task completion time for each visualization technique is varied across the tasks as shown in Figure 10. There is no major difference

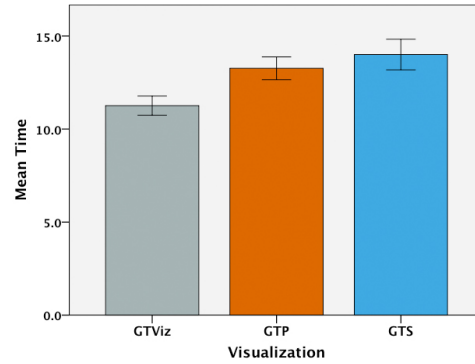


Figure 18: Mean task completion time for each visualization.

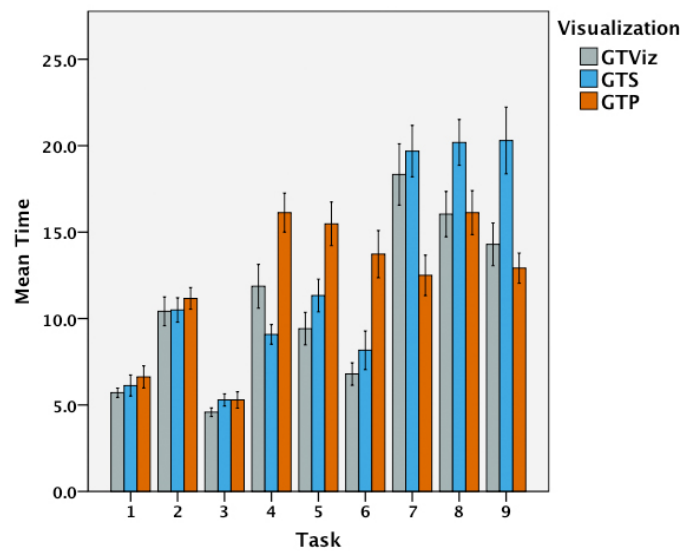


Figure 19: Mean task completion time for each task.

in task completion time across the first set of tasks (Task 1 to Task 3). GTViz and Geo-SparkClouds (GTS) performed faster than Geo-Temporal Pies (GTP) across Task 4 to Task 6 with GTS exhibiting better completion time for Task 4. GTP performs faster on Task 7 but this is almost matched by GTViz on Task 8 and Task 9.

To further investigate the performance of the visualizations across these tasks, I performed additional ANOVA tests within task category or complexity (Specific, Trends, Comparisons) and within task

type considering the unknown element of the tasks (topic, area or time period).

The results yield a significant interaction of Visualization \times Task Category ($F_{4,44} = 27.8, p < 0.001$) and a main effect of Task Category ($F_{2,540} = 156, p < 0.001$). Again, we can see no significant difference in the task completion time of the visualizations for the first category of tasks (Specific). For the second category (Trends), both GTViz and GTS perform faster than GTP. GTP and GTViz perform faster than GTS for the third category of tasks (Comparison). These results are shown in Figure 20.

We also find a main effect of Task Group ($F_{2,22} = 9.6, p = 0.001$) but no significant interaction of Visualization \times Task Group ($F_{4,44} = 2.8, p = 0.100$). These results are shown in Figure 21.

