Exploring the Geographic Uncertainty Associated with Crowdsourced Crisis Information: a Geovisualisation Approach

Iain Dillingham

Thesis submission for admission to the degree of
Doctor of Philosophy in Geographic Information Science

City University London, Department of Computer Science

August 2013

CITY
UNIVERSITY OF LONDON
EST 1894

City, University of London
Northampton Square
London
EC1V 0HB
United Kingdom

T +44 (0)20 7040 5060

THE FOLLOWING PARTS OF THIS THESIS HAVE BEEN REDACTED FOR COPYRIGHT REASONS:

p. 142-162, Journal Articles

Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank Jason Dykes and Jo Wood for agreeing to supervise my studies. Over the last four years—from MSc to PhD—their enthusiasm, encouragement, and patient questioning have helped me immensely.

My colleagues from the giCentre and the Centre for Human–Computer Interaction Design at City University London also deserve special mention. Foremost among them, thank you to Aidan Slingsby, for his sound advice on handling the day-to-day issues encountered by a young researcher. As well as Aidan, it has been a pleasure to work with Donia Badawood, Roger Beecham, Graham Dove, Sarah Goodwin, Alexander Kachkaev, and Ali Ramathan: each has contributed to my thinking during our weekly group meetings, our collaborative projects, and while waiting for the kettle to boil.

Finally, thank you to John, Sue, Adam, Fiona, Aisla, and Allison for their encouragement and, most importantly, for reminding me of the things that matter most in life.

Iain Dillingham
London, August 2013

About the text

The text was prepared using the LaTeX markup language. Body text and headings are set in Charter BT. Monospaced text is set in Inconsolata. Figure text is set in Tuffy. British English spelling is used, although references adhere to the original source: for this reason, both *visualization* (American English) and *visualisation* (British English) can be found.

Finally, for issues of semantics, syntax, and style, there is no better reference than *BUGS in Writing: A Guide to Debugging Your Prose* (Dupré 1998). It was consulted regularly.

Declaration

I grant powers of discretion to the University Librarian to allow this thesis to be copied in whole or in part without further reference to me. This permission covers only single copies made for study purposes, subject to the normal conditions of acknowledgement.

Abstract

New information and communications technologies, such as mobile phones and social media, have presented the humanitarian community with a dilemma: how should humanitarian organisations integrate information from crisis-affected communities into their decision-making processes whilst guarding against inaccurate information from untrustworthy sources? Advocates of crisis mapping claim that, under certain circumstances, crowdsourcing can increase the accuracy of crisis information. However, whilst previous research has studied the geography of crisis information, the motivations of people who create crisis map mashups, and the motivations of people who crowdsource crisis information, the geography of, and the uncertainty associated with, *crowdsourced* crisis information has been ignored. As such, the current research is motivated by the desire to explore the geographic uncertainty associated with, and to contribute a better understanding of, crowdsourced crisis information.

The current research contributes to the fields of GISc (Geographic Information Science) and crisis informatics; crisis mapping; and geovisualisation specifically and information visualisation more generally. These contributions can be summarised as an approach to, and an understanding of, the geographic uncertainty associated with crowdsourced crisis information; three geovisualisation software prototypes that can be used to identify meaningful patterns in crisis information; and the design, analysis, and evaluation model, which situates the activities associated with designing a software artefact—and using it to undertake analysis—within an evaluative framework. The approach to the geographic uncertainty associated with crowdsourced crisis information synthesised techniques from GISc, geovisualisation, and natural language processing. By following this approach, it was found that location descriptions from the Haiti crisis map did not 'fit' an existing conceptual model, and, consequently, that there is a need for new or enhanced georeferencing methods that attempt to estimate the uncertainty associated with free-text location descriptions from sources of crowdsourced crisis information.

Chapter 1

Introduction

New information and communications technologies, such as mobile phones and social media, have allowed members of crisis-affected communities to exchange information with each other, and with the outside world, more effectively than ever before (Coyle & Meier 2009). However, the shift from one-to-many to many-to-many forms of communication during a crisis, sometimes at a global scale, has presented the humanitarian community with a dilemma (Coyle & Meier 2009). On the one hand, humanitarian organisations need to do more to integrate information from crisis-affected communities into their decision-making processes (Heinzelman & Waters 2010). An important criticism of the response to the 2010 Haiti earthquake, for example, was that traditional information gathering mechanisms did not, and could not, integrate information from ordinary Haitians, which resulted in many people suffering considerable (and unnecessary) hardship (Heinzelman & Waters 2010). On the other, humanitarian organisations do not wish to base important—sometimes life-threatening—decisions on what may turn out to be inaccurate information from untrustworthy sources: the risks are simply too great (Coyle & Meier 2009, Tapia et al. 2011). Consequently, whilst some argue that the humanitarian community needs to change its approach to new information and communications technologies (Coyle & Meier 2009), the implications of this change are far from clear.

In addition to promoting information exchange, social media have also popularised new types of applications, such as *mashups*, which merge separate data sources into single user interfaces (Zang et al. 2008). In the case of a *crisis map mashup*, the single user interface is a web-based interactive map and the separate data sources are APIs (Application Programming Interfaces) to social media, the content of which is related to the crisis at hand. People create crisis map mashups for a variety of reasons: the Ushahidi crisis map, for example, was created by

Kenyan activists following the general election in December 2007 to persuade and mobilise others (Liu & Palen 2010, Meier 2012). This crisis map was especially significant because its success popularised *crisis mapping*, an interdisciplinary field of study that is emerging at the intersection of geography, computer science, and disaster and emergency management (Meier 2009, Ziemke 2012). Advocates of crisis mapping claim that, under certain circumstances, the crowdsourcing of crisis information can increase its accuracy (see, for example, Meier 2013). As such, crisis mapping can be seen as an attempt to bridge the gaps between humanitarian organisations and crisis-affected communities: by facilitating the crowdsourcing of crisis information, crisis mapping claims to offer a solution to the problem of guarding against inaccurate information from untrustworthy sources.

Previous research has studied the geography of crisis information (Vieweg et al. 2010, Gelernter & Mushegian 2011), the motivations of people who create crisis map mashups (Liu & Palen 2010), and the motivations of people who crowdsource crisis information (Starbird & Palen 2011). However, the geography of, and the uncertainty associated with, *crowdsourced* crisis information has been ignored. As such, the current research is motivated by the desire to explore the geographic uncertainty associated with, and to contribute a better understanding of, crowdsourced crisis information and, ultimately, to test the claim that crisis mapping offers a solution to the problem of guarding against inaccurate information from untrustworthy sources.

1.1 Research scope

The current research is a case study[1] of the Haiti crisis map (Ushahidi 2010). The Haiti crisis map was created following the 7.0 magnitude earthquake that struck the country on 12 January 2010 (Disasters Emergency Committee 2013) and contains 3,606 reports that span the period from 12 January 2010 to 1 February 2011. These reports were downloaded on 10 February 2011. An example report is shown in Table 1.1. The Haiti crisis map and the Ushahidi crisis map are based on the same software platform, which has since adopted the Ushahidi name. Since the Haiti crisis map, Ushahidi has become a *de facto* standard for crisis mapping specifically and for mapping crowdsourced information more generally. Indeed, there were 12,795 Ushahidi-based maps by July 2012 (Crowdglobe 2012), a number that can only have increased in the intervening period.

Despite the existence of a *de facto* standard for crisis mapping, exploring the

[1] A *case study* can be defined as "an intensive study of a single unit for the purpose of understanding a larger class of (similar) units" (Gerring 2004, p.342).

Attribute	Value
id	15
title	Karibe Hotel Collapsed
date	2010-01-13 10:57:00
location	Karibe Hotel, Juvenat 7 Petion-Ville, ...
description	The Karibe Hotel and ajoining apartments ...
category	5a. Structure effondres | Collapsed structure,...
latitude	18.51933
longitude	-72.301626
approved	YES
verified	NO

Table 1.1: Example report from the Haiti crisis map.

geographic uncertainty associated with crowdsourced crisis information is problematic because the data are relatively sparse: unlike other projects that crowdsource information, such as OpenStreetMap (OpenStreetMap 2013) and Wikipedia (Wikipedia 2013), the Haiti crisis map specifically and Ushahidi-based maps more generally do not record how reports have been transformed by the crowdsourcing process. In other words, reports do not indicate either the lineage or the credibility of the information (MacEachren et al. 2005). Nevertheless, each report is located in two ways: as a free-text location description and as a crowdsourced point location. This characteristic is why the current research focuses on the *geographic* uncertainty associated with crowdsourced crisis information: not only is crowdsourced crisis information inherently geographic, but because location is expressed as a place (a free-text location description) and in space (a crowdsourced point location) the geography provides a basis for comparison.

1.2 Research problems

The current research addresses three interrelated problems:

1. How should free-text location descriptions from the Haiti crisis map be characterised?

2. How should free-text location descriptions from the Haiti crisis map be compared to their respective crowdsourced point locations?

3. What can be inferred about the the geographic uncertainty associated with the Haiti crisis map from these characteristics and comparisons?

1.3 Research contributions

The current research contributes to the fields of GISc (Geographic Information Science) and crisis informatics; crisis mapping; and geovisualisation specifically and information visualisation more generally.

- *GISc and crisis informatics.* The current research contributes an approach to, and an understanding of, the geographic uncertainty associated with crowdsourced crisis information. From a crisis informatics perspective, Palen et al. (2010, p.6) call for research into the characteristics of crisis information. The authors focus on accuracy, which can be seen as a component of uncertainty (Fisher 1999, MacEachren et al. 2005). From a GISc perspective, Devillers et al. (2010, p.396) call for research into new sources of spatial data, such as volunteered geographic information (Goodchild 2007). Crowdsourced crisis information can be thought of as volunteered geographic information. Consequently, the current research contributes to the research agendas proposed by Palen et al. (2010) and Devillers et al. (2010).

- *Crisis mapping.* The current research contributes three geovisualisation software prototypes that can be used to identify meaningful patterns in crisis information. This contribution is discussed in Chapter 5, Sections 5.4.1, 5.4.2, and 5.4.3. Meier (2009) and Ziemke (2012) call for research into techniques that support the visual analysis of crisis information; techniques that help validate crisis information, uncover patterns, and test hypotheses. Consequently, the current research contributes to the crisis mapping research agenda proposed by Meier (2009) and Ziemke (2012).

- *Geovisualisation.* The current research contributes the design, analysis, and evaluation model, which situates the activities associated with designing a software artefact—and using it to undertake analysis—within an evaluative framework: it is an iterative approach to the design, implementation, and evaluation of a software artefact that attempts to bridge the gap between the context of use and the context of design. This contribution is discussed in Chapter 5, Sections 5.1.4 and 5.5.

Preliminary findings have been presented at several national and international workshops and conferences.

- The 2011 IEEE Conference on Visual Analytics Science and Technology, Providence, RI, USA, 23–28 October 2011 (Dillingham, Dykes & Wood 2011).

- The 2012 GIS Research UK Annual Conference, Lancaster University, Lancaster, UK, 11–13 April 2012 (Dillingham et al. 2012*a*).
- The 2012 EuroVis Workshop on Visual Analytics, Vienna, Austria, 4–5 June 2012 (Dillingham et al. 2012*c*).
- The 2012 IEEE Conference on Information Visualization, Seattle, WA, USA, 14–19 October 2012 (Dillingham et al. 2012*b*).
- GeoViz Hamburg: Interactive Maps that Help People Think, HafenCity University Hamburg, Hamburg, Germany, 6–8 March 2013 (Dillingham, Dykes & Wood 2013).
- The 2013 GIS Research UK Annual Conference, The University of Liverpool, Liverpool, UK, 3–5 April 2013 (Dillingham, Wood & Dykes 2013).

Three parallel projects have contributed to the current research. Again, these have been presented at several national and international workshops and conferences.

- 'Exploring Road Incident Data with Heat Maps'. The 2011 GIS Research UK Annual Conference, The University of Portsmouth, Portsmouth, UK, 27–29 April 2011 (Dillingham, Mills & Dykes 2011).
- 'Monitoring the health of computer networks with visualization: VAST 2012 Mini Challenge 1 award: "Efficient use of visualization"'. The 2012 IEEE Conference on Visual Analytics Science and Technology, Seattle, WA, USA, 14–19 October 2012 (Kachkaev et al. 2012).
- 'Creative User-Centered Visualization Design for Energy Analysts and Modellers'. The 2013 IEEE Conference on Information Visualization, Atlanta, GA, USA, 13–18 October 2013 (Goodwin et al. 2013).

1.4 Research overview

The research overview, which is shown in Figure 1.1, provides a graphical summary of the current research. It is also shown at the beginning of each chapter, with the chapter highlighted in grey. The research overview emphasises that there are two strands to the current research. The first strand, which relates to the nature of crowdsourced crisis information, is positioned to the left of the research overview and encompasses Chapters 2, 3, and 4. The second strand, which relates to the visualisation of crowdsourced crisis information, is positioned to the right of the research overview and encompasses

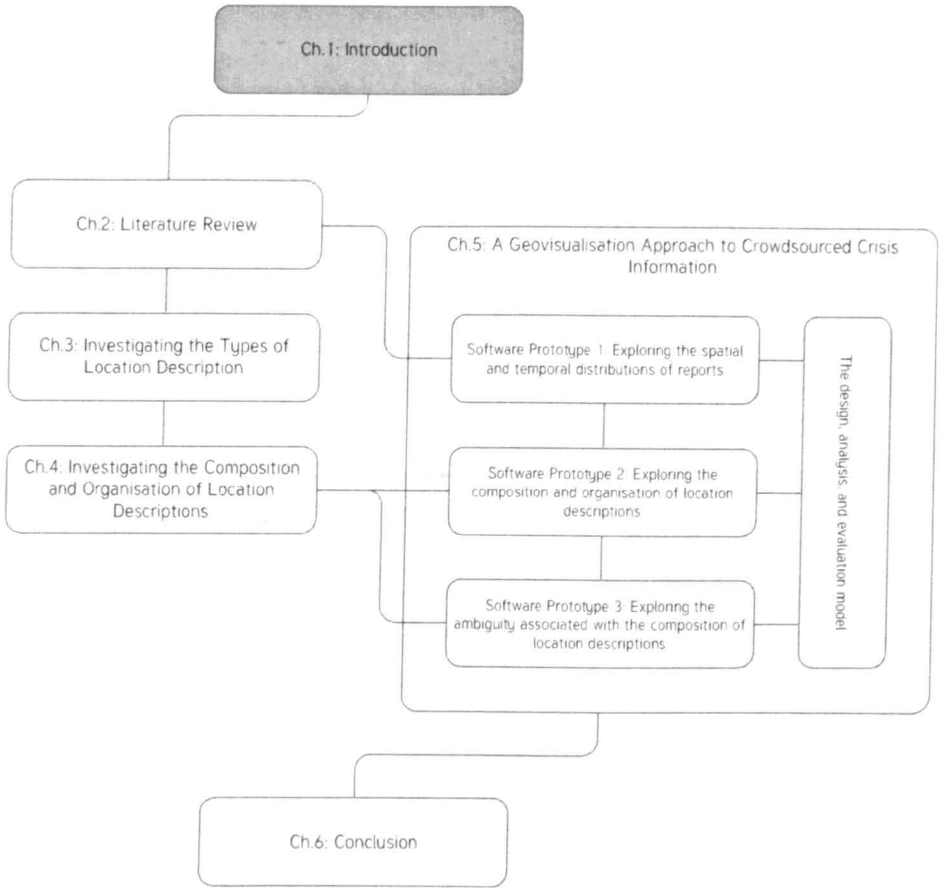

Figure 1.1: The research overview provides a graphical summary of the current research. It is displayed at the beginning of each chapter, with the current chapter highlighted in grey.

Chapter 5. The Introduction and Conclusion—Chapters 1 and 6—join these two strands. Consequently, these chapters are positioned at the centre of the research overview.

Chapter 2

Literature Review

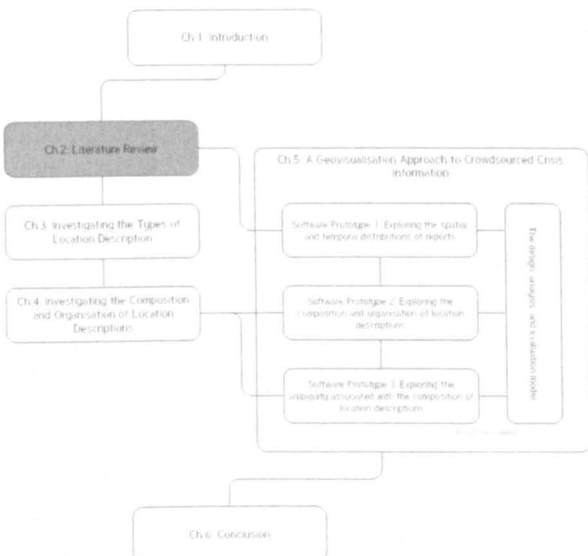

It is important to understand the nature of crisis information, as well as the practices that are associated with its creation, collection, dissemination, and use, from different perspectives. However, the literature is highly fragmented. On the one hand, crisis informatics has sought to address the criticism that information science has tended to ignore the role of the public in times of crisis (Palen & Liu 2007) by investigating the practices of individuals and groups who participate in crisis events, especially when their participation is mediated by new communications technologies (see, for example, Palen & Liu 2007, Liu & Palen 2010, Palen et al. 2010, Vieweg et al. 2010). On the other, there are divergent approaches to crisis information within geographic information science (GISc). These approaches can be summarised as those that emphasise volunteered geographic information

(Goodchild 2007, 2009) and the notion of space, and those that emphasise user-generated content and the notion of place. This chapter brings together both perspectives. First, crisis information is considered from the perspective of crisis informatics. The discussion covers the geography of communications within the crisis-affected community, the Ushahidi crisis map and crisis mapping, and the concerns formal response organisations have with using information sourced directly from the crisis-affected community. Second, crisis information is considered from the GISc perspective. The discussion starts with a consideration of crisis information as volunteered geographic information, and then broadens to encompass notions of space and place, and issues of geographic information uncertainty more generally.

2.1 Crisis information from the crisis informatics perspective

Palen & Liu (2007) identify three pathways along which information travels during a crisis: first, communications within the crisis-affected community; second, communications between the crisis-affected community and the wider world; and third, communications between the crisis-affected community and formal response organisations. These three pathways structure the following discussion of crisis information from the crisis informatics perspective. First, existing research into the geography of communications within the crisis-affected community is reviewed. Second, the Ushahidi crisis map—an example of a crisis map mashup—and the transition from the Ushahidi crisis map to crisis mapping is situated within existing research into communications between the crisis-affected community and the wider world. This section also briefly considers common motivations for creating crisis map mashups. Third, existing research into communications between the crisis-affected community and formal response organisations is considered. It is shown that, whilst formal response organisations are willing to use social media under certain circumstances, they see information sourced directly from the crisis-affected community as not trustworthy and not accurate. The degree to which information is trustworthy and accurate is returned to in Section 2.2.2, where issues of geographic information uncertainty are discussed.

2.1.1 Communications within the crisis-affected community

Several authors have highlighted the potential benefits of using communications within the crisis-affected community, such as *tweets* (messages shared on the Twitter social media website) and images, as sources of information. Vieweg et al.

(2010), for example, argue these communications can enhance *situational awareness*, which MacEachren et al. (2011) define as both the process and the result of determining the state of the world relevant to the goals of crisis management, and the use of this result to make decisions or predictions. Nevertheless, the nature of communications within the crisis-affected community has changed rapidly (and will continue to change rapidly) because of the influence of new technologies (Palen & Liu 2007, Palen et al. 2010).

Two studies from the crisis informatics perspective have investigated the geography of communications within the crisis-affected community. In the first, Vieweg et al. (2010) investigated tweets that related to the Red River floods in the northern United States and Canada, and to the wildfires in Oklahoma, that took place in March and April 2009. In the second, Gelernter & Mushegian (2011) investigated the types of location information that were contained in tweets that related to the earthquake that struck Christchurch, New Zealand, in February 2011.

In the first study, Vieweg et al. (2010) collected 19,162 flood-related tweets from 49 Twitter users and 2,779 wildfire-related tweets from 46 Twitter users. In both cases, the Twitter users were located in the immediate vicinity of either the floods or the wildfires. The authors found that whilst only 18% of flood-related tweets and 40% of wildfire-related tweets contained location information, 86% of flood-affected users and 78% of wildfire-affected users sent at least one tweet that contained at least one location. The authors also found that 6% of flood-related tweets and 8% of wildfire-related tweets contained relative references (e.g. "*x* miles from *y*"). Vieweg et al. (2010) argue that although location information was important to both groups of Twitter users (as evidenced by the high proportion who sent at least one tweet that contained at least one location), the unpredictable nature of wildfires meant that locations were communicated more frequently. Consequently, Vieweg et al. (2010) conclude that the predictability of the crisis (i.e. the degree to which the location of the crisis was predictable) was related to the amount of location information that was communicated by members of the crisis-affected community.

In the second study, Gelernter & Mushegian (2011) collected 1,407 earthquake-related tweets. Unlike Vieweg et al. (2010), the authors do not report how many users sent these tweets, or whether or not these users were located in the immediate vicinity of the earthquake. However, based on a sample of 300 (21%) of these tweets, the authors developed a typology of locations. The classes in this typology described locations at a wide range of geographic scales, from country-level to address-level (Table 2.1). The authors then applied this typology to all 1,407 earthquake-related tweets. Whilst they do not report how frequently the classes were identified, they do report that 253 locations were identified 1,207 times.