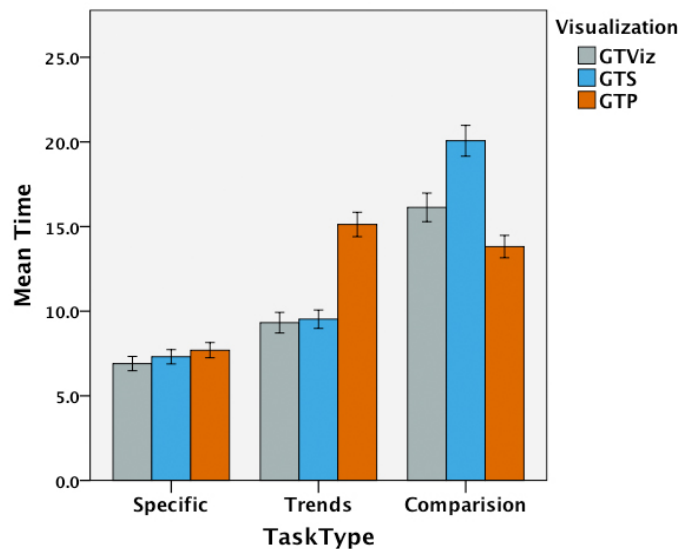


Figure 20: Mean task completion time for each task type.

With post-hoc comparisons, we find that GTViz generally performs better, followed by GTP and GTS respectively, which is reflected in

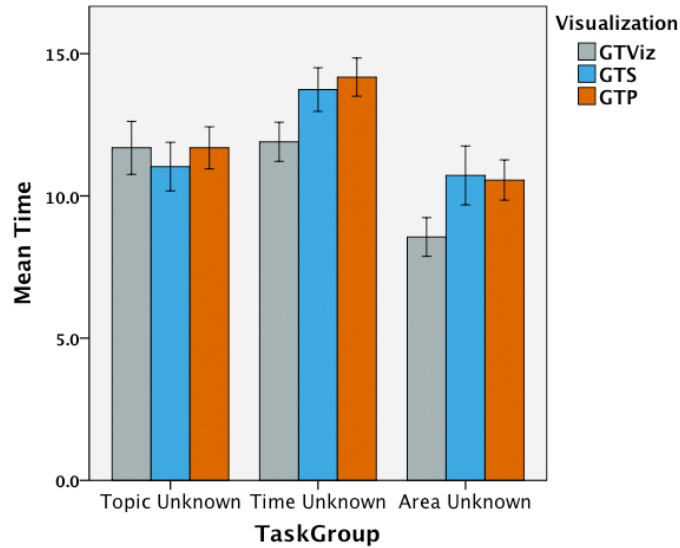


Figure 21: Mean task completion time for each group.

the mean overall task completion time for each visualization as shown in Figure 18.

6.2.2 Error

Since the participants were focused on the tasks, the overall error rate was very low at 0.9%. I conducted a non-parametric Friedman test to evaluate differences in mean error rate across the three visualizations. The test did not yield significance ($\chi^2(2) = 2.33, p = 0.311$). In fact, the last set of tasks (Task 7 to Task 8) has all the recorded errors with four of them on Task 7 and one each on Task 8 and Task 9.

6.2.3 *Subjective Preferences*

Participants were asked to provide subjective feedback on their preference of the visualization after completing each task. Responses were on a 5 Point Likert Scale, with 1=Not Preferred and 5=Most Preferred. Two final questions asked the user to provide a rank on the overall preference and the ease of use of each visualization (1=Very Difficult, and 5=Very Easy). I analyzed these ratings using Friedman tests and Wilcoxon tests for post-hoc analyses.

An overall test on the questions yields a significant difference ($\chi^2(2) = 15.4, p < 0.001$). Follow-up pair-wise Wilcoxon test reveal GTViz to be preferred over GTS and GTP. The same is reflected in the answer to the question about over-all preference, where we get a significance ($\chi^2(2) = 17.9, p < 0.001$). GTViz is preferred over GTS. There is no major preference of GTViz over GTP or GTS over GTP. I analyze each question separately as well. We don't see any significant preference of visualizations across Task 1 to Task 3. For Task 4, users showed a preference of GTS. For Task 5 and 6, users preferred GTViz over GTP and GTS. For Task 7 to Task 9, users preferred GTP and GTViz over GTS.

FUTURE WORK AND CONCLUSION

7.1 LIMITATIONS AND FUTURE WORK

The visualization tools that I presented address the issue of presenting the contents of geo-temporal documents on the map. It is a major step forward from the existing visualizations applicable to Crisis Maps. However, there are certain shortcomings and challenges that still need a solution. I mention some these in this section.

The list of colours distinguishable by people with or without colour blindness is very small. Though one can extend the visualizations to include as many words, it becomes very difficult to choose distinguishable colours for the words as the number of words increases.

Moreover, the current visualization shows only a particular number of line charts at a time. The chart becomes unreadable and cluttered as the number of activated words increases. One of the main goals of the visualization is to retain the geographic information, which we achieve to some degree. However, some areas on the map are still covered where the visualization components are placed. This limitation can be solved by making the component movable but still attached to the area of its origin by a thin plumb line. I intend to address these limitations in the future work.

Visualizing geo-spatial textual documents still presents several challenges. With geo-temporal tag cloud, we focus on documents or reports coming from crisis hit areas. With it, we provide an alternative way to visualize Crisis Maps. We are left with an interesting research challenge to see how this visualization technique can be extended to incorporate geo-temporal documents of a different nature. Moreover, in the case of crowd-sourced data, we rely on human expertise to tag the data accurately. However, looking at methods of automating this process presents an interesting research challenge.

As part of the future work, I will be improving and building upon the geo-temporal tag cloud and overcome the current limitations.

Data aggregation poses another research question in the context of visualizing the content of geo-temporal documents. Depending on the nature of the documents, choosing a different point and a different radius of aggregation could give a different view of the region. I also intend to explore the effects of different aggregation approaches and also look at providing the user tools to select areas of interest.

7.2 CONCLUSION

In my thesis, I presented novel techniques for the visualization of geo-temporal textual reports from a crisis hit area to support quick sense-making. The designs are inspired from the simplicity and usefulness of tag clouds, which I incorporate with temporal charts to allow analysis of the textual corpus over time. I also explained the sources of the datasets and the analysis process employed to convert

the data into a format suitable to be utilized by the visualization. I presented two implementations of the visualization using real-life datasets to demonstrate the capabilities of geo-temporal tag cloud and the benefits of going beyond dots and icons in crisis-maps. In the evaluation of the visualizations, I identified interesting phenomenon and trends in the data that one could note at a glance without having to go through all the reports. I also conducted a controlled user study to understand the performance and efficiency of the visualizations. The results showed each visualization having an advantage in at least one type of task with GTViz performing better overall.

BIBLIOGRAPHY

- [1] Ronja Addams-Moring, Markku Kekkonen, and Shushan Zhao. A simple taxonomy for mobile emergency announcement systems. In *Proceedings of the 2nd International ISCRAM Conference*, 2005. (Cited on pages 8 and 17.)
- [2] Hina Aman, Pourang Irani, and Fereshteh Amini. Revisiting crisis maps with geo-temporal tag visualization. In *Pacific Visualization Symposium (PacificVis) 2014*, Yokohama, Japan, 2014. IEEE. (Cited on page 29.)
- [3] Hina Aman, Pourang Irani, and Hai-Ning Liang. A review of common tasks supported by information communication technology for times of emergency. In *Proceedings of the 9th international conference on Information Systems for Crisis Response and Management (ISCRAM 2012)*, Vancouver, Canada, 2012. ISCRAM. (Cited on page 7.)
- [4] Llyr ap Cenydd, Rick Walker, Serban Pop, Helen Miles, Chris Hughes, William Teahan, and Jonathan C Roberts. epspread-storyboarding for visual analytics. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 311–312. IEEE, 2011. (Cited on pages vii, 27, and 28.)
- [5] Holly Ashley et al. *Change at hand: Web 2.0 for development*. Number 59. IIED, 2009. (Cited on page 8.)
- [6] John Barratt. Real-time local twitter trends. <http://trendsmap.com>, January 2013. (Cited on pages vii, 4, 23, and 25.)
- [7] M van Berlo, Richelle van Rijk, and Eric Buiël. A pc-based virtual environment for training team decision-making in high-risk situations. In *Proceedings of the 2nd International ISCRAM Conference April*, pages 195–201, 2005. (Cited on page 18.)
- [8] Mohamed Bishr and Lefteris Mantelas. A trust and reputation model for filtering and classifying knowledge about urban growth. *GeoJournal*, 72(3-4):229–237, 2008. (Cited on page 13.)
- [9] Harald Bosch, Dennis Thom, M Worner, Steffen Koch, Edwin Puttmann, Dominik Jackle, and Thomas Ertl. Scatterblogs: Geospatial document analysis. In *Visual Analytics Science and Technology (VAST), 2011 IEEE Conference on*, pages 309–310. IEEE, 2011. (Cited on pages vii, 27, and 28.)