Without knowing how frequently the classes were identified, it is hard to know what proportion of earthquake-related tweets contained location information. Consequently, it is hard to compare flood-related, wildfire-related, and earthquake-related tweets. Nevertheless, it seems that it was relatively common for earthquake-related tweets to contain location information. Furthermore, it is clear that the types of location identified by Gelernter & Mushegian (2011) in earthquake-related tweets are similar to the types of location identified by Vieweg et al. (2010) in both flood-related and wildfire-related tweets. In each case, the types of location seemed to refer to well-defined features, such as administrative divisions, addresses, or readily identifiable buildings (Table 2.1). That is to say, the types of location seemed not to refer to *vague places*, or areas that have vernacular names and vague spatial extents (Jones et al. 2008).

In summary, existing research into the geography of communications within the crisis-affected community suggests that there is an inverse relationship between the predictability of the crisis and the amount of location information that is communicated by members of the crisis-affected community. Furthermore, locations tend to refer to well-defined features and tend not to be referred to in terms of relative references. The discussion now turns to the second pathway along which—according to Palen & Liu (2007)—information travels during a crisis: communications between the crisis-affected community and the wider world.

2.1.2 Communications between the crisis-affected community and the wider world: from the Ushahidi crisis map to crisis mapping

Zang et al. (2008) define a *mashup* as a single user interface to multiple data sources. In the case of a *crisis map mashup*, the single user interface is a web-based interactive map and the multiple data sources are APIs (application programming interfaces) to social media, the content of which is related to the crisis at hand. Consequently, a crisis map mashup could be said to exemplify communications between the crisis-affected community and the wider world. Liu & Palen (2010) highlight five common motivations for creating crisis map mashups:

1. *To expedite communication.* Creators were motivated by the desire to communicate with a large number of people.

2. *To persuade and mobilise others.* Creators were motivated by the desire to make crisis information more compelling.

Floods and wildfires (Vieweg et al. 2010)
County name
Place name
City name
Address
Highway

Earthquake (Gelernter & Mushegian 2011)
Country
State, region, city
Abbreviated place
Neighbourhood or district
Topographical or infrastructure feature
Cluster of buildings
Geolocatable building, area, or organisation
What and where
Street address
Multiple places
Generic place
Place with hashtag

Table 2.1: Types of location in crisis-related tweets.

3. *To make information more accessible and easy to use.* Creators were motivated by their dissatisfaction with geographic information systems.

4. *To explore the potential benefits of visualisation.* Creators were motivated by the desire to explore the potential benefits of visualisation.

5. *For personal interest and personal gain.* Creators were motivated by the desire to demonstrate their technical prowess to others.

Liu & Palen (2010) discuss the Ushahidi crisis map, which Kenyan activists created following the general election in December 2007 (Meier 2012). Motivated by the desire to persuade and mobilise others (Liu & Palen 2010), these activists encouraged ordinary Kenyans to report their experiences of violence and intimidation; these reports were then represented on a web-based interactive map (Meier 2012). The Ushahidi crisis map led to a software platform and a non-profit company[1] (Ushahidi 2013*b*). The success of Ushahidi also popularised *crisis mapping*, an interdisciplinary field of study that is emerging at the intersection of geography, computer science, and disaster and emergency management (Meier 2009, Ziemke 2012).

[1]For convenience, *Ushahidi* will refer to the software platform, rather than the non-profit company. However, the two are closely related. *The Ushahidi crisis map* will refer to the crisis map that was created by Kenyan activists following the general election in December 2007.

The Ushahidi crisis map and crisis mapping are closely associated with the notion of *crowdsourcing*, which can be defined as the process whereby a group of individuals (the crowd) completes tasks in response to an open call (Howe 2006, 2009). However, the transition from the Ushahidi crisis map to crisis mapping saw important changes to the crowd. Indeed, whereas the Ushahidi crisis map was composed of one crowd—ordinary Kenyans, or the crowd of contributors—crisis mapping is composed of two crowds: first, the crowd of contributors; and second, the crisis mappers. The second crowd, the crisis mappers, are often "remote operators", far from the crisis-affected area (Starbird & Palen 2011, p.1076), and gather, translate, georeference, and verify information from the crowd of contributors (Meier & Munro 2010). On the one hand, advocates of crisis mapping claim that the work of crisis mappers increases the accuracy of crisis information (see, for example, Meier 2013). This claim is supported by praise for the Haiti crisis map from FEMA (the Federal Emergency Management Agency) and by anecdotal evidence that it was used by search and rescue teams, the US Marines, and the US Coastguard (Heinzelman & Waters 2010, McClendon & Robinson 2012). However, this claim is often reported as fact: see, for example, Heinzelman & Waters (2010) for non-academic, and Starbird (2012) for academic, examples. On the other, Sutherlin (2013) argues—from her experience of the crises in Haiti in 2010, Libya and Egypt in 2011, and Somalia in 2011–2012—that separating the crowd of contributors from their information makes it almost impossible to verify how well the the crowd of crisis mappers translate this information. Consequently, it is hard to test the claim that the work of crisis mappers increases the accuracy of crisis information.

In summary, the Ushahidi crisis map and crisis mapping are closely associated with the notion of crowdsourcing. However, the transition from the Ushahidi crisis map to crisis mapping saw important changes to the crowd. Existing research into communications between the crisis-affected community and the wider world highlights the limitations of the claim that the work of crisis mappers increases the accuracy of crisis information. The discussion now turns to the third pathway along which—according to Palen & Liu (2007)—information travels during a crisis: communications between the crisis-affected community and formal response organisations.

2.1.3 Communications between the crisis-affected community and formal response organisations

The collection, dissemination, and use of information sourced *indirectly* from the crisis-affected community by formal response organisations is not new: MapAction, for example, have acted as an interface between the crisis-affected community and formal response organisations since 2003 (MapAction 2011). However, there have been increasing calls from NGOs (Non-Governmental Organisations) for formal response organisations to use information sourced directly from the crisis-affected community in their decision-making processes, by, for example, gathering information from social media (see, for example, Coyle & Meier 2009, Heinzelman & Waters 2010). Nevertheless, despite their willingness to use social media to support their public relations activities, formal response organisations are reluctant to use information sourced directly from the crisis-affected community in their decision-making processes because they see this information as not trustworthy and not accurate (Tapia et al. 2011). Indeed, when Tapia et al. (2011) interviewed representatives of 13 humanitarian organisations, they found that these representatives were unwilling to 'trade' more timely information for less accurate information. This finding appears to challenge the the GISc perspective on crisis information, which is discussed in the next section.

In summary, despite increasing calls from NGOs and their willingness to use social media to support their public relations activities, formal response organisations are reluctant to use information sourced directly from the crisis-affected community in their decision-making processes because they see this information as not trustworthy and not accurate.

In this section, crisis information was considered from the perspective of crisis informatics. The discussion covered the geography of communications within the crisis-affected community, the Ushahidi crisis map and crisis mapping, and the concerns formal response organisations have with using information sourced directly from the crisis-affected community. In the next section, crisis information is considered from the GISc perspective. The discussion starts with a consideration of crisis information as volunteered geographic information, and then broadens to encompass notions of space and place, and issues of geographic information uncertainty more generally.

2.2 Crisis information from the GISc perspective

GISc has long been interested in crisis information. Indeed, Ushahidi can be seen as a *public participation GIS* as it was designed to promote the interests of grass roots, non-governmental organisations (Sieber 2006). Nevertheless, there are divergent approaches to crisis information within GISc. These approaches can be summarised as those that emphasise volunteered geographic information (Goodchild 2007,

2009) and the notion of space, and those that emphasise user-generated content and the notion of place. The implications of expressing location in terms of place, rather than in terms of space, are returned to in Section 2.2.1. In this section, research that considers crisis information as volunteered geographic information is reviewed.

Goodchild & Glennon (2010) offer a research agenda for investigating the role of volunteered geographic information in times of crisis. In common with previous assertions (Goodchild 2007, 2009), the authors argue that volunteered geographic information is often more timely, but less accurate than asserted geographic information. However, according to Goodchild & Glennon (2010), false negatives (not reacting when something is true), which are often a result of waiting more time for more accurate information, are more costly than false positives (reacting when something is not true), which are often a result of waiting less time for less accurate information. Consequently, Goodchild & Glennon (2010) claim that in times of crisis, the advantages of volunteered geographic information outweigh the limitations. Nevertheless, the authors call for research into the two aspects of volunteered geographic information—the geographic information aspect and the volunteered aspect—that is directed towards better understanding the accuracy of volunteered geographic information in times of crisis.

A small number of studies support the claim that in times of crisis, the advantages of volunteered geographic information outweigh the limitations (Goodchild & Glennon 2010). Three examples are discussed in this section: Corbane et al. (2012), Crooks et al. (2013), and Crooks & Wise (2013).

Corbane et al. (2012) used SMS messages that were sent to Mission 4636—an initiative that was closely associated with the Haiti crisis map—to investigate the relationship between the spatial distribution of these messages and building damage. The authors found a strong association between the spatial distribution of SMS messages and building damage and argue that, more generally, the former can be used as an early indicator of the latter. Nevertheless, the study area was a 9km by 9km area centred on Port-au-Prince, the Haitian capital. Not only were there more buildings and more people in this area than in other parts of the country, but this area was very close to the epicentre of the earthquake. Consequently, the potential for damage to be caused and reported was considerable: it is unclear whether the association between the spatial distribution of SMS messages and building damage would be as strong in areas of fewer buildings and fewer people.

Crooks et al. (2013) compared two forms of crisis information that related to a 5.8 magnitude earthquake that struck the eastern US on 23 August 2011: the first, 21,362 tweets that contained the #earthquake or #quake hashtags and were located by the devices that sent the tweets, rather than by applying natural-language processing techniques to the tweets themselves; the second, approximately 125,000 reports submitted to the USGS' 'Did you feel it?' website. Whilst the comparison was limited to the spatial and temporal distributions of the two forms of crisis information, the authors reiterated the claim that volunteered geographic information from Twitter was faster than volunteered geographic information from the USGS' 'Did you feel it?' website.

Crooks & Wise (2013) argue that agent-based models, which can be used by members of the crisis management community to explore how a population might react to a crisis, can be initialised with crisis information. However, although the authors suggest that reports from the Haiti crisis map could be used to initialise an agent-based model, they conduct their investigation using non-crisis information. Consequently, the strengths and limitations of using crisis information to initialise agent-based models are unclear.

In summary, a small number of studies support the claim that in times of crisis, the advantages of volunteered geographic information outweigh the limitations (Goodchild & Glennon 2010). Nevertheless, these studies rely on volunteered geographic information being located in terms of space, rather then in terms of place. The familiar space–time–attribute model of geographic information (see, for example, Andrienko & Andrienko 2006, Longley et al. 2011) has recently been challenged and a broader conception of geographic information—"information that links names and descriptive information to particular places, features, or locations on the Earth's surface" (Elwood et al. 2012, p.572)—has emerged. Consequently, the discussion now turns to notions of space and place.

2.2.1 Notions of space and place

GISc does not capture the notion of socially-produced, populated place as well as it captures the notion of Euclidean space (Fisher & Unwin 2005). However, the notion of place can be captured in a free-text location description, such as a tweet, and then this free-text location description can be georeferenced to give a point location (a coordinate) and hence the notion of space. As such, a free-text location description, such as "Outside the National Gallery, on Trafalgar Square" can be transformed into a point location, such as "51.508046, -0.128908".

To understand what is lost when a free-text location description is transformed into a point location, it is necessary to estimate the uncertainty associated with the free-text location description. At least two georeferencing methods have been proposed that attempt to estimate this uncertainty: the point-radius and the probability distribution methods (Wieczorek et al. 2004, Guo et al. 2008). In the former case, the location is represented as a point and the associated uncertainty as a radius; in the latter, the associated uncertainty is represented as a field. These methods have been applied to free-text location descriptions stored in the Mammal Networked Information System (Wieczorek et al. 2004, Guo et al. 2008). In addition, the point-radius method has been applied to free-text location descriptions that relate to search and rescue incidents (Doherty et al. 2011), which have clear similarities with those from the Haiti crisis map.

Based on this research, Guo et al. (2008) propose a conceptual model of a free-text location description. In this model, the place described by the free-text location description is known as the *target object*. However, the actual terms used are known as the *reference objects*. The reference objects may be related to the target object by one or more spatial relationships. For example, the target object "Outside the National Gallery, on Trafalgar Square" is a combination of two reference objects—"the National Gallery" and "Trafalgar Square"—and two spatial relationships—"Outside" and "on". Consequently, georeferencing is the process of estimating the location of the target object based on the reference objects and the spatial relationships (Guo et al. 2008).

In summary, when georeferencing free-text location descriptions, such as those from the Haiti crisis map, it is necessary to estimate the uncertainty associated with the location to understand what is lost when a place-based representation is transformed into a space-based representation. At least two georeferencing methods have been proposed that attempt to estimate this uncertainty. These methods are based on a conceptual model of a free-text location description that will prove useful when considering free-text location descriptions from the Haiti crisis map. The discussion now turns to issues of geographic information uncertainty more generally.

2.2.2 Geographic information uncertainty

In Section 2.1.3, it was noted that formal response organisations are reluctant to use information sourced directly from the crisis-affected community in their decision-making processes because they see the information as not trustworthy and not accurate (Tapia et al. 2011). In the GISc literature, the degree to which information is trustworthy and accurate is often seen as a matter of the uncertainty associated with the information. Interest in geographic information uncertainty grew from efforts to establish spatial data quality standards in the 1980s and 1990s (Devillers et al. 2010) and in the intervening 30 years, many researchers have discussed different classes of geographic information uncertainty. Fisher (1999), for example, discusses ambiguity and vagueness, the distinction between which is perhaps best explained by Plewe (2002), who argues there is a difference between the uncertainty associated with the conceptualisation of a phenomenon and the uncertainty associated with the measurement of a phenomenon. On the one hand, *ambiguity* can be defined as the uncertainty associated with the conceptualisation of a phenomenon and, as such, is a matter of subjective judgement (Fisher 1999). On the other, *vagueness* can be defined as the uncertainty associated with the measurement of a phenomenon and, as such, is the degree to which an object is a member of a class of objects (Fisher 1999).

MacEachren et al. (2005) provide a comprehensive typology of geographic information uncertainty, which has been validated in the domain of intelligence analysis (MacEachren et al. 2005, Thomson et al. 2005) and the domain of floodplain mapping (Roth 2009). The authors identify nine classes of geographic information uncertainty, which are as follows:

1. *Accuracy (error)*. The difference between a measurement and the reality. This value is usually estimated based on knowledge of the underlying phenomena.

2. *Precision*. The exactness of a measurement. This value is usually derived from the measurement device.

3. *Completeness*. The degree to which the information includes everything that is of interest.

4. *Consistency*. The degree to which the components of the information are in mutual agreement.

5. *Lineage*. The conduit through which the information has passed. This type has multiple sub-types (e.g. number and nature of processes).

6. *Currency*. The time between information collection, processing, and use. This type is often a function of context.

7. *Credibility*. The perceived reliability of the information source. This perception might be based on previous experience.

8. *Subjectivity*. The degree to which the information resulted from judgement rather than measurement.

9. *Interrelatedness*. The degree to which the information source is independent of other information sources.

Accuracy is often discussed in the GISc literature (Fisher 1999) and several methods have been developed to measure the error associated with point, line, and area features (see, for example, Goodchild & Hunter 1997). These methods are often based on the comparison of a tested source of unknown accuracy to a reference source of known accuracy (Goodchild & Hunter 1997). Consequently, whilst MacEachren et al. (2005) define accuracy as the difference between *a measurement* and *the reality*, it should be remembered that there is a difference between the uncertainty associated with the conceptualisation of a phenomenon and the uncertainty associated with the measurement of a phenomenon (Plewe 2002). In other words, *the reality* is actually a conceptual model (Veregin 1999).

Returning to the reluctance of formal response organisations to use information sourced directly from the crisis-affected community in their decision-making processes, which was discussed in Section 2.1.3, it is clear that the degree to which information is trustworthy incorporates elements of lineage and credibility from the MacEachren et al. (2005) typology, whilst the degree to which information is accurate incorporates elements of accuracy and precision from the MacEachren et al. (2005) typology. Unfortunately, reports from the Haiti crisis map do not indicate either the lineage or the credibility of the information, which makes an investigation into the trustworthiness of crowdsourced crisis information impossible. Estimating the accuracy of reports from the Haiti crisis map is is also problematic because there are few, if any, reference sources. Nevertheless, each report from the Haiti crisis map is located in two ways: as a free-text location description and as a crowdsourced point location. Treating the crowdsourced point location as the reference source and the free-text location description as the tested source (or vice versa) is a promising means of understanding the uncertainty associated with this information. Indeed, this approach is essentially the same as investigating the consistency of crowdsourced crisis information, as *the reality* (a conceptual model) is simply an alternative component of the information.

The concept of precision is similar to the concept of *granularity*, or the degree to which space can be partitioned into a series of cells (Bittner & Smith 2001). It seems reasonable to assume that crowdsourced crisis information is intended for situations where high precision (fine granularity) is required: Heinzelman & Waters (2010) and McClendon & Robinson (2012), for example, state that the Haiti crisis map was used by search and rescue teams, the US Marines, and the US Coastguard, who would clearly require very precise information to complete their missions successfully. Whilst the precision of a point location may be spurious (an artefact of the computation, rather than of the measurement, process) the granularity of a free-text location description is likely to give a better indication of its precision. Consequently, investigating the precision of crowdsourced crisis information is one means of understanding the uncertainty associated with this information.

In summary, MacEachren et al. (2005) provide a comprehensive typology of geographic information uncertainty that will prove useful when considering free-text location descriptions from the Haiti crisis map. Indeed, it is argued that the reluctance of formal response organisations to use information sourced directly from the crisis-affected community in their decision-making processes (Tapia et al. 2011) can be seen as a matter of the lineage, credibility, accuracy, and precision of this information. Given the limitations of reports from the Haiti crisis map, however, it is clear that the current research should focus on the accuracy and precision of reports.

In this section, crisis information was considered from the GISc perspective. The discussion started with a consideration of crisis information as volunteered geographic information, and then broadened to encompass notions of space and place, and issues of geographic information uncertainty more generally.

2.3 Conclusion

This chapter sought to understand crisis information, as well as the practices that surround its creation, collection, dissemination, and use, from the perspectives of crisis informatics and GISc.

The discussion of crisis information from the crisis informatics perspective was structured according to three pathways along which—according to Palen & Liu (2007)—information travels during a crisis: first, communications within the crisis-affected community; second, communications between the crisis-affected community and the wider world; and third, communications between the crisis-affected community and formal response organisations. Treating each pathway in turn, existing research into the geography of communications within the crisis-affected community suggests that there is an inverse relationship between the predictability of the crisis and the amount of location information that is communicated by members of the crisis-affected community. Furthermore, locations tend to refer to well-defined features and tend not to be referred to in terms of relative references. Existing research into communications between the crisis-affected community and the wider world highlights the limitations of the claim that the work of crisis mappers increases the accuracy of crisis information. Finally, existing research into communications between the crisis-affected community and formal response organisations suggests that formal response organisations see information sourced directly from the crisis-affected community as not trustworthy and not accurate. Consequently, formal response organisations are reluctant to use information sourced directly from the crisis-affected community in their decision-making processes.

The discussion of crisis information from the GISc perspective started with a consideration of crisis information as volunteered geographic information, and then broadened to encompass notions of space and place, and issues of geographic information uncertainty more generally; issues that intersect with the concerns formal response organisations have with using information sourced directly from the crisis-affected community. Existing research into crisis information as volunteered geographic information supports the claim that in times of crisis, the advantages of volunteered geographic information outweigh the limitations. However, these studies tend not to consider the uncertainty associated with the location component of crisis information. MacEachren et al. (2005) provide a comprehensive typology of geographic information uncertainty that will prove useful when considering free-text location descriptions from the Haiti crisis map. Indeed, it is argued that the reluctance of formal response organisations to use information sourced directly from the crisis-affected community in their decision-making processes (Tapia et al. 2011) can be seen as a matter of the lineage, credibility, accuracy, and precision of this information.

The current research now splits into two strands. The first strand relates to the nature of crowdsourced crisis information and is continued in Chapter 3, where an investigation into the types of location description from the Haiti crisis map is reported. The second strand relates to the visualisation of crowdsourced crisis information and is continued in Chapter 5, Section 5.4.1, where a geovisualisation software prototype for exploring the spatial and temporal distributions of reports is discussed.

Chapter 3

Investigating the Types of Location Descriptions

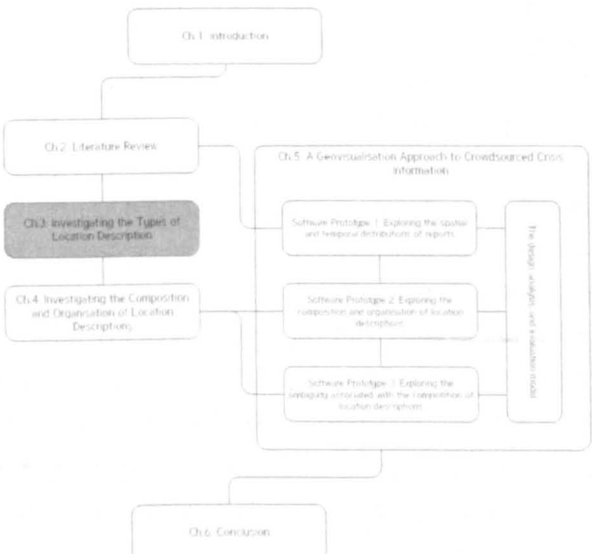

In Chapter 2, Section 2.2.2 it was noted that an investigation into the consistency of crowdsourced crisis information was a promising means of understanding the uncertainty associated with this information. As with accuracy, consistency involves a comparison: in this case, a comparison of each point location with its associated location description. In turn, this comparison requires either the point location to be transformed into a location description, or the location description to be transformed into a point location. Clearly, the latter is more straightforward than the former and in Chapter 2, Section 2.2.1 two georeferencing methods were discussed that, as well as transform the location description into a point location, also estimate the uncertainty associated with this point location. These were the point-radius and the probability distribution methods of Wieczorek et al. (2004) and Guo et al. (2008).

The point-radius and the probability distribution methods have their strengths and limitations, and these shaped the primary and secondary aims of the research that is reported in this chapter and elsewhere (Dillingham et al. 2012*a*). An important limitation of the point-radius and the probability distribution methods is that they only account for certain types of uncertainty, namely, those types associated with reference objects and spatial relationships. The former group contains spatial extent, an unknown datum, imprecise coordinate measurements, and map scale; the latter group contains imprecise distance and direction measurements (Guo et al. 2008). Consequently, the primary aim of the research reported in this chapter was to investigate whether the types of location descriptions that were identified by Wieczorek et al. (2004) and Guo et al. (2008) could be found in location descriptions from the Haiti crisis map. If they could, then either of these methods could be applied to these location descriptions. Each report would then be associated with either a circle or a field of uncertainty (depending on the chosen method) as well as the crowdsourced point location. Using the circle of uncertainty as an example, the crowdsourced point location could then be compared to the circle of uncertainty. If the crowdsourced point location was contained by the circle of uncertainty, then the crowdsourced point location and its associated location description could be said to be consistent with each other. If the crowdsourced point location was not contained by the circle of uncertainty, then the distance between the crowdsourced point location and the centre of the circle of uncertainty, minus the radius, could be said to be proportional to the inconsistency between the crowdsourced point location and its associated location description.

An important strength of the point-radius and the probability distribution methods is that they have been applied to free-text location descriptions stored in the Mammal Networked Information System (MaNIS) (Wieczorek et al. 2004, Guo et al. 2008). In addition, the point-radius method has been applied to free-text location descriptions that relate to search and rescue incidents (Doherty et al. 2011), which have clear similarities with those from the Haiti crisis map. Consequently, the secondary aim of the research reported in this chapter was to compare the degree to which the Haiti crisis map and the MaNIS datasets were similar.

In summary, the research reported in this chapter had both primary and secondary aims. The former was to investigate whether the types of location descriptions that were identified by Wieczorek et al. (2004) and Guo et al. (2008) could be found in location descriptions from the Haiti crisis map. The latter was to compare the degree to which the Haiti crisis map and the MaNIS datasets were similar. The methods that were used to address these aims are described in Section 3.1. The results are presented and discussed in Sections 3.2 and 3.3. Finally, conclusions are drawn in Section 3.4.

3.1 Methods

Three participants (P1, P2, and P3) independently classified all 3,606 location descriptions from the Haiti crisis map. In other words, each participant classified how each location was described. Each participant was given a table of 12 classes of location description (Table 3.1), which were derived from Wieczorek et al. (2004) and Guo et al. (2008). Each participant was also given a spreadsheet that contained all 3,606 location descriptions in random order. If a participant was unsure as to how to classify a location description, they were asked to classify it as *uncertain* and to comment on their reason for doing so. On average, each participant took four hours to complete the classification task.

Following the classification task, results were tabulated and Fleiss' kappa was used to compute the overall degree of agreement between participants, as well as that by class, corrected for agreement by chance alone (Fleiss 1971). The kappa statistic (κ) was computed using the *irr* library for R (Gamer et al. 2012). Although interpretations of κ are arbitrary, for convenience the ranges suggested by Landis & Koch (1977) are reported in this chapter.[1]

Finally, the overall degree of agreement between participants was considered sufficient enough to classify location descriptions by simple majority vote. When all participants disagreed, location descriptions were classified as *uncertain*. To allow comparison with those stored in the Mammal Networked Information System, all *uncertain* and *coordinates* cases were removed. This removal was unfortunate, but necessary: neither Wieczorek et al. (2004) nor Guo et al. (2008) report class frequency for the MaNIS dataset, and whilst Guo et al. (2008) report class proportion for the MaNIS dataset, the authors do so for the 12 classes in Table 3.1, minus uncertain and coordinates classes.

3.2 Results

Starting with class frequency by participant, all participants classified the majority of location descriptions as *feature* (F). In other words, all participants concluded

[1] Landis & Koch (1977) suggest six ranges of κ: *poor* ($\kappa < 0$); *slight* ($0 \leq \kappa < 0.2$); *fair* ($0.2 \leq \kappa < 0.4$); *moderate* ($0.4 \leq \kappa < 0.6$); *substantial* ($0.6 \leq \kappa < 0.8$); and *almost perfect* ($0.8 \leq \kappa \leq 1$).

Code	Class	Example
U	Uncertain	Presumably central Chile
C	Coordinates	42.4532 84.8429
F	Feature	Springfield
P	Path	Hwy. 1
J	Junction	Confluence of Labarge Creek and South Labarge Creek
FOH	Offset from a feature or path at a heading	10km N of Kuala Lumpur
NF	Near a feature or path	Big Bay vicinity
FS	Subdivision of a feature or path	N part of Mono Lake
FOO	Orthogonal offsets from a feature	1 miles N, 3 miles W of Fairview
FH	Heading from a feature, no offset	W of Tucson
FO	Offset from a feature or path, no heading	5km outside Calgary
BF	Between features or paths	Between Point Reyes and Inverness

Table 3.1: Classes of location description. Classes were derived from Wieczorek et al. (2004) and Guo et al. (2008).

Class	κ	Strength of agreement
C	0.776	Substantial
NF	0.532	Moderate
FOH	0.500	Moderate
FH	0.480	Moderate
BF	0.473	Moderate
F	0.459	Moderate
J	0.438	Moderate
P	0.378	Fair
FS	0.299	Fair
FO	0.186	Slight
U	0.169	Slight
FOO	0.000	Slight

Table 3.2: Degree of agreement between participants by class. Values in the *Strength of agreement* column follow Landis & Koch (1977).

that the majority of locations were described in terms of features alone (Figure 3.1). The next most common class for all participants was *path* (P). However, the differences between *feature* and *path* were considerable: P1, for example, classified 3,060 (85%) of location descriptions as *feature* but only 288 (8%) as *path*. *Near a feature or path* (NF) and *subdivision of a feature or path* (FS) were ranked either third or fourth for all participants (with one exception) but again the sizes of these classes were small in comparison to *feature*.

Fleiss' kappa indicated that there was *moderate* overall agreement between participants ($\kappa = 0.418$), although there was some variation between classes: agreement was *substantial* on the *coordinates* class ($\kappa = 0.776$) but less on all other classes (Table 3.2). This indicated that there was some *ambiguity*, or the doubt associated with the classification of a phenomenon (Fisher 1999), associated with the classification task. However, the overall degree of agreement between participants was considered sufficient enough to classify location descriptions by simple majority vote.

When location descriptions were classified by simple majority vote, *feature* was ranked first and *uncertain* (U) was ranked second (Figure 3.2). Once again, the difference between first and second ranked classes was considerable: 2,591 (72%) of location descriptions were classified as *feature* but only 376 (10%) were classified as *uncertain*. Very few location descriptions were classified as either *heading from a feature, no offset* (FH), *offset from a feature or path at a heading* (FOH) or *offset from a feature or path, no heading* (FO). No location descriptions were classified as *orthogonal offsets from a feature* (FOO). In other words, very few locations were described in terms of offsets and headings: the majority were described in terms of features alone.

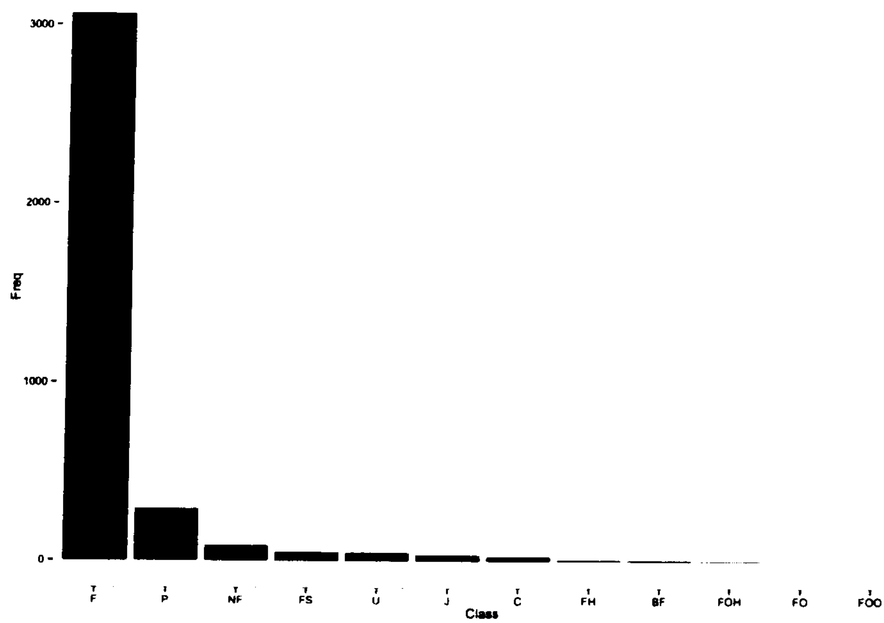

(a) P1

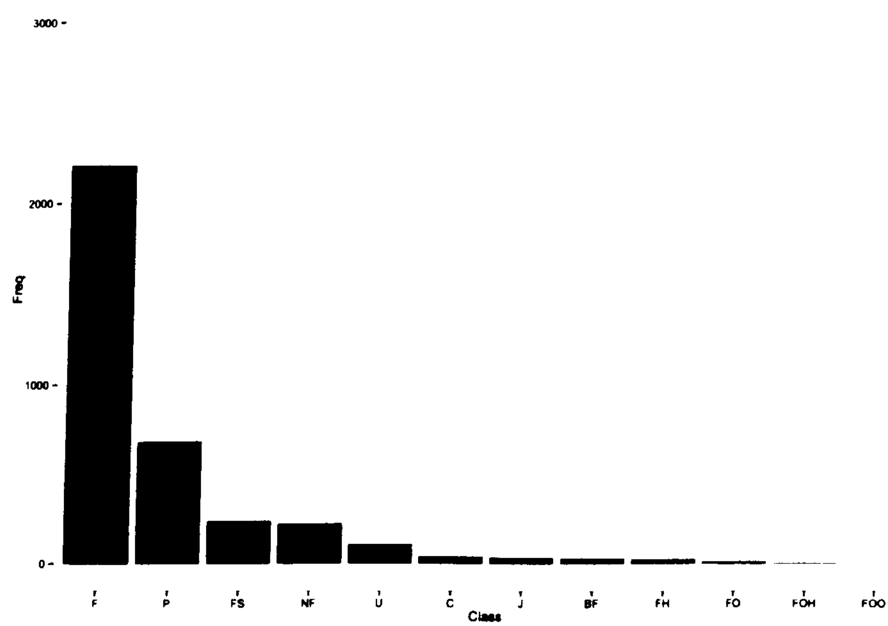

(b) P2

Figure 3.1: Class frequency by participant, Haiti crisis map.

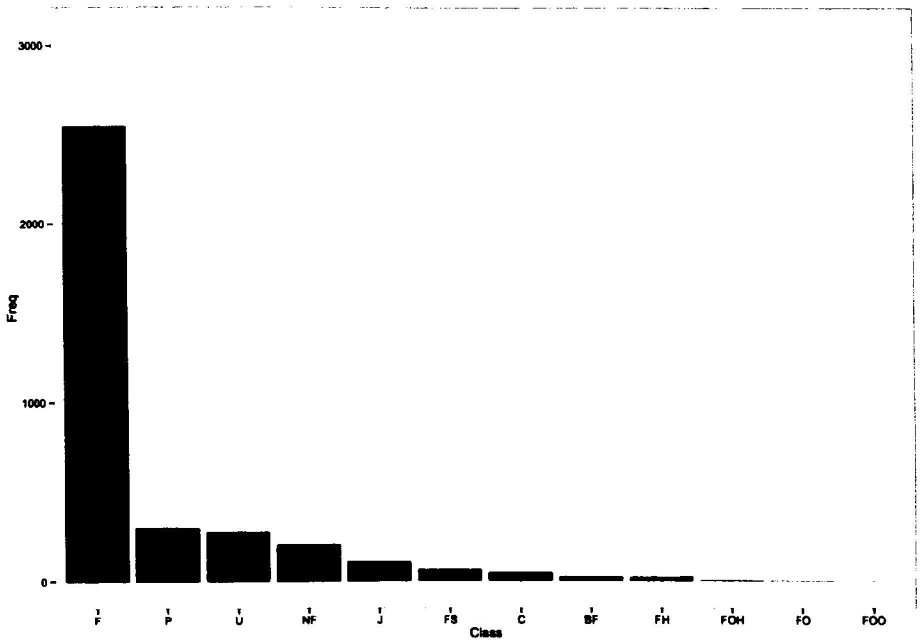

(c) P3

Figure 3.1: Class frequency by participant, Haiti crisis map.

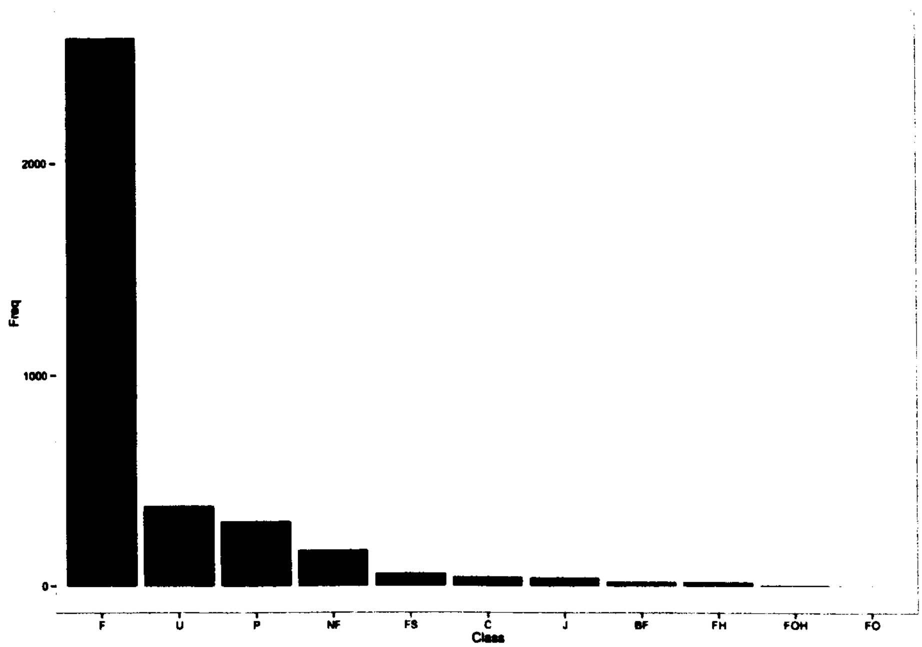

Figure 3.2: Class frequency, Haiti crisis map.

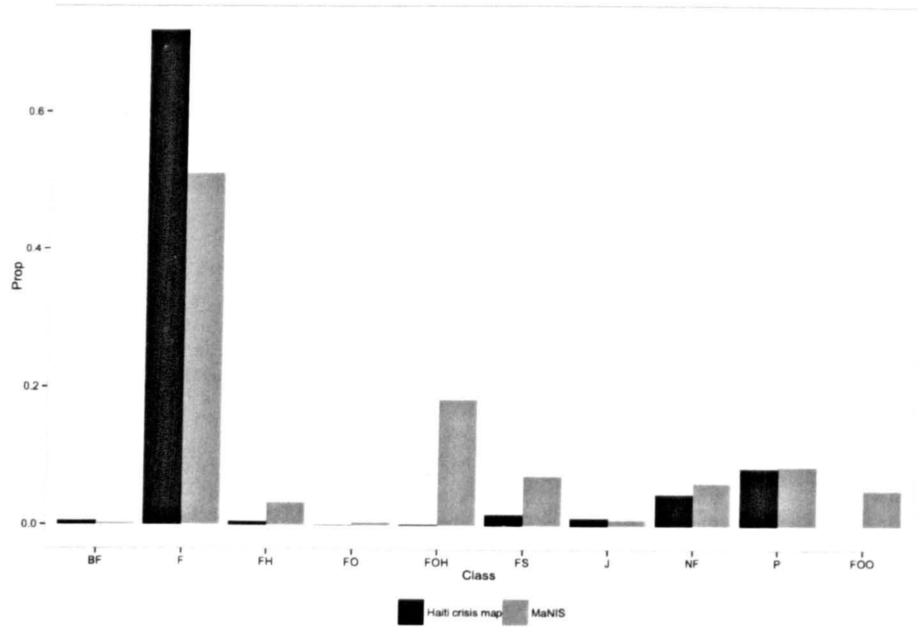

Figure 3.3: Proportion of location descriptions in each class, Haiti crisis map and the MaNIS dataset.

Turning to the question of the degree to which the Haiti crisis map and the MaNIS datasets were similar, it was clear that while *feature* was the largest class in both datasets, the proportion of locations that were described in terms of features was far higher in the Haiti crisis map than in the MaNIS dataset (Figure 3.3). Furthermore, the proportions of locations that were described in terms of offsets and headings (FH, FOH, FO, and FOO) were considerably lower in the Haiti crisis map than in the MaNIS dataset.

3.3 Discussion

An important limitation of the point-radius and the probability distribution methods is that they only account for certain types of uncertainty, namely, those types associated with reference objects and spatial relationships. The former group contains spatial extent, an unknown datum, imprecise coordinate measurements, and map scale; the latter group contains imprecise distance and direction measurements (Guo et al. 2008). The majority of locations from the Haiti crisis map were described in terms of features alone, and seldom in terms of offsets or headings. Consequently, the only type of uncertainty that was likely to be associated with location descriptions from the Haiti crisis map was that associated with the spatial extents of the reference objects. To be specific, the uncertainty associated with imprecise coordinate, distance, or direction measurements would rarely be associated with location descriptions, because so few locations were described in terms of coordinates, offsets, or headings. Similarly, the uncertainty associated with an unknown datum and map scale would rarely be associated with location descriptions, because the same web-based mapping service would almost certainly have been used to georeference location descriptions. For these reasons, applying the point-radius or probability distribution methods to location descriptions from the Haiti crisis map was clearly not worthwhile.

That the uncertainty associated with the spatial extents of the reference objects was likely to be important was, however, an interesting finding. More needed to be known about the composition of location descriptions from the Haiti crisis map. In other words, more needed to be known about the types of geographic objects (i.e. place names and spatial relationships) that were contained in these location descriptions and how these types were organised with respect to each other.

Finally, because only the proportions of location descriptions in each class were available for records stored in the Mammal Networked Information System, the degree to which the two datasets were similar could not be measured statistically. However, it was clear that the differences between the two datasets were sufficient to judge them to be dissimilar. The proportion of locations that were described in terms of features was far higher in the Haiti crisis map than in the MaNIS dataset, for example. Furthermore, the former had proportionally fewer location descriptions that contained offsets and headings.

3.4 Conclusion

The findings reported in this chapter suggest that location descriptions from the Haiti crisis map do not match the types of location descriptions identified by Wieczorek et al. (2004) and Guo et al. (2008) and, consequently, that location descriptions from the Haiti crisis map do not 'fit' the conceptual model proposed by Guo et al. (2008), which was discussed in Chapter 2, Section 2.2.1. Nevertheless, it is still useful to think of a location description as a collection of geographic objects and spatial relationships, although more information about the nature of these objects and relationships, and about how they are organised with respect to each other, is required to understand the uncertainty associated with crowdsourced crisis information. For this reason, an investigation into the composition of organisation of location descriptions from the Haiti crisis map is reported in Chapter 4.