- [10] Marcia Brooks. Challenges for warning populations with sensory disabilities. In *Proceedings of the 3rd international conference on Information Systems for Crisis Response and Management (IS-CRAM 2006)*, Newark, USA, 2006. ISCRAM. (Cited on page 17.)
- [11] Bill P Buckles, Laura Steinberg, Xiaohui Yuan, Xiaoping Liu, Liangmei Hu, and Yassine Belkhouche. Analysis, modeling, and rendering of urban flood events. In *Proceedings of the 2008 international conference on Digital government research*, pages 431–432. Digital Government Society of North America, 2008. (Cited on page 15.)
- [12] Monika Büscher, Preben Holst Mogensen, and Margit Kristensen. When and how (not) to trust it? *Crisis Response and Management and Emerging Information Systems: Critical Applications*, page 72, 2011. (Cited on page 11.)
- [13] Mifan Careem, David Bitner, and R Silva. Gis integration in the sahana disaster management system. In *Proceedings the International Conference on Information Systems for Crisis Response and Management, Delft, The Netherlands, 2007*. (Cited on page 14.)
- [14] Myriam Dunn Cavelty and Jennifer Giroux. Crisis mapping: A phenomenon and tool in emergencies. *CSS Analysis in Security Policy*, (103):1–4, 2011. (Cited on pages 13 and 29.)
- [15] Myriam Dunn Cavelty and Jennifer Giroux Morrow. Crisis mapping: A phenomenon and tool in emergencies. *CSS Analysis in Security Policy*, 103, 2011. (Cited on page 2.)
- [16] Christopher Collins, Fernanda B Viegas, and Martin Wattenberg. Parallel tag clouds to explore and analyze faceted text corpora. In *Visual Analytics Science and Technology, 2009. VAST 2009. IEEE Symposium on*, pages 91–98. IEEE, 2009. (Cited on page 26.)
- [17] Weiwei Cui, Shixia Liu, Li Tan, Conglei Shi, Yangqiu Song, Zekai Gao, Huamin Qu, and Xin Tong. Textflow: Towards better understanding of evolving topics in text. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2412–2421, 2011. (Cited on page 25.)
- [18] Sharon S Dawes, Anthony M Cresswell, and Bruce B Cahan. Learning from crisis lessons in human and information infrastructure from the world trade center response. *Social Science Computer Review*, 22(1):52–66, 2004. (Cited on page 10.)