Chapter 4

Investigating the Composition and Organisation of Location Descriptions

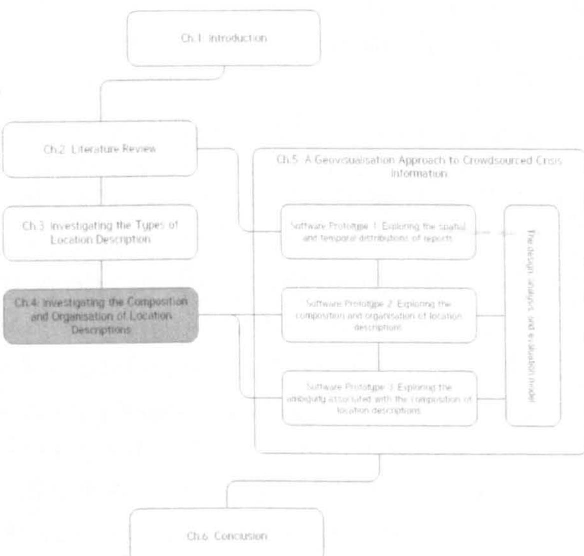

In Chapter 3, Section 3.3 it was noted that more needed to be known about the composition and organisation of free-text location descriptions from the Haiti crisis map. Consequently, the aim of the research reported in this chapter and elsewhere (Dillingham, Wood & Dykes 2013) was to investigate the types of geographic objects (e.g. villages, towns, and cities) and spatial relationships (e.g. terms such as *northeast*, *southernmost*, and *surrounding*) that were contained in free-text location descriptions, and to investigate how these types were organised with respect to each other (e.g. did villages precede towns? did spatial relationships succeed towns but precede cities?). This aim was decomposed into two objectives. The first was to identify and classify geographic objects and spatial relationships. The second was to summarise how these objects and relationships are organised with respect to each other. The methods that were used to address these objectives are described in Section 4.1. The results are presented and discussed in Sections 4.2 and 4.3. Finally, conclusions are drawn in Section 4.4.

4.1 Methods

The research reported in Chapter 3 adopted a manual approach to classifying all 3,606 location descriptions from the Haiti crisis map. However, classifying the more numerous geographic objects and spatial relationships contained in these location descriptions was likely to be prohibitively time-consuming and, as such, an automatic approach was used to undertake the research reported in this chapter. This approach is discussed in Section 4.1.2, where the various stages of a geoparsing system are described. As a precursor to this discussion, automatic approaches to geographic classification from natural language processing and geographic information science perspectives are described in Section 4.1.1.

4.1.1 Automatic approaches to geographic classification

Automatic approaches to geographic classification are based on two processes: identifying place names and resolving them to their referents (Grover et al. 2010). However, the terminology that describes these processes varies between disciplines. From the natural language processing perspective, both processes are called *named entity recognition* (Leidner 2007), which is part of the *information extraction* process (Jurafsky & Martin 2009). From the geographic information science perspective, both processes are called *geoparsing*, which is part of the *geocoding* and *georeferencing* processes (Grover et al. 2010). The former attempts to assign coordinates to address data; the latter, to "geographically relevant text" (Goldberg 2011, p.727).

Bird et al. (2009, p.261) present an architecture for a simple information extraction system (Figure 4.1). The key elements in this architecture are *tokenization*, where words are separated from the surrounding text (Jurafsky & Martin 2009), and *entity detection*, where the identification and resolution processes occur. During the identification process, words or word sequences are looked up in a place names dictionary. The resolution process is similar. However, rather than looking up words or word sequences in a place names dictionary, the place names themselves

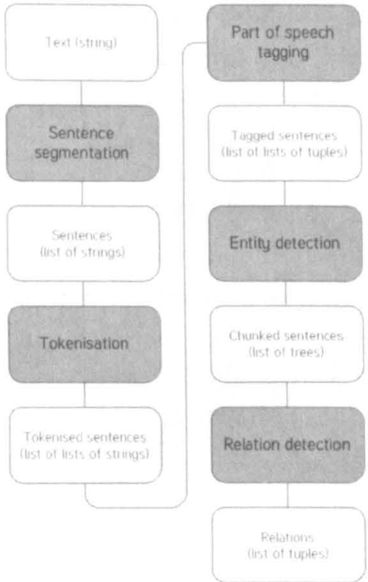

Figure 4.1: A simple information extraction system (Bird et al. 2009, p.261).

are looked up in a gazetteer. The distinction between a place names dictionary and a gazetteer is subtle but important: whilst the place names are unique in the former, they are not necessarily unique in the latter. Consequently, a gazetteer may contain several candidate entries for a single place name. The resolution process attempts to resolve this place name to a single candidate entry.

Whilst the resolution process can be complex (see, for example, Leidner 2007), Grover et al. (2010) suggest a simple approach: they compute a score using several heuristics, rank candidate entries by this score, and select the candidate entry that is ranked first. The heuristics suggested by Grover et al. (2010) are as follows:

- *Type of geographic object.* Some types are preferred to others; for example, *populated place* is preferred to *facility*.

- *Population.* Places with larger populations are preferred to those with smaller populations.

- *Context features.* Features in the surrounding text might provide a geographic focus (see below). A spatial relationship, such as a containment relationship, is an example of a context feature.

- *Geographic focus.* Candidates that are nearer to this focus are preferred to candidates that are further from it. The geographic focus may be provided by either a context feature (see above) or human judgement.

- *Geographic clustering.* Like Jones et al. (2008), Grover et al. (2010) argue that places contained in the same text tend to cluster. Candidates that are nearer to these clusters are preferred to candidates that are further from them.

In summary, automatic approaches to geographic classification require place names to be successfully identified and resolved. These processes can be situated within an architecture for a simple information extraction system, where they comprise the entity detection stage. Whilst the resolution process can be complex, heuristics have been developed to reduce this complexity. The discussion now turns to the implementation of these ideas: the automatic approach to geographic classification that was used to undertake the research reported in this chapter.

4.1.2 The geoparsing system

The architecture for a simple information extraction system presented by Bird et al. (2009, p.261) was used as the model for the geoparsing system. However, only the tokenization and entity detection elements of this system were implemented (Figure 4.2). This decision was made for two reasons. First, visual inspection of location descriptions suggested that locations were not described in full sentences. Furthermore, the tokens contained in location descriptions were likely to be geographic objects; that is, they were likely to be *proper nouns*, or "names of specific persons or entities" (Jurafsky & Martin 2009, p.159). For these reasons, the sentence segmentation and part of speech tagging elements were deemed unnecessary. Second, as the system was designed to geoparse—rather than to georeference—location descriptions, the relation detection element was also deemed unnecessary: the objectives of the research reported in this chapter were to investigate the composition and organisation of location descriptions, rather than to assign coordinates to location descriptions.

Despite the apparent simplicity of the geoparsing system, it nevertheless required several stages to identify and resolve place names successfully. The first and second stages were the creation of the gazetteer and the spatial relationships files, which can be thought of as the 'raw materials' for the geoparsing system. The third stage was the system for geoparsing location descriptions. The fourth stage was the system for determining the order of geographic objects in location descriptions. These stages—as well as the normalisation process, which was common to several stages—are described in the following sections.

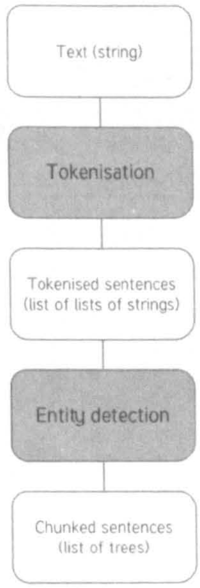

Figure 4.2: The geoparsing system.

Creating the gazetteer file

A subset of the GeoNames database (GeoNames 2013*a*) was used to create the gazetteer file. Data for all countries (*allCountries.zip*) and alternate place names (*alternateNames.zip*) were loaded into a database. The database schema can be found in Appendix A.

The gazetteer file was created in three stages. First, the name, GeoNames ID, and latitude and longitude coordinates for all GeoNames entries that were located in Haiti (`countryCode = "HT"`) were written to a file. Second, any English or French alternate names for these entries were written to a file.[1] Finally, the two files were concatenated. Table 4.1 shows the structure of the gazetteer file. It should be noted that there was not necessarily a one-to-one mapping between rows in this file and GeoNames entries. For example, GeoNames entry 3,723,988 (`id = 3723988`) appeared five times in the gazetteer file: once in its canonical form (*Republic of Haiti*), twice in its English form (*Haiti* and *Republic of Haiti*) and twice in its French form (*Haïti* and *République d'Haïti*). Indeed, there was not necessarily a one-to-one mapping between rows in this file and place names, as the previous example demonstrates. However, the gazetteer file referred to all place names in their canonical, and in their alternate English and French, forms.

The gazetteer file contained 15,730 entries, of which only 17 referred to alter-

[1] Haiti has two official languages: French and Haitian Creole (CIA 2013). However, only place names in the former were present in the GeoNames database. English was selected because it was the language of the Haiti crisis map.

Attribute	Description	Example
`id`	GeoNames ID	3731260
`name`	Place name	Abricots
`latitude`	Latitude coordinate	18.65000
`longitude`	Longitude coordinate	-74.30000
`featureCode`	GeoNames feature code	ADM3

Table 4.1: The structure of the gazetteer file.

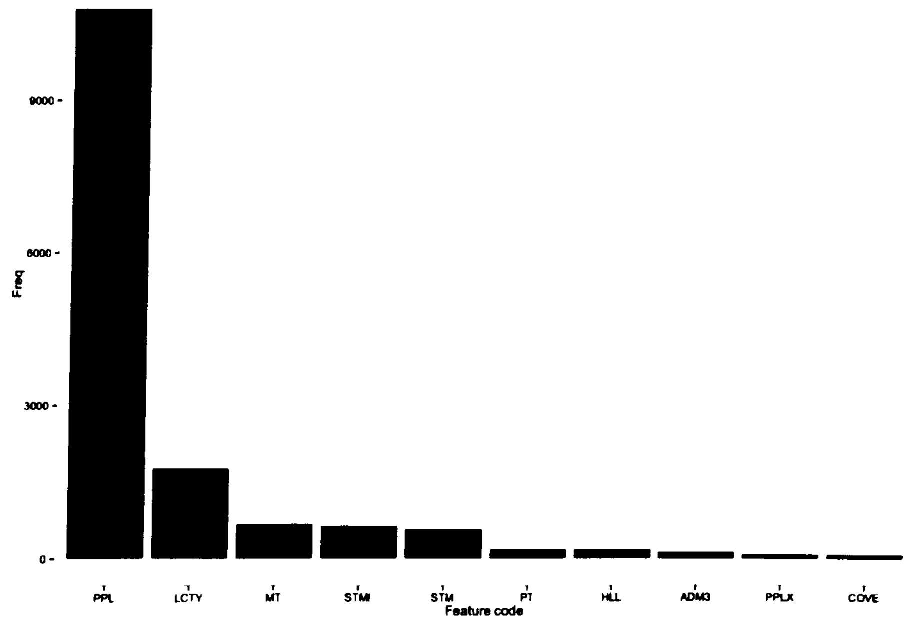

Figure 4.3: Frequency of the ten most common gazetteer feature codes.

nate place names. As can be seen in Table 4.1, each entry referenced a GeoNames feature code. Whilst there were 668 GeoNames feature codes at time of writing (GeoNames 2013*b*), only 94 were referenced by entries in the gazetteer file. The frequency of the ten most common gazetteer feature codes is shown in Figure 4.3. An alphabetical index to feature codes is shown in Table 4.2.

Creating the spatial relationships file

A Java application (*TokenFrequencyTool.java*) was written to normalise and tokenise location descriptions, and then to count the occurrences of each token. The normalisation process used by the application is described below. Source code for the application, which made use of several components from the LingPipe and Processing libraries (alias-i 2013, Processing 2013), can be found in Appendix A.

Visual inspection of the output of the application identified 22 terms that related to spatial relationships, such as prepositions (e.g. *above* and *after*) and directions (e.g. *northeast* and *southernmost*), as well as vaguer terms such as *surrounding* and *vicinity*. These terms were copied into the spatial relationships file.

The entries in the gazetteer and the spatial relationships files can be thought of as the 'raw materials' for the geoparsing system. The third and fourth stages of this system—a system for geoparsing location descriptions and a system for determining the organisation of geographic objects in location descriptions—are discussed below.

Geoparsing location descriptions

A Java application (*GeoparsingTool.java*) was written to identify the geographic objects and spatial relationships contained in location descriptions, and to resolve place names to single entries in the gazetteer file. The normalisation process used by the application is described below. Source code for the application, which made use of several components from the LingPipe and Processing libraries (alias-i 2013, Processing 2013), can be found in Appendix A. The application executed as follows:

1. Parse reports from the Haiti crisis map, and entries in the gazetteer and the spatial relationships files. Normalise location descriptions, place names, and spatial relationships.
2. Construct a dictionary from the normalised place names and the normalised spatial relationships.
3. Construct a chunker and chunk each normalised location description. The *chunker* was responsible for tokenising each normalised location description and for looking up each token in the dictionary. It was instructed to return only longest matches: for example, given dictionary entries of *Prince* and *Port-au-Prince*, and a token of *Port-au-Prince*, the chunker would return *Port-au-Prince*. It was instructed to ignore case: for example, the chunker would treat the tokens *Delmas* and *delmas* as if they were equal.
4. If the match is a place name, look up entries in the gazetteer file with the same place name. Resolve the match to a single entry in the gazetteer file.

To resolve the match to a single entry in the gazetteer file, matches were first compared according to three criteria and ranked in order of quality (best–worst). The match that was ranked first (best) was then returned. The three criteria were as follows:

1. *Complete equality.* If the matches referred to the same GeoNames entry, the matches were equal and the comparison was considered complete. If not, the comparison continued.

2. *Distance.* The distance in meters from each match to the point location of the report was computed and the distances were compared. If the difference between them was greater than ten meters, shorter distances were preferred to longer distances and the comparison was considered complete. If the difference between them was less than or equal to ten meters, the comparison continued. This threshold value was intended to prevent small differences—possibly artefacts of computation—affecting the ranking: in cases where two matches were less than or equal to this threshold value, the feature codes criterion would be applied (see below). The *distance* criterion is analogous to the *geographic focus* heuristic proposed by Grover et al. (2010).

3. *Feature codes.* The feature codes were compared. If the feature codes were equal, the matches were equal. If the feature codes were different, *locality* was preferred to *populated place*. After both outcomes, the comparison was considered complete. The *feature codes* criterion is analogous to the *type of geographic object* heuristic proposed by Grover et al. (2010).

The output of the application was a text file. Each line in this file contained the report ID, the start and end character index of the geographic object as it appeared in the location description, and the GeoNames feature code (the RLTNSHP feature code was used to reference a spatial relationship). Lines in this file were ordered by report ID and by start index. Consequently, given lines l and $l+1$, line $l+1$ would contain either details of a geographic object to the right of that contained in line l, or details of a geographic object contained in another location description.

Determining the organisation of geographic objects in location descriptions

A Java application (*PairTool.java*) was written to count the occurrences of pairs of feature codes contained in each location description. For example, given five geographic objects referenced by two location descriptions

```
id      start   end     featureCode
1022    0       6       RLTNSHP
1022    29      36      PPL
1024    24      30      ADM3
1024    32      46      PPLC
1024    48      53      PCLI
```

the application would count three pairs of feature codes

featureCode1	featureCode2	freq
RLTNSHP	PPL	1
ADM3	PPLC	1
PPLC	PCLI	1

Source code for the application, which made use of several components from the LingPipe and Processing libraries (alias-i 2013, Processing 2013), can be found in Appendix A.

Normalising location descriptions and place names

Several stages in the geoparsing system described above required location descriptions, place names, and spatial relationships to be normalised before they were processed. This was because visual inspection revealed that location descriptions contained escaped HTML characters (e.g. "&" for "&") and that accented characters had been used inconsistently in location descriptions: *Cité Soleil* (acute accent), for example, also appeared as *Cite Soleil* (no acute accent). To increase the number of matches when tokens were looked up in the dictionary, escaped HTML characters were converted to their unescaped equivalents and accented characters were converted to their unaccented equivalents. The `StringEscapeUtils.unescapeHtml4 (String input)` and `StringUtils.stripAccents(String input)` methods from the Apache Commons Lang library were used to perform these conversions systematically (The Apache Software Foundation 2013).

4.2 Results

Returning to the first research objective—to identify and classify geographic objects and spatial relationships—the most frequent feature code was clearly *populated place* (PPL, 2,309), with *third-order administrative division*, *independent political entity*, and *section of a populated place* (ADM3, PCLI, and PPLX; 775, 750, and 745) being less frequent. The feature code *capital of a political entity* (PPLC, 586) was less frequent still. However, the remaining feature codes were relatively infrequent. This group included *spatial relationship* (RLTNSHP, 115) and *locality* (LCTY, 110). Figure 4.4 shows the frequency of types of geographic objects contained in location descriptions from the Haiti crisis map in full (an alphabetical index to feature codes is given in Table 4.2).

Moving to the second research objective—to summarise how geographic objects and spatial relationships are organised with respect to each other—the most

Feature code	Description
ADM1	first-order administrative division
ADM3	third-order administrative division
AIRF	airfield
AIRP	airport
AIRQ	abandoned airfield
BAY	bay
CAPE	cape
CLF	cliff
COVE	cove
FLLS	waterfall
FT	fort
HLL	hill
HTL	hotel
ISL	island
LCTY	locality
LK	lake
MT	mountain
PCLI	independent political entity
PPL	populated place
PPLA	seat of a first-order administrative division
PPLC	capital of a political entity
PPLX	section of populated place
PT	point
ST	street
STM	stream
STMI	intermittent stream

Table 4.2: Alphabetical index to feature codes (GeoNames 2013*b*).

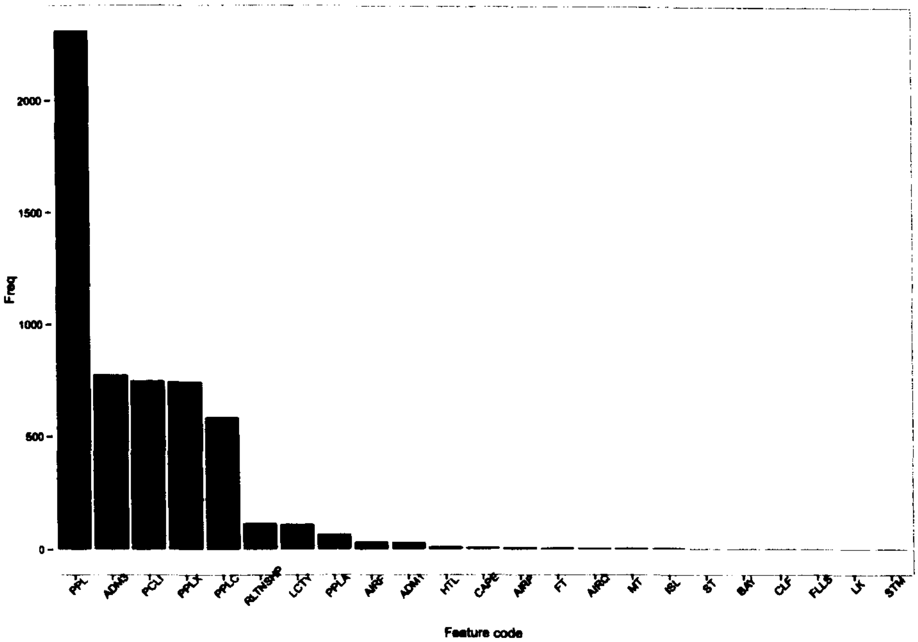

Figure 4.4: Frequency of types of geographic objects contained in location descriptions from the Haiti crisis map (an alphabetical index to feature codes is given in Table 4.2).

frequent pair of either objects or relationships was *populated place* paired with *populated place* (PPL–PPL, 454), such as “La Rica, Caldas” or “Delmas, Petion-Ville”; then *populated place* paired with *independent political entity* (PPL–PCLI, 332), such as “Citi Soleil, Haiti” or “Deschapelles, Haiti”; *section of a populated place* paired with *populated place* (PPLX–PPL, 129), such as “the 1st section of petit-goave, petit-goave” or “32 Impass Albin, Pelerin 5”; *capital of a political entity* paired with *independent political entity* (PPLC–PCLI, 122), such as “Port-au-Prince, Haiti”; and *populated place* paired with *third-order administrative division* (PPL–ADM3, 101), such as “Place Boyer, Petionville” or “Delmas 33, Petionville”. Figure 4.5 shows the frequency of paired types of geographic objects contained in location descriptions from the Haiti crisis map in full (an alphabetical index to feature codes is given in Table 4.2).

4.3 Discussion

Returning to the first research objective—to identify and classify geographic objects and spatial relationships—it is evident that location descriptions from the Haiti crisis map tended not to contain localities and, as such, vague places. The terms *local-*

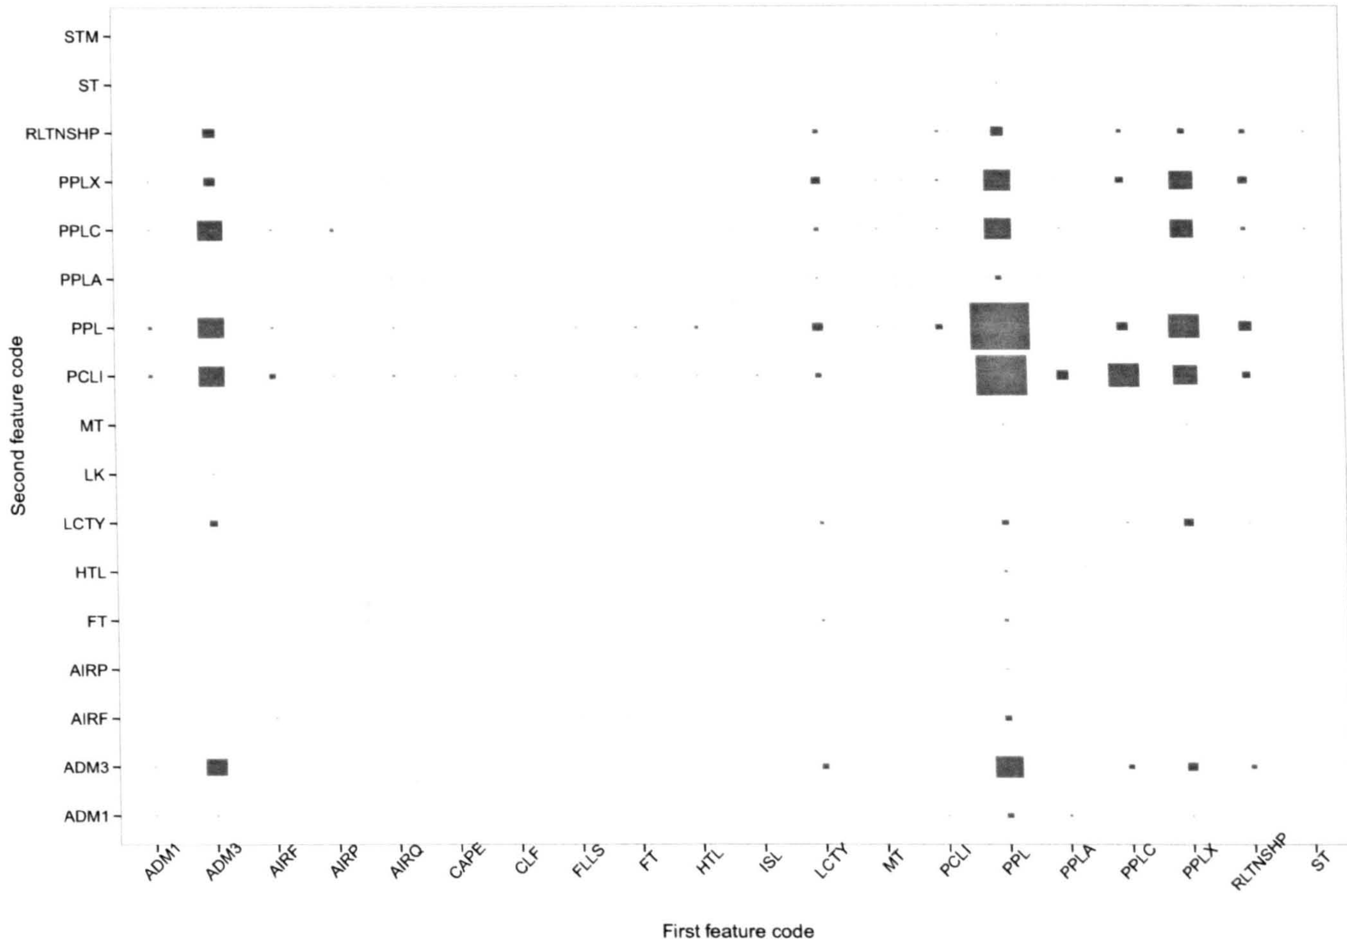

Figure 4.5: Frequency of paired types of geographic objects contained in location descriptions from the Haiti crisis map (an alphabetical index to feature codes is given in Table 4.2). The majority of geographic objects contained in location descriptions were populated places, as the largest rectangles lie on the PPL axes. Furthermore, populated places were often paired with countries, as the second-largest rectangle lies at the intersection of the PPL–PCLI axes.

ity and *vague place* refer, in essence, to the same type of geographic object: whilst a *locality* is defined as "a minor area or place of unspecified or mixed character and indefinite boundaries" (GeoNames 2013*b*), a *vague place* is defined as an area that has a vernacular name and a vague spatial extent (Jones et al. 2008). In Section 4.2, it was noted that location descriptions from the Haiti crisis map contained 110 localities, or 2.0% of the total number of geographic objects. In contrast, the Haiti subset of the GeoNames database contained 1,772 localities, or 11.3% of the total number of entries. However, comparing location descriptions from the Haiti crisis map to the Haiti subset of the GeoNames database is problematic: whilst the latter is a 'set' of the places in Haiti (duplicate places are not permitted), the former is simply a 'list' (duplicate places are permitted). Consequently, using the Haiti subset of the GeoNames database as a basis for comparison with location descriptions from the Haiti crisis map—using the Haiti subset of the GeoNames database to calculate expected values in a chi-square goodness of fit test, for example—would be erroneous. All that can be said is that location descriptions from the Haiti crisis map tended not to contain vague places.

In contrast to vague places, location descriptions from the Haiti crisis map tended to contain more, relatively coarse granularity types of geographic objects (i.e. more than relatively fine types of geographic objects). However, because a subset of the GeoNames database was used to create the gazetteer file, all the geographic objects that were classified as countries referred to Haiti. That repeated references to Haiti were made in the Haiti crisis map was surprising, as the geographic focus seemed clear. Nevertheless, whilst the additional information seemed redundant from the perspective of communications between the crisis-affected community and formal response organisations, it may have enhanced communications between the crisis-affected community and the wider world; perspectives that are discussed in Chapter 2.

Moving to the second research objective—to summarise how geographic objects and spatial relationships are organised with respect to each other—it is evident from Figure 4.5 that the majority of geographic objects contained in location descriptions were populated places and that populated places were often paired with countries. This pairing hinted at a finer-to-coarser granularity pattern, which was also exhibited in the populated places that were paired with capitals and the populated places that were paired with third-order administrative divisions. However, the existence of a finer-to-coarser granularity pattern between pairs of geographic objects did not necessarily indicate the existence of a finer-to-coarser granularity pattern across multiple geographic objects. This limitation raised two important questions:

1. To what degree are different patterns of geographic objects evident in location descriptions?

2. To what degree are the spatial and temporal distributions of location descriptions that contain different geographic objects different to each other?

These questions are addressed in Chapter 5 Section 5.4.2, where a geovisualisation software prototype for exploring the composition and organisation of location descriptions is discussed.

The research reported in this chapter highlighted two limitations with the Haiti subset of the GeoNames database. The first limitation was that many duplicate entries existed in the *geonames* table (i.e. entries with similar place names and point locations). Whilst it was clear from the database schema that the intention was to store the canonical form of a place name in the *geonames* table and any alternate forms in the *alternatenames* table, alternate forms were often stored alongside their respective canonical forms in the *geonames* table. These alternate forms were often classified as populated places, possibly because this was a more general feature code. The second limitation was that the feature codes did not explicitly measure spatial extent. Instead, spatial extent was inferred: populated places were seen as finer granularity than countries, for example. A better approach would have been to use the Haiti subset of the GeoNames database in combination with a classification of granularity (see, for example, Richter et al. 2012, Richter, Vasardani, Stirling, Richter & Winter 2013). This approach would have involved classifying the entries in the Haiti subset of the GeoNames database and geoparsing location descriptions using the natural language processing techniques described in Section 4.1.2. Alternatively, Richter, Winter, Richter & Stirling (2013) describe two algorithms for identifying the granularity of place names, which have the advantages of having being designed for, and evaluated using, location descriptions. Whilst the use of these algorithms would have been preferred, it should be noted that neither describes a geoparsing process. Consequently, an approach that combined the natural language processing techniques described in Section 4.1.2 with the two algorithms described by Richter, Winter, Richter & Stirling (2013) would still have been required.

A major limitation of the research reported in this chapter was the lack of an evaluation of the the geoparsing system described in Section 4.1.2. On the one hand, the lack of an evaluation meant that it was hard to assess the degree to which the results were sensitive to the identification and resolution processes: different threshold values or preferences could have changed the results, but whether for better or for worse was not known. On the other, it was not clear how the results should have been evaluated: whilst evaluation methods are well-documented in the natural language processing literature,[2] these evaluation methods rely on the creation of an accurate, or *gold standard*, dataset. Creating a gold standard dataset for location descriptions from the Haiti crisis map would have meant manually classifying a representative sample of location descriptions. However, in the absence of information about the types of geographic objects that were contained in these location descriptions, it was hard to determine how large a representative sample of location descriptions should be. Manually classifying all 3,606 location descriptions clearly removed the need for a gold standard dataset. Furthermore, there was no reason to believe that location descriptions from the Haiti crisis map were representative of all sources of crowdsourced crisis information.

In response to the major limitation described above, it is argued that there is a role for geovisualisation for investigating how well geographic objects were classified; that is, for evaluating the geoparsing system that is described in Section 4.1.2. This role is described in Chapter 5 Section 5.4.3, where a geovisualisation software prototype for exploring the *ambiguity*, or the doubt associated with the classification of a phenomenon (Fisher 1999), associated with the composition of location descriptions is discussed.

4.4 Conclusion

The findings reported in this chapter suggest that location descriptions from the Haiti crisis map tended not to contain vague places; tended to contain more, relatively coarse granularity types of geographic objects (i.e. more than relatively fine types of geographic objects); and tended to be paired in a finer-to-coarser granularity pattern. However, these findings are extremely tentative and this chapter raises more questions than it answers. In Section 4.3 two geovisualisation software prototypes are referred to that seek to address some of these questions. These applications are discussed in Chapter 5, Sections 5.4.2 and 5.4.3.

[2]Measures of precision and recall are often used to evaluate named entity recognition systems. *Precision* is the percentage of chunks that the system identified correctly, where correctness is measured in terms of the start and end character index and the class for each chunk. *Recall* is the percentage of chunks that were present in the text that the system identified correctly (Jurafsky & Martin 2009).

Chapter 5

A Geovisualisation Approach to Crowdsourced Crisis Information

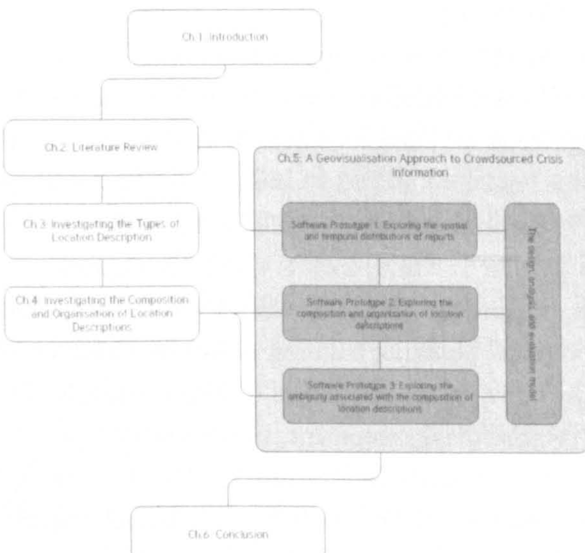

The merits of the interactive visual exploration of information are well documented (see, for example, MacEachren 1994*b*, Card et al. 1999, Dykes et al. 2005, Thomas & Cook 2005, 2006, Keim et al. 2010). However, traditional geographic information systems (GISs), which would normally support the interactive visual exploration of geographic information, were developed for space-based, rather than place-based, representations of geographic information (Fisher & Unwin 2005, Gregory & Hardie 2011) and so are poorly suited to crowdsourced crisis information. Consequently, three geovisualisation software prototypes were developed to support the interactive visual exploration of reports from the Haiti crisis map: their goals were to explore the spatial and temporal distributions of reports; to explore the composition and organisation of location descriptions; and to explore the ambiguity associated with the composition of location descriptions. These software prototypes—and the processes, technologies, principles, and techniques that surrounded them—are the subject of this chapter.

Collectively, the software prototypes represent a geovisualisation approach to crowdsourced crisis information. This approach is characterised by the overlap between the context of use and the context of design: a situation where the designer is also the user of a software artefact. These contexts are discussed in Section 5.1. This situation raises the important question of how a software artefact should be evaluated when the context of use and the context of design overlap. The design, analysis, and evaluation model, which is discussed in Section 5.1.4 and elsewhere (Dillingham et al. 2012*b*), is offered as an answer to this question: it is a synthesis of ideas from the fields of action research, human–computer interaction, and visualisation, and is an iterative approach to the design, implementation, and evaluation of a software artefact.

The development of a software artefact should be grounded in an appraisal of appropriate technologies, principles, and techniques. Consequently, the two visualisation technologies that were considered for developing the software prototypes—Processing (Reas & Fry 2006) and D3 (Bostock et al. 2011)—are discussed in Section 5.2. Similarly, the principles and techniques that unify the software prototypes are discussed in Section 5.3. In Section 5.4, the design, analysis, and evaluation model is used to structure the discussion of each software prototype. In each case, the working hypotheses, scenario, features, design decisions, and findings are reported as evidence of an iterative approach to design, implementation, and evaluation. However, because the development of each software prototype was itself an evaluation of the design, analysis, and evaluation model, the utility of the model is reflected on in Section 5.5. Finally, conclusions are drawn in Section 5.6.

5.1 Evaluation and design

The relationship between evaluation and design is discussed in this section. First, this relationship is discussed from the visualisation perspective in Section 5.1.1. The design study methodology (Sedlmair et al. 2012) is presented and examples of design studies are highlighted. Second, this relationship is discussed from the action research and human–computer interaction perspectives in Section 5.1.2. Similarities between action research, scenario-based design, and the design study methodology are then identified in Section 5.1.3. Finally, the design, analysis, and evaluation model is discussed in Section 5.1.4. The model situates the activities associated with designing a software artefact—and using it to undertake analysis—within an evaluative framework: it is an iterative approach to the design, implementation, and evaluation of a software artefact that attempts to bridge the gap between the context of use and the context of design.

5.1.1 The design study methodology: bridging the gap between the context of use and the context of design?

The concept of evaluation encompasses the concepts of *validation* and *verification*: concepts that are summarised by the questions *Are we building the right product?* and *Are we building the product right?* (Munzner 2009). Addressing these questions helps the visualisation researcher avoid the many pitfalls of visualisation research (Munzner 2008). However, the visualisation literature has emphasised the concept of validation, possibly because it has been influenced by the human–computer interaction literature. Consequently, many evaluation methods are user-centred: they range from quantitative approaches based on controlled experiments, such as usability evaluation (Greenberg & Buxton 2008), user studies (Kosara et al. 2003), and insight reports (Saraiya et al. 2005, North 2006) to qualitative approaches, such as grounded evaluation (Isenberg et al. 2008).

Recent interest in qualitative approaches to evaluation has emphasised the context of use (Isenberg et al. 2008). Indeed, Greenberg & Buxton (2008, p.111) argue that controlled experiments that ignore the context of use can be "harmful". However, whilst participatory approaches to design attempt to narrow the gap between the context of use and the context of design (see, for example, Koh et al. 2011, Lloyd & Dykes 2011), qualitative approaches to evaluation maintain this gap: the two contexts, and their respective practitioners, are seen as completely separate. Sedlmair et al. (2012), for example, argue that a design study should report how visualisation researchers analysed a problem that was faced by domain experts; how they designed a visualisation to address this problem; how they validated the design with the domain experts; and, through reflection, how they identified guidelines that can be applied to similar problems in different domains in the future. Sedlmair et al. (2012) locate these activities in nine stages, which they group in three phases: they link each successive stage to each preceding stage (Figure 5.1). The three phases and nine stages of the design study methodology (Sedlmair et al. 2012) are as follows:

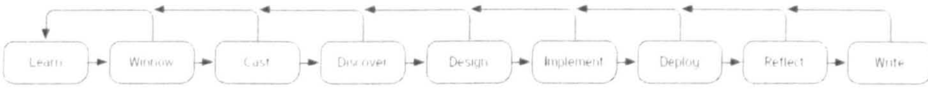

Figure 5.1: The design study methodology (Sedlmair et al. 2012). Each successive stage is linked to each preceding stage. For clarity, only links to the *learn* stage are shown.

1. *Precondition*

    (a) *Learn.* The researcher assimilates the visualisation literature related to, for example, visual encodings and interactions.

    (b) *Winnow.* The researcher identifies promising collaborations. Practical, intellectual and interpersonal issues should be considered.

    (c) *Cast.* The researcher identifies roles: front line analysts, gatekeepers, connectors, translators, co-authors, and tool builders, for example.

    (d) *Discover.* The researcher learns enough about the domain to be able to characterise and abstract it.

2. *Core*

    (a) *Design.* The researcher designs a wide range of visual encodings and interactions for implementation.

    (b) *Implement.* The researcher implements a narrow range of visual encodings and interactions for deployment.

    (c) *Deploy.* The researcher elicits feedback from the users of the visualisation tool in a real world context.

3. *Analysis*

    (a) *Reflect.* The researcher confirms, refines, extends, or rejects existing, or proposes new, guidelines.

    (b) *Write.* The researcher revisits characterisations and abstractions from the discover stage.

The visualisation literature contains several examples of design studies that follow the methodology proposed by Sedlmair et al. (2012). For example, Nielsen et al. (2009) discuss a software application they developed to support *genome sequencing*, or the process of identifying the DNA sequence of an organism's genome. Similarly, Meyer et al. (2009, 2010) discuss two software applications—MizBee and MulteeSum—they developed to support *comparative genomics*, or the study of how genome structure and function differ between species. Finally, working with historical geographers, Weaver et al. (2007) developed a software application for exploring spatio-temporal patterns in human behaviour. In each case, the domain experts' research questions were used to characterise their domains. These research questions were then abstracted and these abstractions were used to inform the choice of visual encodings and interactions.[1] Each design was then validated by the domain experts: Weaver et al. (2007) and Nielsen et al. (2009) were guided by informal feedback, whilst Meyer et al. (2009, 2010) discuss how MizBee and MulteeSum caused the domain experts to change their approach to their research. The authors then identify guidelines. Weaver et al. (2007), for example, describe a technique that supports the detection of regular and irregular cyclical events. By drawing parallels between the problems faced by historical geographers, and the problems faced by intelligence analysts and emergency managers, Weaver et al. (2007) demonstrate how this technique can be applied to similar problems in different domains in the future.

The visualisation literature also contains several examples of design studies that do not follow the methodology proposed by Sedlmair et al. (2012) in the sense that the context of use and the context of design, and their respective practitioners, overlap. In other words, whilst the visualisation researchers are not necessarily domain experts, they are nevertheless sufficiently expert in the domain to be able to contribute to it. In these cases, the visualisation users are the visualisation designers. For example, Wood et al. (2011) used spatially-ordered hierarchical layouts (Wood & Dykes 2008, Slingsby et al. 2009) to find evidence of position bias in the results of local elections. The authors found that, under certain circumstances, the position (alphabetical order) of a candidate on the ballot paper affected the position (numerical order) of the candidate in the local election. Whilst none of the authors are political scientists, they argue that the findings contradict the political science literature and, consequently, validate spatially-ordered hierarchical layouts.

The approach adopted by Wood et al. (2011) supports the argument made by van Wijk (2006) that visualisation researchers can successfully bridge the gap between the domain and the visualisation areas of expertise (i.e. the context of use and the context of design) by exploring the domain area of expertise themselves. Indeed, the reverse is also true: as libraries such as Processing (Reas & Fry 2006) and D3 (Bostock et al. 2011) make it easier to develop visualisations, domain experts are increasingly able (and willing) to design and implement tools that address their own research questions (see, for example, Foley & Demšar 2013). However,

[1] The progression from *domain characterisation*, to *domain abstraction*, to *visual encoding and interaction design* is the basis of Munzner's nested model for visualisation design and evaluation (Munzner 2009), which is discussed in Section 5.5.

the visualisation literature is less clear about how these tools should be evaluated. Munzner (2009) suggests that upstream, a design could be justified with reference to what is known about human perception and cognition, using expert reviews (Tory & Moller 2005) or heuristic evaluations (Zuk et al. 2006), for example. Alternatively, a design could be justified in a more discursive fashion: this approach is taken by Wood et al. (2007) to reflect on the strengths and weaknesses of their map mashup and could be seen as part of the critical dialogue advocated by Kosara et al. (2008). Munzner (2009) suggests that downstream, results (i.e. images or videos) could be presented and discussed as a case study. This approach is taken by Wood et al. (2011), who used spatially-ordered hierarchical layouts (Wood & Dykes 2008, Slingsby et al. 2009) to find evidence of position bias in the results of local elections. However, the implications of these suggestions are not considered in the visualisation literature.

Addressing the question of how a software artefact should be evaluated when the context of use and the context of design overlap was deemed worthwhile as the interactive visualisation of reports from the Haiti crisis map was central to the current research. Two perspectives—action research and scenario-based design—connected with the visualisation literature and offered some insight into this question. These perspectives are considered in the next section. Similarities between these perspectives and the design study methodology are discussed in Section 5.1.3.

5.1.2 Action research and scenario-based design

As Baskerville & Wood-Harper (1996) explain, action research originated independently in the USA and the UK in the immediate aftermath of the Second World War. In the USA, Kurt Levin and colleagues at the Research Centre for Group Dynamics at the University of Michigan were developing a general theory of social change. In the UK, researchers at the Tavistock Institute in London were treating patients suffering from disorders that related to their experiences of combat and prisoner-of-war camps. The Tavistock Institute researchers found that universal treatments were illusive, as the causes of the disorders were very different. Consequently, they developed a body of knowledge by reflecting on their own practice.

Contemporary forms of action research are often adopted by practitioners who wish to improve either their own practice, or that of others with whom they collaborate (Oates 2006). Action research is especially useful where the problem is hard to define, as it is approached through iterative cycles of planning, acting, and reflecting (Oates 2006). Baskerville & Wood-Harper (1996) state that a typical action research cycle has five stages (Figure 5.2). These stages are as follows:

Figure 5.2: A typical action research cycle (Baskerville & Wood-Harper 1996).