- [19] T De Groeve and P Riva. Early flood detection and mapping for humanitarian response. In *Proceedings of the 6th international ISCRAM conference*, 2009. (Cited on pages 15 and 16.)
- [20] Sebastian Deneff, Leonardo Ramirez, Tobias Dyrks, and Gunnar Stevens. Handy navigation in ever-changing spaces: an ethnographic study of firefighting practices. In *Proceedings of the 7th ACM conference on Designing interactive systems*, pages 184–192. ACM, 2008. (Cited on page 18.)
- [21] Wenwen Dou, Xiaoyu Wang, Drew Skau, William Ribarsky, and Michelle X Zhou. Leadline: Interactive visual analysis of text data through event identification and exploration. In *Visual Analytics Science and Technology (VAST), 2012 IEEE Conference on*, pages 93–102. IEEE, 2012. (Cited on pages 4 and 24.)
- [22] Gabrielle Dunlevy. Flood rescue helicopter pilot relied on iphone for navigation. <http://www.couriermail.com.au/news/flood-rescue-helicopter-pilot-relied-on-iphone-for-navigation/story-e6freon6-1226045049383>, April 2011. (Cited on page 19.)
- [23] Ko Fujimura, Shigeru Fujimura, Tatsushi Matsubayashi, Takeshi Yamada, and Hidenori Okuda. Topigraphy: visualization for large-scale tag clouds. In *Proceedings of the 17th international conference on World Wide Web*, pages 1087–1088. ACM, 2008. (Cited on pages vii and 23.)
- [24] Georg Gartner, David A Bennett, and Takashi Morita. Towards ubiquitous cartography. *Cartography and Geographic Information Science*, 34(4):247–257, 2007. (Cited on page 13.)
- [25] Michael F Goodchild. Citizens as sensors: the world of volunteered geography. *GeoJournal*, 69(4):211–221, 2007. (Cited on page 13.)
- [26] Rebecca Goolsby. Social media as crisis platform: The future of community maps/crisis maps. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 1(1):7, 2010. (Cited on page 14.)
- [27] Martin J Halvey and Mark T Keane. An assessment of tag presentation techniques. In *Proceedings of the 16th international conference on World Wide Web*, pages 1313–1314. ACM, 2007. (Cited on page 29.)
- [28] Tony Hammond, Timo Hannay, Ben Lund, and Joanna Scott. Social bookmarking tools (i) a general review. *D-lib Magazine*, 2(4), 2005. (Cited on page 4.)

- [29] Susan Havre, Beth Hetzler, and Lucy Nowell. Themeriver: Visualizing theme changes over time. In *Information Visualization, 2000. InfoVis 2000. IEEE Symposium on*, pages 115–123. IEEE, 2000. (Cited on page 25.)
- [30] Thomas Heverin. Microblogging for distributed surveillance in response to violent crises: ethical considerations. In *Proceedings of the 2011 iConference*, pages 827–828. ACM, 2011. (Cited on page 12.)
- [31] Kevin Hoffman. In search of... the perfect tag cloud. *Whitepaper, August*, 2006. (Cited on page 37.)
- [32] Bernardo Huberman, Daniel Romero, and Fang Wu. Social networks that matter: Twitter under the microscope. *Available at SSRN 1313405*, 2008. (Cited on page 2.)
- [33] Amanda Lee Hughes and Leysia Palen. Twitter adoption and use in mass convergence and emergency events. *International Journal of Emergency Management*, 6(3):248–260, 2009. (Cited on page 12.)
- [34] Yasuyo G Ichihara, Masataka Okabe, Koichi Iga, Yosuke Tanaka, Kohei Musha, and Kei Ito. Color universal design: the selection of four easily distinguishable colors for all color vision types. In *Electronic Imaging 2008*, pages 68070O–68070O. International Society for Optics and Photonics, 2008. (Cited on page 37.)
- [35] CW Johnson. Applying the lessons of the attack on the world trade center, 11th september 2001, to the design and use of interactive evacuation simulations. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 651–660. ACM, 2005. (Cited on page 19.)
- [36] Gary Klein, Brian Moon, and Robert R Hoffman. Making sense of sensemaking 1: Alternative perspectives. *Intelligent Systems, IEEE*, 21(4):70–73, 2006. (Cited on page 30.)
- [37] Milos Krstajic, Enrico Bertini, and Daniel Keim. Cloudlines: compact display of event episodes in multiple time-series. *Visualization and Computer Graphics, IEEE Transactions on*, 17(12):2432–2439, 2011. (Cited on page 24.)
- [38] Jonas Landgren and Urban Nulden. A study of emergency response work: patterns of mobile phone interaction. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 1323–1332. ACM, 2007. (Cited on page 10.)