1. *Diagnosing.* The research team defines a problem (a working hypothesis), situating the problem within a theoretical framework.
2. *Planning.* The research team, guided by the theoretical framework, considers alternative courses of action for solving the problem.
3. *Taking action.* The research team works together to implement the planned course of action.
4. *Evaluating action.* The research team study the consequences of taking action with respect to the theoretical framework.
5. *Specifying learning.* The research team identify general principles that can inform future action research cycles, as well as the wider scientific community.

Like action research, scenario-based design attempts to bridge the gap between research and practice (Carroll & Rosson 1992). In this approach to design, which can be seen as the application of action research to the field of human–computer interaction (Carroll & Rosson 1992), the designer writes a scenario (a narrative) that describes how a user interacts with a system (Rosson & Carroll 2002). The scenario helps the designer identify the features of the system. The designer can then make claims about these features, where a claim establishes a causal relationship between a feature and its positive or negative psychological consequences (Carroll & Rosson 1992). The designer can also justify these features with reference to the literature (Carroll & Rosson 1992). Carroll (2000) identifies several benefits of scenario-based design, the most significant of which is that it encourages "reflection in the context of design" (Carroll 2000, p.47). In other words, a design can be developed and reflected upon simultaneously: design-activity does not compete with reflection-activity and the focus remains on the context of use (Carroll 2000).

5.1.3 Similarities between action research, scenario-based design, and the design study methodology

Sedlmair et al. (2012) argue that the goal of the design study methodology is to produce transferable, rather than reproducible, findings. A researcher engaged in an action research project would welcome this goal: action research is an "interventionist" approach and, as such, "to an arch positivist it should seem very unscientific" (Baskerville & Wood-Harper 1996, p.236). This apparent lack of rigour, however, is seen as a considerable advantage by researchers engaged in action research projects: Susman & Evered (1978) argue that science—in the positivist tradition—is not suited to situations where the research subjects can reflect on the analysis, where they are influenced by the definition of the problem, and where the research is directed towards solving a problem the research subjects have helped to define. Consequently, a researcher engaged in an action research project would be quite happy to accept the role played by judgement, by conjecture—"leaps of the imagination" (Susman & Evered 1978, p.598)—and by the researcher and the research subjects collaborating with each other (Susman & Evered 1978). These characteristics are also fundamental to the design study methodology.

Like a typical action research cycle, the problems addressed by the design study methodology often lack definition and are approached through iterative cycles of planning, acting, and reflecting. However, Sedlmair et al. (2012) identify nine stages, which they group in three phases, and link each successive stage to each preceding stage (Figure 5.1). In contrast, Baskerville & Wood-Harper (1996) state that a typical action research cycle has five stages, and link each successive stage only to its predecessor (Figure 5.2). Consequently, there is far more iteration within the design study methodology than within a typical action research cycle. Nevertheless, it is clear that the three phases proposed by Sedlmair et al. (2012) within the context of the design study methodology closely correspond to the five stages proposed by Baskerville & Wood-Harper (1996) within the context of a typical action research cycle. The precondition phase, where literature is assimilated and the domain is characterised and abstracted, is similar to the diagnosing stage; the core phase, where designs are proposed, selected, and validated, is similar to the planning, taking action, and evaluating action stages; and the analysis phase, where designs are reflected upon and guidelines are communicated, is similar to the evaluating action and specifying learning stages.

There are differences, however, between the design study methodology and a typical action research cycle. The first, and arguably the most important, difference concerns the question of whose practice is being improved. In a typical action research cycle, the *problem situation* (Oates 2006) is not abstracted. Consequently, the researcher and the domain expert share the problem situation. In contrast, in the design study methodology there are two areas of expertise—the domain area and the visualisation area—and, as such, two problem situations. On the one hand, this duality explicitly recognises that the domain expert and the visualisation researcher come from different backgrounds and scientific traditions. On the other, this duality could cause tension.

The second difference between the design study methodology and a typical action research cycle is that the design study methodology requires the sustained, collaborative involvement of domain experts and visualisation researchers: according to Sedlmair et al. (2012, p.2432), the roles are distinct:

> *"While strong problem-driven work can result from situations where the same person holds both of these roles, we do not address this case here."*

In contrast, whilst an action research project is often characterised by a researcher and a practitioner working collaboratively (Baskerville & Wood-Harper 1996), it is possible for one person to hold both roles simultaneously (Susman & Evered 1978).

The acceptance within action research that one person can hold both researcher and practitioner roles simultaneously, coupled with the similarities between the design study methodology, action research, and scenario-based design, suggested that ideas from each perspective could be combined to address the question of how a software artefact should be evaluated when the context of use and the context of design overlap. The design, analysis, and evaluation model is offered as an answer to this question: one possible synthesis of ideas from these perspectives.

5.1.4 The design, analysis, and evaluation model

The design, analysis, and evaluation model, which is discussed in this section and elsewhere (Dillingham et al. 2012*b*), situates the activities associated with designing a software artefact—and using it to undertake analysis—within an evaluative framework: it is an iterative approach to the design, implementation, and evaluation of a software artefact that attempts to bridge the gap between the context of use and the context of design. The model combines the concepts of scenarios, features, and justifications from scenario-based design with the stages of a typical action research cycle (Figure 5.3). The activities that comprise this cycle are classified as design, analysis, or evaluation activities. They are as follows:

1. *Hypothesising*. The researcher formulates a working hypothesis (an analysis activity).

2. *Planning action*. The researcher writes a scenario, identifies and justifies features (design activities), and formulates a development plan.

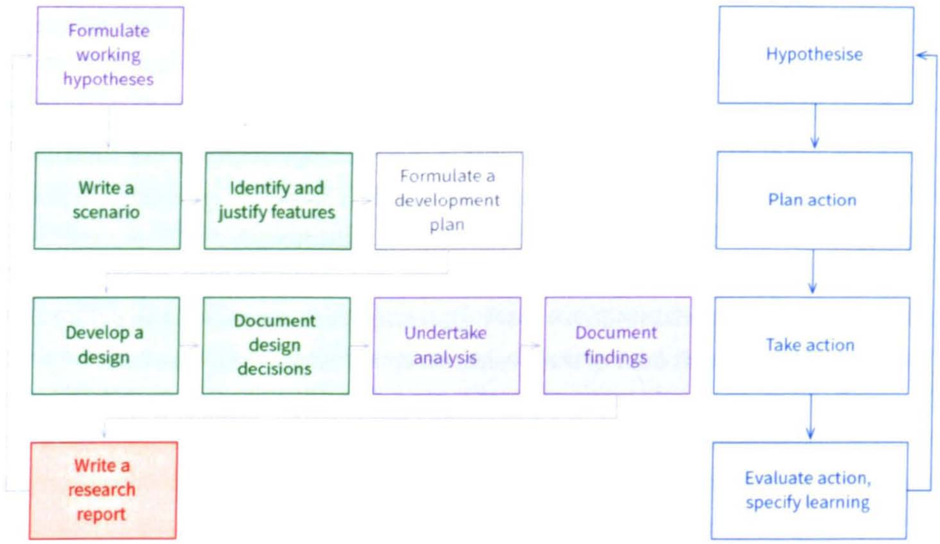

Figure 5.3: The design, analysis, and evaluation model (left), and a typical action research cycle (right) (Baskerville & Wood-Harper 1996). Analysis activities are shown in purple, design activities in green, and evaluation activities in orange.

3. *Taking action.* The researcher develops a design and documents design decisions (design activities). The researcher undertakes analysis and documents findings (analysis activities).

4. *Evaluating action and specifying learning.* The designer writes a research report (an evaluation activity).

The design, analysis, and evaluation model was used to develop—and in Section 5.4 is used to structure the discussion of—the three geovisualisation software prototypes. In other words, the model was adopted as the methodology for designing, implementing, and evaluating the three geovisualisation software prototypes that are discussed in Section 5.4. However, the development of a software artefact should be grounded in an appraisal of appropriate technologies, principles, and techniques. Consequently, the two visualisation technologies that were considered for developing the software prototypes—Processing (Reas & Fry 2006) and D3 (Bostock et al. 2011)—are discussed in Section 5.2. Similarly, the principles and techniques that unify the software prototypes are discussed in Section 5.3.

5.2 Visualisation technologies

The two visualisation technologies that were considered for developing the software prototypes are discussed in this section. The first technology was Processing

(Reas & Fry 2006), which is a programming language and a library for Java. The second technology was D3 (Bostock et al. 2011), which is a library for JavaScript. Processing and D3 can be seen as domain-specific languages (Mernik et al. 2005), as they support the development of visualisations by providing additional functionality that is not readily available in their underlying programming languages (see, for example, Fry 2007, Dewar 2012). Nevertheless, each also enjoys access to these programming languages, which increases their flexibility. Each technology is discussed individually in the following two sections before conclusions are drawn on the choice of visualisation technology in the third section.

5.2.1 Processing/Java

Processing/Java offer several advantages, not least because they are established technologies: Processing was first released in 2001 and Java in 1995. Over the last 18 years, many libraries have been written that provide additional functionality for Java: several components from the LingPipe library (alias-i 2013), for example, were used in Chapter 4. Processing enjoys access to these libraries, as well as to those that have been written specifically for visualisation, such as the giCentre Utilities and GeoMap libraries (Wood et al. 2013, Wood & Dillingham 2012), which provide statistical and geographic functionality.

Processing offers a simple mechanism for creating shapes, such as rectangles and ellipses, and for controlling their properties, such as their colour. Until relatively recently, a further advantage of Processing was the ability to distribute sketches as *applets*, or applications that could be run in a web browser. Recently, however, several popular web browsers have disabled applets, making sketches harder to distribute in this way. Nevertheless, sketches can still be distributed as applications for the three most common families of operating systems.

5.2.2 D3/JavaScript

D3 provides additional functionality for JavaScript that makes it easier to manipulate elements structured using any XML format, such as SVG, and styled using CSS.[2] The library also provides a simple mechanism for transitioning between states (Heer & Robertson 2007). Whilst D3 supports the creation of common visualisation components (e.g. axes) as well as many types of visualisation (e.g. force-directed graphs, chord diagrams, and treemaps), early releases neglected es-

[2] *XML*, or Extensible Markup Language, defines rules for encoding information in a format that can be read by humans and computers alike. *SVG*, or Scaleable Vector Graphics, is an XML format for vector graphics. *CSS*, or Cascading Style Sheets, is a language for encoding the style of elements in XML or SVG documents, such as their colour.

sential geographic functionality. For example, although it was relatively easy to create maps, it was much harder to manipulate them: zooming and panning were not supported at this time.

D3 is the most recent result of work undertaken by the Stanford Visualization Group and was preceded by a similar library called Protovis (Bostock & Heer 2009). The speed at which development on Protovis was halted in favour of development on D3 was of some concern, although this concern was allayed by the speed at which the developer community adopted D3. Indeed, as an open-source project, D3 has benefited from the contributions of individuals from outside the Stanford Visualization Group. These contributions include some of the essential geographic functionality that was lacking in early releases (Bostock & Davies 2013).

The merits of JavaScript are debatable: it certainly takes concerted effort to avoid what Crockford (2008, p.109) calls its "bad parts". In addition to problems associated with its design, there are also problems associated with its various implementations. Browser inconsistencies frequently cause frustration, although some areas are more affected than others. The handling of mouse and key events is one such area. As these events provide the main mechanisms for adding interaction to a visualisation, these inconsistencies were of some concern.

Finally, when compared to Java, JavaScript has limited tool support. Whilst JavaScript IDEs (integrated development environments) are available, none rival their Java counterparts. However, tools are available that support the development of JavaScript in a browser, making for a more interactive development process.

5.2.3 Choosing a visualisation technology

A software prototype was developed to evaluate the suitability of D3/JavaScript. This software prototype is shown in Figures 5.4 and 5.5 and is reported elsewhere as an early introduction to the current research (Dillingham, Dykes & Wood 2011). It was found that the benefits of this technology combination did not outweigh the risks. For example, it became clear that implementing essential geographic functionality, such as zooming and panning, was more time-consuming than anticipated, not least because doing so required a thorough understanding of the D3, JavaScript, and SVG APIs (Application Programming Interfaces). Similarly, as the requirement to use information extraction techniques emerged, the ability to integrate additional functionality, such as components from the LingPipe library (alias-i 2013), became apparent. For these reasons, as well as a positive previous experience of developing geovisualisation software prototypes in Processing/Java (Dillingham, Mills & Dykes 2011), the three geovisualisation software prototypes that are discussed in Section 5.4 were developed in Processing/Java.

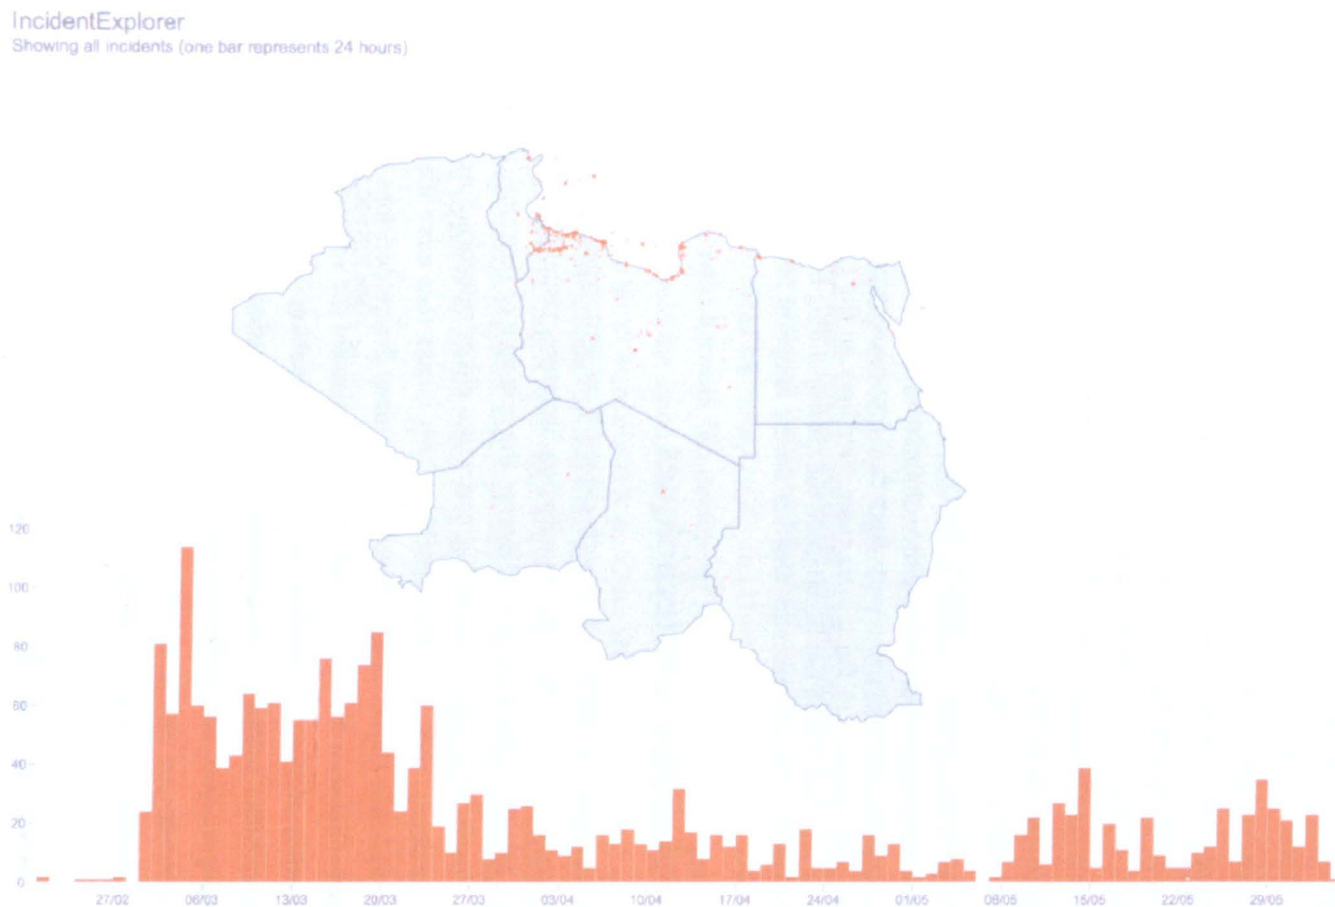

Figure 5.4: Overview of the software prototype that was developed to evaluate the suitability of D3/JavaScript. The map and the temporal histogram depict reports from the Libya crisis map (OCHA 2011). One bar on the temporal histogram represents 24 hours.

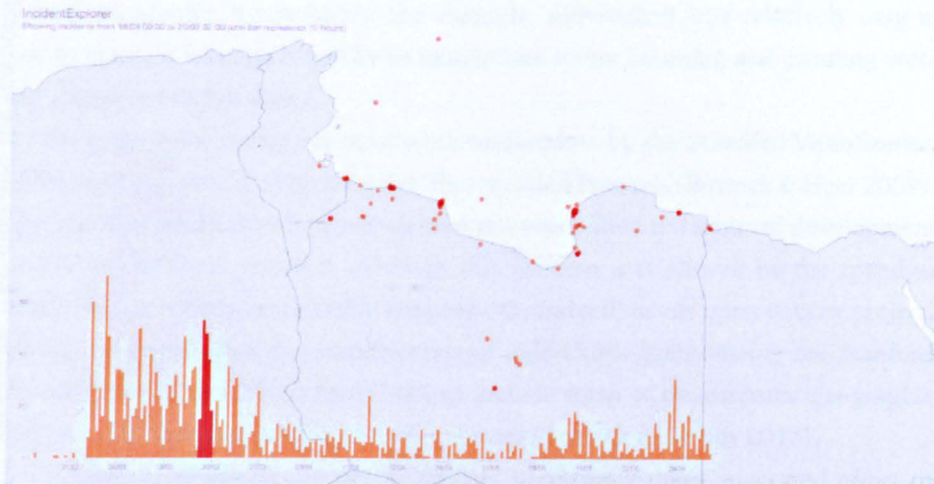

Figure 5.5: Selecting and zooming in the simple geovisualisation software prototype that was developed to evaluate the suitability of D3/JavaScript. The map and the temporal histogram depict reports from the Libya crisis map (OCHA 2011). The map has been zoomed to the Libyan coast. Five bars on the temporal histogram have been selected, which has highlighted the corresponding bars on the temporal histogram and points on the map. One bar on the temporal histogram represents 10 hours.

5.3 Principles and techniques

The principles and techniques that unify the software prototypes are discussed in this section. In common with the visualisation literature, visual encoding and interaction are discussed separately. However, the principle of *sketching*, which Fry (2007) equates to *scripting*, did not fit easily into either category. Speed characterises sketching: code is written quickly, which means that more ideas can be tested in less time (Dykes 2005). Consequently, many design decisions reflected a compromise between the time it would take to implement the various alternatives and the utility of those alternatives. Nevertheless, sketching is consistent with the iterative approach to design, implementation, and evaluation that is at the heart of the design, analysis, and evaluation model. Visual encoding principles and interaction principles are discussed separately in the following two sections before the software prototypes are themselves presented in Section 5.4.

5.3.1 Visual encoding

Three visual encoding principles unify the software prototypes. The first visual encoding principle is the careful use of the *visual variables* (Bertin 2011) to establish a *graphic hierarchy* (MacEachren 1994*a*), which involves the separation of graphic elements into conspicuous *figure* and inconspicuous *ground* (MacEachren 1994*a*). MacEachren (1994*a*) suggests contrasting either colour value (lightness) or colour hue between two graphic elements so that the smaller of the two elements will be figure, for example. This suggestion was followed in all three software prototypes to distinguish points on each map from the map itself. The second visual encoding principle is that of the consistent use of colour value (lightness) and colour saturation both within and between the software prototypes. Several ColorBrewer schemes were used for the graphic elements within each software prototype (Harrower & Brewer 2003), which made this principle easy to follow. The ColorBrewer website allowed these schemes to be evaluated quickly, which is consistent with the principle of sketching. An advantage of the ColorBrewer schemes is that colours from each scheme are comparable: that is, the first orange from the ColorBrewer Oranges scheme is comparable to the first grey from the ColorBrewer Greys scheme. Whilst it would be possible to devise consistent colours by controlling colour hue, and varying colour value (lightness) and—to a lesser extent—colour saturation, doing so would not be consistent with the principle of sketching. The third visual encoding principle is that of removing axis ticks, labels, and lines from charts, and is influenced by Tufte (2001) and Unwin et al. (2008). Tufte (2001) advises that the ratio between the amount of ink—the number of pixels—used to display the data and the amount of ink used for other purposes (the data–ink ratio) should be as small as reasonably possible. Unwin et al. (2008) argue that because exploratory graphics should emphasise the structure of the data, axis elements such as ticks, labels, and lines are less important in exploratory graphics than in presentation graphics. Consequently, the temporal histograms on the first and second software prototypes omit axis ticks, labels, and lines.

5.3.2 Interaction

Several researchers have argued that interaction has not received as much attention as visual encoding in the visualisation literature (Yi et al. 2007), even though interaction is central to knowledge construction (Pike et al. 2009). This imbalance is unfortunate, not least because the science of interaction has the potential to unify the 'high level' interaction between the user and the problem with the 'low level' interaction between the user and the interface (Pike et al. 2009). Nevertheless, this imbalance is being corrected, most notably by Roth (2012, 2013*a*,*b*), who has investigated *cartographic interaction*, or "the dialogue between a human and a map mediated by a computing device" (Roth 2012, p.377). Consequently, cartographic interaction is briefly reviewed in this section as a necessary precursor to a discussion of interaction principles.