- [39] Bongshin Lee, Nathalie Henry Riche, Amy K Karlson, and Sheelagh Carpendale. Sparkclouds: Visualizing trends in tag clouds. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1182–1189, 2010. (Cited on pages vii, viii, 4, 24, 25, 40, and 43.)
- [40] Carson Kai-Sang Leung and Christopher L. Carmichael. Fpviz: a visualizer for frequent pattern mining. In *Proceedings of the ACM SIGKDD Workshop on Visual Analytics and Knowledge Discovery: Integrating Automated Analysis with Interactive Exploration VAKD*, 2009. (Cited on page 33.)
- [41] Carson Kai-Sang Leung, Fan Jiang, and Pourang P. Irani. Fpmapviz: A space-filling visualization for frequent patterns. In *IEEE 11th International Conference on Data Mining Workshops (ICDMW)*, pages 804–811. IEEE, 2011. (Cited on page 33.)
- [42] Sophia B Liu and Leysia Palen. The new cartographers: Crisis map mashups and the emergence of neogeographic practice. *Cartography and Geographic Information Science*, 37(1):69–90, 2010. (Cited on page 4.)
- [43] Sophia B Liu, Leysia Palen, Jeannette Sutton, Amanda L Hughes, and Sarah Vieweg. In search of the bigger picture: The emergent role of on-line photo sharing in times of disaster. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)*, 2008. (Cited on page 12.)
- [44] Ann Majchrzak and Philip HB More. Emergency! web 2.0 to the rescue! *Communications of the ACM*, 54(4):125–132, 2011. (Cited on page 2.)
- [45] Alessio Malizia, Pablo Acuña, Teresa Onorati, and Paloma Díaz. Cap-ones: an emergency notification system for all. *International Journal of Emergency Management*, 6(3):302–316, 2009. (Cited on page 17.)
- [46] Peter Meier. Haiti: Taking stock of how we are doing. <http://blog.ushahidi.com/2010/02/06/ushahidi-how-we-are-doing/>, March 2013. (Cited on page 2.)
- [47] Marcelo Mendoza, Barbara Poblete, and Carlos Castillo. Twitter under crisis: Can we trust what we rt? In *Proceedings of the first workshop on social media analytics*, pages 71–79. ACM, 2010. (Cited on page 12.)

- [48] Sharon Meraz. Citizen journalism, citizen activism, and technology: Positioning technology as a 'second superpower' in times of disasters and terrorism. In *7th International Symposium on Online Journalism*, pages 7–8, 2006. (Cited on page 12.)
- [49] Nathan Morrow, Nancy Mock, Adam Papendieck, and Nicholas Kocmich. Independent evaluation of the ushahidi haiti project. *Development Information Systems International*, 8, 2011. (Cited on page 2.)
- [50] Nathan Morrow, Nancy Mock, Adam Papendieck, and Nicholas Kocmich. Independent evaluation of the ushahidi haiti project. *Development Information Systems International*, 8, 2011. (Cited on page 29.)
- [51] Nancy Obermeyer. Thoughts on volunteered (geo) slavery. In *NCGIA and Vespucci Workshop on Volunteered Geographic Information, December*, pages 13–14, 2007. (Cited on page 13.)
- [52] Jens Ortmann, Minu Limbu, Dong Wang, and Tomi Kauppinen. Crowdsourcing linked open data for disaster management. In *Proceedings of the Terra Cognita Workshop on Foundations, Technologies and Applications of the Geospatial Web in conjunction with the ISWC*, pages 11–22, 2011. (Cited on page 2.)
- [53] Jens Ortmann, Minu Limbu, Dong Wang, and Tomi Kauppinen. Crowdsourcing linked open data for disaster management. In *Proceedings of the Terra Cognita Workshop on Foundations, Technologies and Applications of the Geospatial Web in conjunction with the ISWC*, pages 11–22, 2011. (Cited on page 13.)
- [54] Leysia Palen, Kenneth M Anderson, Gloria Mark, James Martin, Douglas Sicker, Martha Palmer, and Dirk Grunwald. A vision for technology-mediated support for public participation & assistance in mass emergencies & disasters. In *Proceedings of the 2010 ACM-BCS Visions of Computer Science Conference*, page 8. British Computer Society, 2010. (Cited on page 13.)
- [55] Leysia Palen and Sophia B Liu. Citizen communications in crisis: anticipating a future of ict-supported public participation. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 727–736. ACM, 2007. (Cited on pages 8 and 12.)
- [56] Yan Qu, Chen Huang, Pengyi Zhang, and Jun Zhang. Microblogging after a major disaster in china: a case study of the 2010 yushu earthquake. In *Proceedings of the ACM 2011 conference on Computer supported cooperative work*, pages 25–34. ACM, 2011. (Cited on page 11.)

- [57] Leonardo Ramirez and Tobias Dyrks. Designing for high expectations: balancing ambiguity and thorough specification in the design of a wayfinding tool for firefighters. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, pages 390–399. ACM, 2010. (Cited on page 19.)
- [58] Paul Rayson and Roger Garside. Comparing corpora using frequency profiling. In *Proceedings of the workshop on Comparing Corpora*, pages 1–6. Association for Computational Linguistics, 2000. (Cited on page 33.)
- [59] Magnus Rembold and Jurgen Spath. Graphical visualization of text similarities. http://www.munterbund.de/visualisierung_textaehnlichkeiten/essay.php, 2006. (Cited on page 26.)
- [60] AW Rivadeneira, Daniel M Gruen, Michael J Muller, and David R Millen. Getting our head in the clouds: toward evaluation studies of tagclouds. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 995–998. ACM, 2007. (Cited on page 22.)
- [61] Anthony C Robinson, Robert E Roth, and Alan M MacEachren. Challenges for map symbol standardization in crisis management. In *Proceedings of the 7th International ISCRAM Conference—Seattle*, volume 1, 2010. (Cited on page 14.)
- [62] Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, pages 851–860. ACM, 2010. (Cited on page 16.)
- [63] Christoph Schlieder and Olga Yanenko. Spatio-temporal proximity and social distance: a confirmation framework for social reporting. In *Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Location Based Social Networks*, pages 60–67. ACM, 2010. (Cited on page 12.)
- [64] Timothy Schoenharl, Greg Madey, Gábor Szabó, and Albert-László Barabási. Wiper: A multi-agent system for emergency response. In *Proceedings of the 3rd International ISCRAM Conference*, pages 1–7, 2006. (Cited on page 17.)
- [65] Aidan Slingsby, Jason Dykes, Jo Wood, and Keith Clarke. Interactive tag maps and tag clouds for the multiscale exploration of large spatio-temporal datasets. In *Information Visualization, 2007. IV'07. 11th International Conference*, pages 497–504. IEEE, 2007. (Cited on page 22.)