Roth (2012) structures cartographic interaction according to one or more interaction exchanges, where an *interaction exchange* is a question and answer between a user (i.e. a human) and an interface (i.e. a map). In turn, Roth (2012) structures an interaction exchange according to Norman's seven stages of action model (Norman 1998) because this model, unlike others, assigns agency to both the human and the map. Roth (2012) argues that four locations in this model are particularly important. According to Roth (2012), the first stage of action, *forming the goal*, is an important location in this model because *the goal*, or a poorly-defined task, motivates the interaction exchange. Typical goals of visualisation include exploration, confirmation, synthesis, and presentation: the former two contribute to visual thinking, which takes place in the private realm, whilst the latter two contribute to visual communication, which takes place in the public realm (DiBiase 1990, MacEachren 1994*b*). According to Roth (2012), there are three more important locations in this model: at the second stage of action, *forming the intention*; at the third stage of action, *specifying an action*; and between the third and the fourth stages of action, *executing the action* and *perceiving the state of the system*. Roth (2012, 2013*a*,*b*) calls existing taxonomies of interaction at these locations objective-based, operator-based, and operand-based taxonomies of interaction:

- *Objective-based* taxonomies of interaction delimit the space of possible objectives, where an *objective* is a well-defined task.
- *Operator-based* taxonomies of interaction delimit the space of possible operators, where an *operator* is a function that will support the objective.
- *Operand-based* taxonomies delimit the space of possible operands, where an *operand* is an object with which the user interacts.

The value of cartographic interaction, as a concept, is that it provides clarity: by describing interaction in terms of goals, objectives, operators, and operands it is possible to unify the 'high level' interaction between the user and the problem with the 'low level' interaction between the user and the interface (Pike et al. 2009). In this section, two interaction principles are discussed at the operator level. In Section 5.4, the software prototypes are themselves presented, but the emphasis is on the goal level and the objective level.

The first interaction principle is that of the consistent use of panning and zooming between the software prototypes. Panning and zooming are operators that manipulate the viewpoint of the user (Roth 2012, 2013*a*) and have a low degree of interactivity (Crampton 2002). Nevertheless, each software prototype has a map view that can be panned, by moving the mouse cursor with the right mouse button pressed, and zoomed, by moving the mouse cursor with the left mouse button pressed. In each case, zooming does not change the geographic scale of the map; zooming simply enlarges the map using a series of geometric transformations (although points on the map (i.e. operands) are of a constant size—irrespective of zooming—to reduce occlusion, which Lam (2008) regards as an interaction cost). Consequently, zooming—as implemented in the software prototypes—is neither a change in geographic scale nor a change in level-of-detail: the two types of zooming that are identified by Harrower & Sheesley (2005). However, this design decision was consistent with the principle of sketching as it was not apparent that either a change in geographic scale or a change in level-of-detail were required.

The second interaction principle is that of the consistent use of brushing between the software prototypes. Brushing has been applied to linked scatter plots, and to linked scatter plots and maps (Becker & Cleveland 1987, Monmonier 1989). Crampton (2002) refers to the former as *statistical brushing* and to the latter—following Monmonier (1989)—as *geographic brushing*. Gleicher et al. (2011) categorise linked views as a hybrid visual comparison technique that combines juxtaposition and the explicit encoding of relationships. The authors suggest that this technique is especially valuable because juxtaposition and the explicit encoding of relationships are mutually supportive.

Brushing is a multi-step operator—a selection operator and, for example, a resymbolise operator (Roth 2012, 2013*a*)—that has a high degree of interactivity (Crampton 2002). However, the selection operator is far from straightforward. Wills (1996) describes five selection operators (*replace*, *intersect*, *add*, *subtract*, and *toggle*), which are shown in Figure 5.6. The top left panel in Figure 5.6 depicts the initial state of the display: eight circles (four are selected and so are filled grey, four are not selected and so are filled white) in three groups to the front as well as one rectangle (filled grey) to the back. The circles depict the points on the map and the rectangle depicts the brush. However, although the brush (an operand) is depicted, brushing (a multi-step operator) is not complete. The remaining panels in Figure 5.6 depict the state of the display after brushing is complete. Consequently, the brush is not depicted in these panels: it has been removed from the display. In each case, after the points on the map have been selected, they have been resymbolised (filled either grey or white). However, other operators may be invoked.

Wills (1996) argues that *toggle* is the most powerful and the most forgiving selection operator, although *replace* is more common and easier to understand than the alternatives. The five selection operators can be summarised as follows:

- *Replace*. The points on the map that are contained by the brush replace the existing selection.

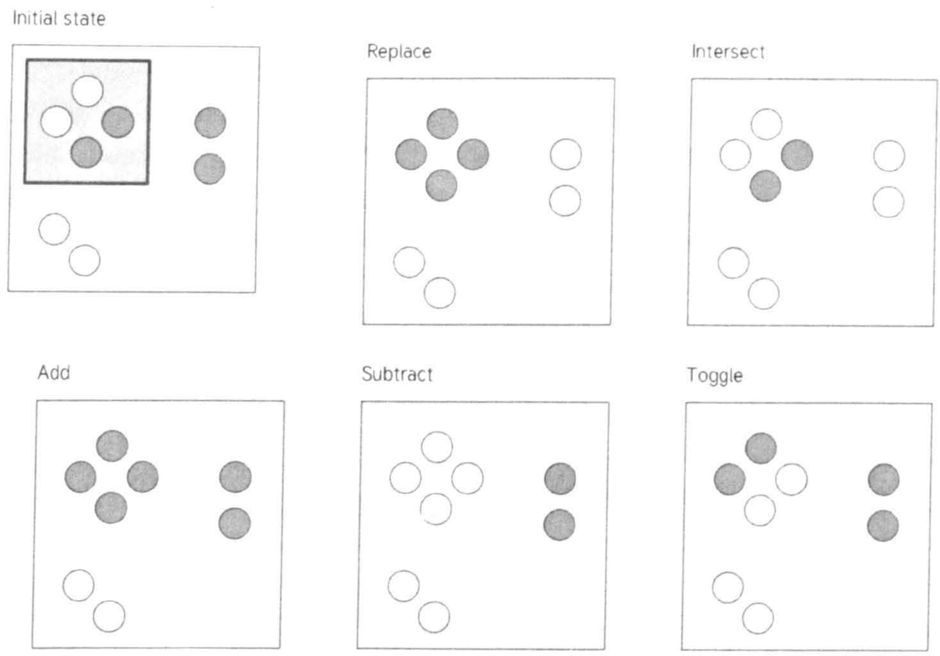

Figure 5.6: Selection operators (Wills 1996). The top left panel depicts the initial state of the display. The remaining panels depict the state of the display after brushing is complete.

- *Intersect.* The points on the map that are contained by the brush intersect the existing selection (logical AND).
- *Add* and *subtract.* The points on the map that are contained by the brush are either added to or subtracted from the existing selection.
- *Toggle.* The points on the map that are contained by the brush are toggled: selected becomes not selected and vice versa (logical NOT).

The *replace* selection operator was implemented in the software prototypes because it is more common and easier to understand than the alternatives (Wills 1996). However, this design decision was consistent with the principle of sketching as it was not apparent that the additional power of the alternatives was required.

In this section, visual encoding principles and interaction principles were discussed. Visual encoding principles included the careful use of the visual variables to establish a graphic hierarchy; the consistent use of colour value (lightness) and colour saturation both within and between the software prototypes; and removing axis ticks, labels, and lines from charts. Interaction principles included the consistent use of panning and zooming and the consistent use of brushing between the software prototypes. Interaction principles were discussed at the operator level. In the following section, the software prototypes are themselves presented, but the emphasis is on the goal level and the objective level.

5.4 Geovisualisation software prototypes

The design, analysis, and evaluation model was used to develop—and is used to structure the discussion of—the three geovisualisation software prototypes. These software prototypes were developed to support the interactive visual exploration of reports from the Haiti crisis map. Exploration is a typical goal of visualisation (DiBiase 1990, MacEachren 1994*b*), where the *goal*, or a poorly-defined task, motivates the *interaction exchange*, or a question and answer between a user (i.e. a human) and an interface (i.e. a map) (Roth 2012). Nevertheless, the goal of each software prototype was given more definition—without it becoming an *objective*, or a well-defined task (Roth 2012)—to meet the initial stage of the design, analysis, and evaluation model: that is, to formulate working hypotheses. Consequently, the goals of the software prototypes were to explore the spatial and temporal distributions of reports; to explore the composition and organisation of location descriptions; and to explore the ambiguity associated with the composition of location descriptions. Each software prototype is now discussed in turn.

5.4.1 Software prototype 1: Exploring the spatial and temporal distributions of reports

The first software prototype was developed early in the current research. Consequently, it addressed two general research questions that related to the spatial and temporal distributions of reports. These research questions were translated into a scenario, and features were then identified and justified.

Research questions

The research questions that motivated the development of the first software prototype are as follows:

- How does the temporal distribution of reports change over space?
- How does the spatial distribution of reports change over time?

Figure 5.7: Overview of the first software prototype, showing the map, the standard ellipse, and the temporal histogram (no spatial or temporal selections).

Scenario

The scenario that motivated the development of the first software prototype is as follows:

John wants to explore how the distribution of reports from the Haiti crisis map changes over space and time. John launches the software prototype and is presented with a map and a temporal histogram. Initially, the map and the temporal histogram show all reports. On the map, each report is represented as a dot positioned according to its longitude and latitude attributes. On the temporal histogram, time is on the horizontal axis and number of reports is on the vertical axis.

By looking at the temporal histogram, John sees that more reports occur earlier rather than later in the crisis. He is interested in where the earlier reports are located, so he brushes several bars on the temporal histogram. The brush is depicted by a rectangle on the temporal histogram whilst brushing is underway. Brushing subsets the reports: it highlights the bars on the temporal histogram and the dots on the map that correspond to the subset. John sees a cluster of reports around Port-au-Prince, so zooms and pans the map to investigate further.

John wants to explore the distribution of this cluster of reports in space and time so he brushes the map. Starting this action returns all map dots and temporal histogram bars to their unhighlighted state. Finishing this action re-subsets the reports: the bars on the temporal histogram and the dots on the map transition to reflect the new subset. The brush is depicted by a rectangle on the map whilst brushing is underway. The brush remains on the map when brushing has completed.

John activates the standard ellipse, which is then drawn inside the brush. John brushes a bar on the temporal histogram (further subsetting the subset of reports that he created when he brushed the map) and the standard ellipse is then redrawn inside the brush. John moves the temporal selection forwards and backwards through time: upon each move, the standard ellipse is redrawn inside the brush and the bars representing the temporal selection are rehighlighted. Having explored the cluster of reports around Port-au-Prince, John resets the map and temporal histogram to their initial state.

Features

The features that were identified in the scenario that motivated the development of the first software prototype are justified as follows:

- *The spatial and temporal views.* The spatial and temporal views reflect an established conceptual framework: geographic information has spatial and temporal components (Andrienko & Andrienko 2006, Longley et al. 2011).
- *Initially, all views depict all reports.* Shneiderman (1996) argues that the overview should precede the zoom, filter, and detail views.
- *The spatial view is a dot map and the temporal view is a temporal histogram.* A dot map is an established visual encoding technique for exploring point patterns (Bailey & Gatrell 1995). However, occlusion can hinder perception (Lam 2008) and it is easy to see point patterns where none exist (Bailey & Gatrell 1995, Wickham et al. 2010).
- *Brushing the temporal histogram highlights the corresponding temporal histogram bars and map dots.* Brushing the temporal histogram applies the select operator. It then applies the resymbolise operator to the temporal histogram and to the map. However, the user may not realise that brushing the temporal histogram does not remove information from the temporal histogram or the map.
- *Brushing the map highlights the corresponding map dots, and transitions the corresponding temporal histogram bars.* Brushing the map applies the select operator. It then applies the resymbolise operator to the map, and the filter operator to the temporal histogram. However, the user may not realise that brushing the map removes information from the temporal histogram but does not remove information from the map.
- *Transitioning between states.* Transitioning between states can "significantly improve graphical perception" (Heer & Robertson 2007, p.1240). However, it is possible to track only eight, slow-moving objects at once (Alvarez & Franconeri 2007).
- *The standard ellipse.* A standard ellipse summarises a point pattern[3] (Kitchin & Tate 2000, O'Sullivan & Unwin 2003). Redrawing the standard ellipse as the temporal selection is moved forwards and backwards through time allows

[3] A *standard ellipse* is located at the mean centre of a set of point locations. Each axis is two standard deviations in length: the major axis extends in the direction of maximum dispersion; the minor axis extends in the direction of minimum dispersion (Kitchin & Tate 2000, O'Sullivan & Unwin 2003).

changes to the point pattern to be compared over time. However, the user may not realise (or may forget) that the standard ellipse is the result of a select operator and a filter operator (i.e. brushing the map).

Design decisions

Two colours from the ColorBrewer Blues scheme are used for the highlighted and non-highlighted points on the map, and the highlighted and non-highlighted bars on the temporal histogram. However, the points on the map are 50% transparent: hence, colour density communicates point density. Two colours from the ColorBrewer Greys scheme are used for the country polygon on the map. Two colours from the ColorBrewer Oranges scheme are used for the standard ellipse: one darker orange for the ellipse and one lighter orange for the axes. The ColorBrewer Oranges scheme is used for the standard ellipse because it contrasts with the ColorBrewer Greys scheme that is used for the country polygon: the relationship within the oranges, and between the oranges and the greys, contributes to a clear graphic hierarchy. The ColorBrewer Blues scheme is used for the points on the map and the bars on the temporal histogram for the same reason: the ground should be the country polygon and, slightly higher in the graphic hierarchy, should be the non-highlighted points and bars. The figure should be the highlighted points and bars and, highest in the graphic hierarchy, should be the standard ellipse.

Brushing the temporal histogram applies the select operator. It then applies the resymbolise operator to the temporal histogram and the map. In other words, brushing the temporal histogram highlights the corresponding bars on the temporal histogram and points on the map: consequently, reports that *are not* in the temporal selection *are* depicted on the map and the temporal histogram. Brushing the map applies the select operator. It then applies the resymbolise operator to the map, and the filter operator to the temporal histogram. In other words, brushing the map highlights the corresponding points on the map and filters the corresponding bars on the temporal histogram: consequently, reports that *are not* in the spatial selection *are not* depicted on the map or on the temporal histogram. In both cases, the select operator is a replace operator. Furthermore, whilst it is possible to brush the temporal histogram and then the map, or the map and then the temporal histogram, neither view can be brushed twice. This design decision was made because although multiple combinations of selection and filter operators can be very powerful, they can also be very confusing, especially when an upstream filter modifies a downstream selection (Weaver 2010).

A temporal selection can be moved right or left using the corresponding arrow keys, which makes it possible to step forwards or backwards through time. For example, Figure 5.8a depicts all reports from the Haiti crisis map. The second bar on the temporal histogram is then brushed, which highlights this bar and the corresponding points on the map, as shown in Figure 5.8b. The right arrow key is then pressed twice, which moves the temporal selection two days forward, as shown in Figure 5.8c. Reports from day 7, day 12, and day 13 of the Haiti crisis map are shown in Figures 5.8d, 5.8e, and 5.8f.

Two sources of base map data were considered for the first and subsequent software prototypes: Global Administrative Areas (GADM 2013) and Natural Earth (Natural Earth 2013). Natural Earth base map data are generalised to 1:10m, 1:50m, and 1:110m geographic scales. Furthermore, as well as countries, data are available for populated places, urban areas, and transport networks. Consequently, Natural Earth base map data are used at the 1:10m geographic scale.

Slocum et al. (2009) suggest that the stereographic projection (*conformal*, or angle-preserving) and the Lambert azimuthal equivalent projection (*equivalent*, or area-preserving) would be suitable for Haiti. Whilst an equivalent projection should be used for a dot map (Slocum et al. 2009), a conformal projection should be used for spatial analysis, such as a standard ellipse (Kitchin & Tate 2000, O'Sullivan & Unwin 2003). Nevertheless, Kitchin & Tate (2000) argue that distortion is minimal over smaller areas. Consequently, the Universal Transverse Mercator coordinate system is used, as neither the stereographic projection nor the Lambert azimuthal equivalent projection were readily available.

Findings

As shown in Figure 5.3, the design, analysis, and evaluation model includes a *document findings* stage, as well as a *write a research report* stage. Whilst the former is an analysis activity, the latter is an evaluation activity. In this section, and in subsequent findings sections, analysis and evaluation are discussed together. In each case, however, a research question is discussed then reflected upon (i.e. analysis then evaluation); only then is the next research question discussed then reflected upon (i.e. the next analysis then the next evaluation).

The first research question that motivated the development of this software prototype was to investigate how the temporal distribution of reports changes over space. The temporal histogram shown in Figure 5.9a indicates that the temporal distribution of all reports is positively skewed and peaks earlier in the crisis. Brushing the map, which filters the temporal histogram, indicates that the temporal distribution of reports remains the same over space: that is, although the heights of the bars change, the shape of the temporal distribution remains the same. For example, compare Figure 5.9a (no spatial selection) to Figure 5.9b (spatial selection).

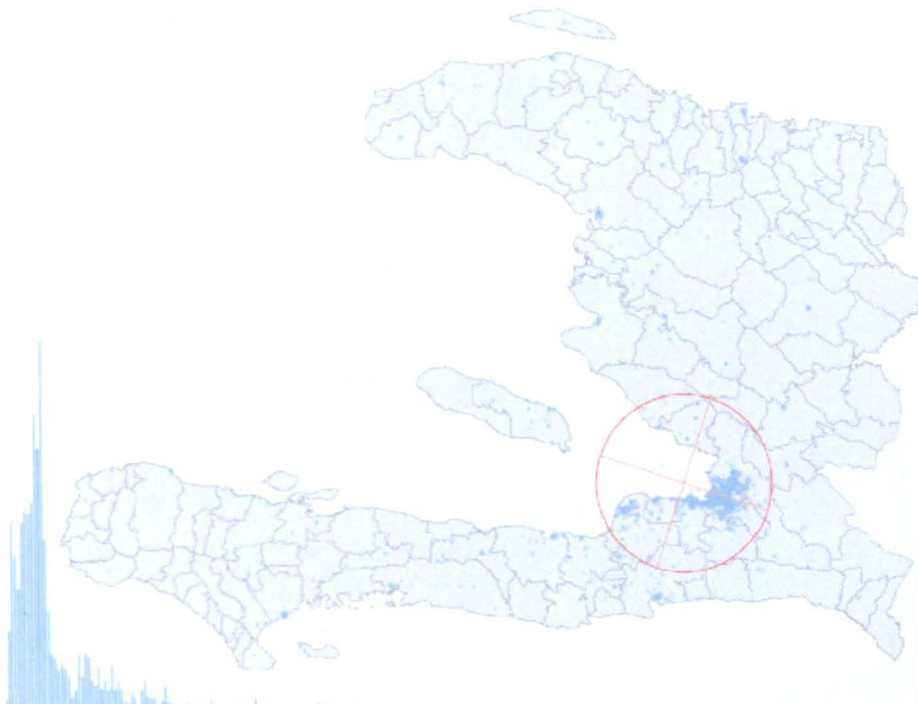

(a) All reports from the Haiti crisis map (no spatial or temporal selections).

(b) Temporal selection of reports from day 2 of the Haiti crisis map (no spatial selection).

Figure 5.8: Stepping forwards through time in the first software prototype.