- [66] MWBE Smeets and Simone Sillem. Intelligent sms as an effective public warning system: the inspiring results of a dutch pilot project. In *Proceedings of the 2nd International ISCRAM Conference*, 2005. (Cited on page 17.)
- [67] Brian M Tomaszewski, Anthony C Robinson, Chris Weaver, Michael Stryker, and Alan M MacEachren. Geovisual analytics and crisis management. In *In Proc. 4th International Information Systems for Crisis Response and Management (ISCRAM) Conference, Delft, the Netherlands, 2007*. (Cited on page 29.)
- [68] Zachary O Toups, Andruid Kerne, William A Hamilton, and Nabeel Shahzad. Zero-fidelity simulation of fire emergency response: improving team coordination learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1959–1968. ACM, 2011. (Cited on page 18.)
- [69] David Troy. A real-time geographic visualization of posts to twitter. <http://twittervision.com>, March 2013. (Cited on page 4.)
- [70] E Tufte. Sparklines: Intense, simple, word-sized graphics. *Beautiful Evidence*, 1:46–63, 2004. (Cited on page 24.)
- [71] IEEE VAST. Ieee vast challenge 2011. <http://hcil.cs.umd.edu/localphp/hcil/vast11/>, March 2011. (Cited on page 26.)
- [72] Fernanda B Viégas, Scott Golder, and Judith Donath. Visualizing email content: portraying relationships from conversational histories. In *Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 979–988. ACM, 2006. (Cited on page 24.)
- [73] Fernanda B Viégas and Martin Wattenberg. Timelines tag clouds and the case for vernacular visualization, interactions, v. 15 n. 4. *July+ August*, 2008. (Cited on page 4.)
- [74] Fernanda B Viegas, Martin Wattenberg, Frank Van Ham, Jesse Kriss, and Matt McKeon. Manyeyes: a site for visualization at internet scale. *Visualization and Computer Graphics, IEEE Transactions on*, 13(6):1121–1128, 2007. (Cited on page 26.)
- [75] Sarah Vieweg, Amanda L Hughes, Kate Starbird, and Leysia Palen. Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1079–1088. ACM, 2010. (Cited on page 11.)

- [76] Sarah Vieweg, Leysia Palen, Sophia B Liu, Amanda L Hughes, and Jeannette Sutton. Collective intelligence in disaster: An examination of the phenomenon in the aftermath of the 2007 virginia tech shootings. In *Proceedings of the Information Systems for Crisis Response and Management Conference (ISCRAM)*, 2008. (Cited on page 12.)
- [77] William L Waugh Jr. Terrorism as disaster. In *Handbook of Disaster Research*, pages 388–404. Springer, 2007. (Cited on page 9.)
- [78] Furu Wei, Shixia Liu, Yangqiu Song, Shimei Pan, Michelle X Zhou, Weihong Qian, Lei Shi, Li Tan, and Qiang Zhang. Tiara: a visual exploratory text analytic system. In *Proceedings of the 16th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 153–162. ACM, 2010. (Cited on page 24.)
- [79] M Wu, Annie Liu, and K Mani Chandy. Virtual environments for developing strategies for interdicting terrorists carrying dirty bombs. In *International Society of Crisis Response and Management Conference*, pages 1–5, 2008. (Cited on page 16.)
- [80] Max Wyss. Earthquake loss estimates applied in real time and to megacity risk assessment. In *Proceedings of the Second International ISCRAM Conference, Brussels, Belgium*, pages 297–299, 2005. (Cited on pages 15 and 16.)
- [81] Max Wyss. The kashmir m7. 6 shock of 8 october 2005 calibrates estimates of losses in future himalayan earthquakes. In *Proceedings of the Third International ISCRAM Conference*, page 5, 2006. (Cited on pages 15 and 16.)
- [82] Dave Yates and Scott Paquette. Emergency knowledge management and social media technologies: A case study of the 2010 haitian earthquake. *International Journal of Information Management*, 31(1):6–13, 2011. (Cited on pages 10 and 11.)

COLOPHON

This thesis was typeset with the pdf_latex $\LaTeX 2_{\epsilon}$ interpreter using Hermann Zapf's *Palatino* type face for text and math and *Euler* for chapter numbers. The listings were set in *Bera Mono*.

The typographic style of the thesis was based on André Miede's wonderful classicthesis \LaTeX style available from CTAN. My modifications were limited to those required to satisfy the constraints imposed by my university, mainly 12pt font on letter-size paper with extra leading. Miede's original style was inspired by Robert Bringhurst's classic *The Elements of Typographic Style* [?]. I hope my naïve, yet carefully considered changes are consistent with Miede's original intentions.

Final Version as of April 28, 2014 at 2:48.