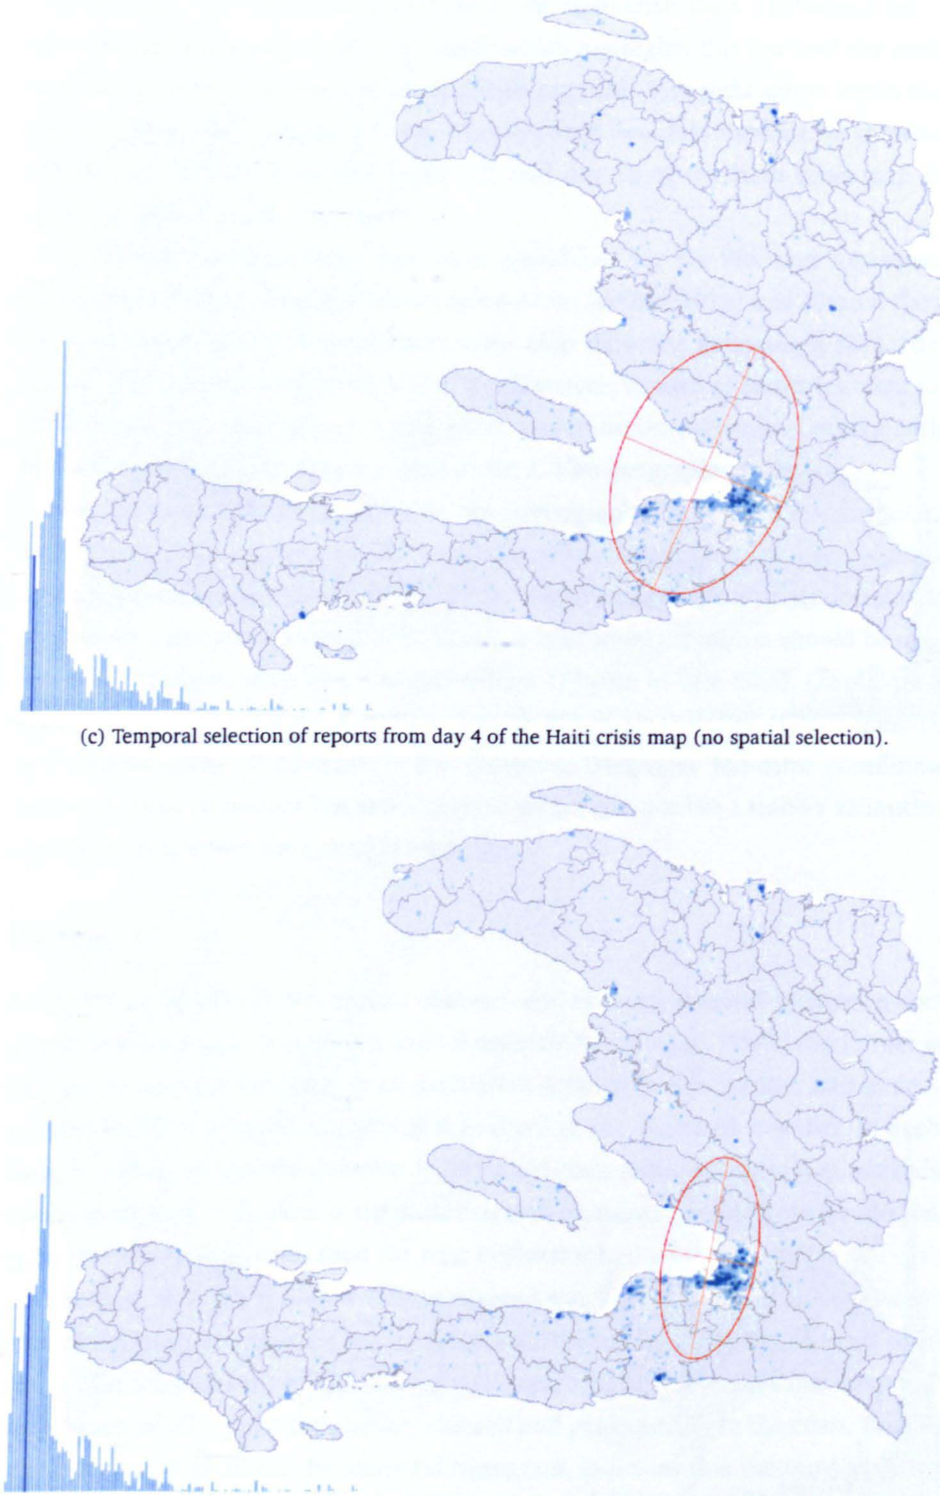

(c) Temporal selection of reports from day 4 of the Haiti crisis map (no spatial selection).

(d) Temporal selection of reports from day 7 of the Haiti crisis map (no spatial selection).

Figure 5.8: Stepping forwards through time in the first software prototype.

(e) Temporal selection of reports from day 12 of the Haiti crisis map (no spatial selection).

(f) Temporal selection of reports from day 13 of the Haiti crisis map (no spatial selection).

Figure 5.8: Stepping forwards through time in the first software prototype.

On reflection (i.e. to evaluate), the design of this software prototype did not fully support the investigation of how the temporal distribution of reports changes over space because brushing the map filters—and so removes information from—the temporal histogram. Consequently, to compare two temporal distributions, such as 'all reports' and 'the spatial subset of reports' as illustrated in Figures 5.9a and 5.9b, the user must remember the first temporal distribution when investigating the second temporal distribution. A better approach would have been to have global–local and local variants of the temporal histogram. These variants are discussed in Section 5.4.2, which can be thought of as a second iteration of the design of this software prototype, and of the design, analysis, and evaluation model.

The second research question that motivated the development of this software prototype was to investigate how the spatial distribution of reports changes over time. The standard ellipse shown in Figure 5.8a indicates that the spatial distribution of all reports is approximately even (i.e. the major and minor axes are of equal length) and is centred on the Port-au-Prince area. The standard ellipses shown in Figures 5.8b to 5.8f indicate that the spatial distribution of reports exhibits considerable variation from day-to-day, although remains centred on the Port-au-Prince area. Nevertheless, it is hard to make meaningful comparisons from day-to-day because the temporal distribution of all reports is positively skewed. In other words, these comparisons are often based on a very small number of reports.

On reflection (i.e. to evaluate), the design of this software prototype did not fully support the investigation of how the spatial distribution of reports changes over time because the standard ellipse is redrawn as the temporal selection is moved forwards and backwards through time. Consequently, to compare multiple spatial distributions, such as those illustrated in Figures 5.8b to 5.8f, the user must remember each spatial distribution, or the standard ellipse that summarises each spatial distribution. A better approach would have been to fade out the standard ellipse that summarises each spatial distribution as the temporal selection is moved forwards and backwards through time (Harrower 2002). However, because the second software prototype, which is discussed in Section 5.4.2, did not implement a standard ellipse, fading out the standard ellipse was not considered further.

The first software prototype addressed two general research questions that related to the spatial and temporal distributions of reports. It was also the product of the first iteration of the design, analysis, and evaluation model. However, because the first software prototype was developed early in the current research, it did not explore the relationships between the spatial and temporal distributions of reports, and their corresponding location descriptions. Given that the current research proceeded to investigate the types of location descriptions in Chapter 3 and

(a) The temporal histogram without a corresponding spatial selection (map zoomed to the Port-au-Prince area). The temporal distribution of all reports is positively skewed and peaks earlier in the crisis.

(b) The temporal histogram with a corresponding spatial selection (map zoomed to the Port-au-Prince area). The temporal distribution of the spatial subset of reports is still positively skewed and still peaks earlier in the crisis.

Figure 5.9: The temporal histogram in the first software prototype, with and without a corresponding spatial selection.

the composition and organisation of location descriptions in Chapter 4, on reflection (i.e. to evaluate), a better approach would have been to explore these relationships further. For example, in this software prototype, two colours from the ColorBrewer Blues scheme are used for the highlighted and non-highlighted points on the map. Instead, several colours from this scheme could be used either discretely or continuously to depict the number of characters in the location descriptions that correspond to the points on the map; different ColorBrewer schemes—such as Blues and Oranges—could then be used for the highlighted and non-highlighted points on the map. Here, number of characters is a proxy for complexity: longer location descriptions are assumed to be more complex. It would be interesting to see whether longer, more complex location descriptions are clustered in space and time. Furthermore, it would be useful to filter reports by the number of characters in the location descriptions and display the location descriptions themselves, to test the assumption that longer location descriptions are more complex. A similar approach is discussed in Section 5.4.3, which can be thought of as a third iteration of the design of this software prototype, and of the design, analysis, and evaluation model.

Having completed one iteration of the design, analysis, and evaluation model—and having discussed analysis and evaluation activities—it is important to consider why the first software prototype did not fully support the research questions. As noted, the first software prototype addressed two general research questions. Furthermore, the scenario that motivated its development was clearly written in terms of what Munzner (2009) calls *operations*, or tasks specified in the language of visualisation, rather than *problems*, or tasks specified in the language of the domain. Consequently, a possible reason why the first software prototype did not fully support the research questions is that it became detached from the underlying problem, which was not specified in enough detail; in other words, the first software prototype was not adequately grounded (Isenberg et al. 2008). As is discussed in Sections 5.4.2 and 5.4.3, a potential solution was identified: namely, to omit the *write a scenario* stage of the model and so to undertake the *identify and justify features* stage following the *formulate working hypotheses* stage. The implications of this potential solution are discussed in Section 5.5, where several sources of ideas for challenging the designer-user are suggested.

5.4.2 Software prototype 2: Exploring the composition and organisation of location descriptions

The aim of the research reported in Chapter 4 was to investigate the types of geographic objects and spatial relationships that were contained in location descriptions, and to investigate how these types were organised with respect to each other. It was found that location descriptions tended not to contain vague places; tended to contain more, relatively coarse granularity types of geographic objects (i.e. more than relatively fine types of geographic objects); and tended to be paired in a finer-to-coarser granularity pattern. Nevertheless, these findings were extremely tentative. Consequently, the second software prototype was developed to extend the research reported in Chapter 4.

Research questions

The research questions that motivated the development of the second software prototype are as follows:

- To what degree are different patterns of geographic objects evident in location descriptions?
- To what degree are the spatial and temporal distributions of location descriptions that contain different geographic objects different to each other?

Scenario

Following the development of the first software prototype, the *write a scenario* stage of the design, analysis, and evaluation model was not deemed worthwhile. This issue is discussed further in Section 5.5. Consequently, having completed the *formulate working hypotheses* stage, development proceeded to the *identify and justify features* stage.

Features

This software prototype can be thought of as a second iteration of the design, analysis, and evaluation model. Consequently, many of the features of this software prototype were described in Section 5.4.2. However, the new features of the second software prototype are justified as follows:

- *The thematic view.* The thematic view depicts each location description as a row and each geographic object contained in each location description as a cell. Cells in the thematic view are coloured by their GeoNames feature class.

Design decisions

The map and the thematic view depict location descriptions that contain the selected GeoNames feature class. This class can be changed by pressing the up and

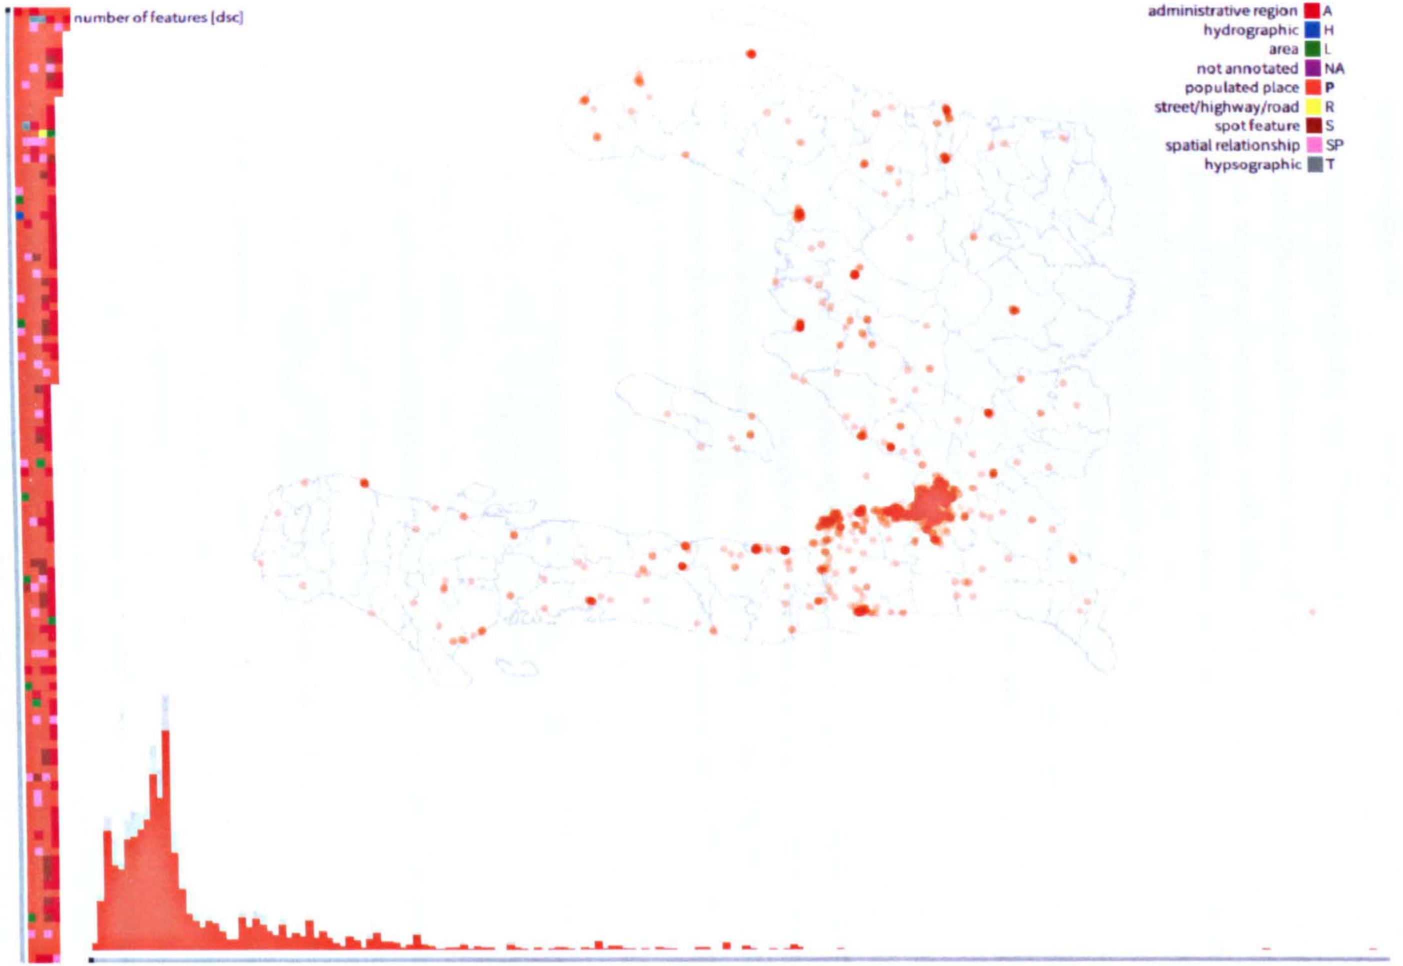

Figure 5.10: Overview of the second software prototype, showing the map, the temporal histogram, and the thematic view. The thematic view depicts each location description as a row and each geographic object contained in each location description as a cell. Rows in the thematic view (i.e. location descriptions) can be ordered by number of geographic objects, by time, or by GeoNames feature class (either ascending or descending).

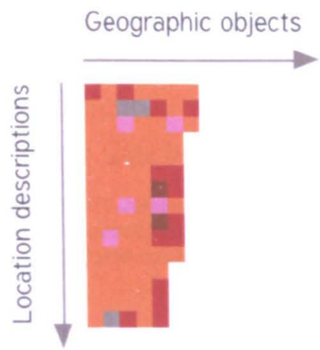

Figure 5.11: A section of the thematic view in the second software prototype. The thematic view depicts each location description as a row and each geographic object contained in each location description as a cell. Cells in the thematic view are coloured by their GeoNames feature class.

down arrow keys, and has a bold label on the legend. Pressing the l (label) key switches between abbreviated and full labels on the legend. Nevertheless, the number of rows (i.e. the number of location descriptions) in the thematic view often exceeds the size of the display. Consequently, the thematic view can be scrolled. Like panning and zooming, scrolling can be thought of as an operator that manipulates the viewpoint of the user (Roth 2012, 2013*a*) and has a low degree of interactivity (Crampton 2002).

Rows in the thematic view (i.e. location descriptions) can be ordered by number of geographic objects, by time, or by GeoNames feature class (either ascending or descending). The row order, which is displayed to the right of the thematic view, can be changed by pressing a number key (from 1 to 6). The row order can be thought of as a sequence operator (Roth 2012, 2013*a*) and has a moderate degree of interactivity (Crampton 2002).

Colours from the ColorBrewer Set1 scheme are used for the GeoNames feature classes and, consequently, for points on the map and cells in the thematic view. This design decision was made for two reasons. First, a qualitative scheme, such as the ColorBrewer Set1 scheme, does not imply an order between its elements (Harrower & Brewer 2003). Second, colours from the ColorBrewer Set1 scheme are saturated, which makes them more visible when used for small areas, such as points on the map or cells in the thematic view. Less saturated colours, such as those from the ColorBrewer Set3 scheme, would be less visible when used for small areas. To avoid interference (simultaneous contrast), light grey and white are used for the border and the fill of the country polygon. However, it can be hard to distinguish the country polygon, which is white, from the background, which is also white, especially after zooming. Lam (2008) regards the problem of associating a

(a) The global–local variant of the temporal histogram. The global layer, which depicts the temporal distribution of all location descriptions, is in the background and has a light grey fill. The local layer, which depicts the temporal distribution of only location descriptions that contain the selected GeoNames feature class, is in the foreground and has a red fill.

(b) The local variant of the temporal histogram. The local layer, which depicts the temporal distribution of only location descriptions that contain the selected GeoNames feature class, is in the foreground and has a red fill.

Figure 5.12: The global–local and local variants of the temporal histogram in the second software prototype.

global representation (i.e. zoomed-out) with a local representation (i.e. zoomed-in) as an interaction cost. A better design decision would have been to use dark grey for the border and light grey for the fill of the country polygon, and to adjust the colours from the ColorBrewer Set1 scheme, which are used for the GeoNames feature classes, according to whether they are used on a light grey background (i.e. used on the map) or on a white background (i.e. used on the legend). Brewer (1997) offers guidance on the nature of this adjustment.

The temporal histogram has global–local and local variants. These variants can be toggled by pressing the s (scale) key. The global–local variant depicts the temporal distribution of all location descriptions in the background, which is the global layer, and the temporal distribution of only location descriptions that contain the selected GeoNames feature class in the foreground, which is the local layer. The global–local variant is shown in Figure 5.12a, where the the global layer has a light grey fill and the local layer has a red fill. The local variant depicts the temporal distribution of only location descriptions that contain the selected GeoNames feature class in the foreground, which is the local layer. The local variant is shown in Figure 5.12b, where the local layer has a red fill. Gleicher et al. (2011) regard the use of global and local layers as *superposition*, or showing objects together in space and time.

For consistency, the same light grey that is used for the border of the country polygon is used for the fill of the bars in the global layer of the temporal histogram.

However, in the graphic hierarchy the border of the country polygon is below the bars because the bars occupy a larger amount of screen space than the border of the country polygon. The colour from the ColorBrewer Set1 scheme that is used for the selected GeoNames feature class is used for the fill of the bars of the local layer of the temporal histogram.

Findings

The first research question that motivated the development of this software prototype was to investigate the degree to which different patterns of geographic objects are evident in location descriptions. Three such patterns were identified by using this software prototype. These patterns lend support to the finding, which was reported in Chapter 4, that geographic objects tend to be paired in a finer-coarser-granularity pattern:

- For almost all location descriptions that contain administrative regions, the administrative region is the last geographic object in the location description and is preceded by one or more populated places. The thematic view in Figure 5.13 illustrates this finding.
- For almost all location descriptions that contain spot features, the spot feature is preceded by a populated place. In most cases, it is also followed by an administrative region. The thematic view in Figure 5.14 illustrates this finding.
- For many location descriptions that contain spatial relationships, the spatial relationship is both preceded and followed by a populated place (i.e. the spatial relationship is between two populated places). However, in a similar number of location descriptions, the spatial relationship is the first geographic object in the location description and is followed by one or more populated places. The thematic view in Figure 5.15 illustrates this finding.

On reflection, the design of this software prototype could have better supported the investigation of the degree to which different patterns of geographic objects are evident in location descriptions. To illustrate how, consider an example of two location descriptions, each of which contains three geographic objects: the first location description's geographic objects are classed as 'PPA'; the second location description's geographic objects are classed 'PPP'. Arguably, these location descriptions are more similar to each other than to a third location description, whose geographic objects are classed 'SRA', for example. Nevertheless, because the map and the thematic view only depict location descriptions that contain the selected GeoNames

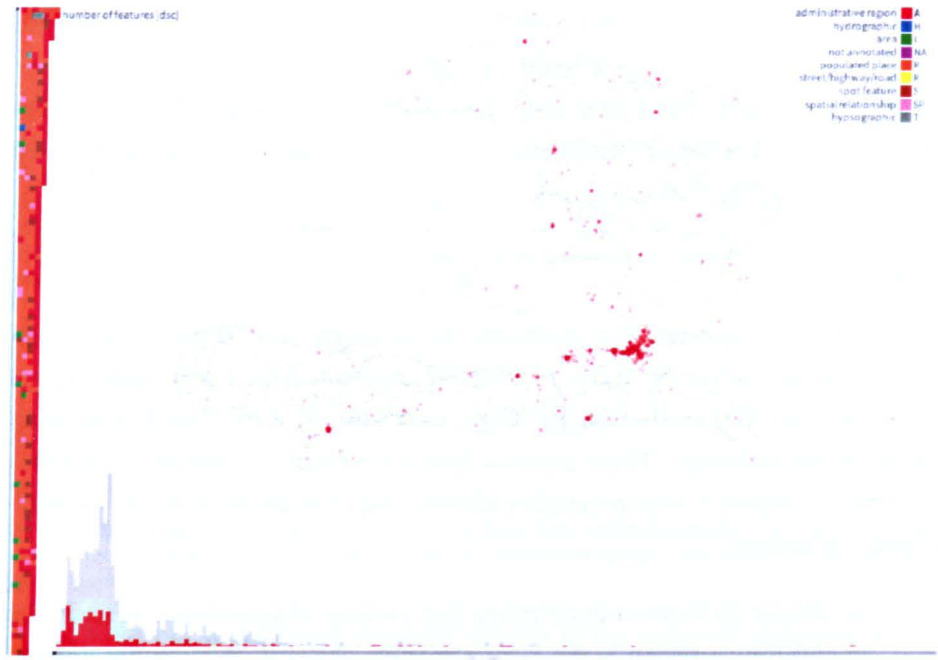

Figure 5.13: Administrative regions. The thematic view indicates that for almost all location descriptions that contain administrative regions, the administrative region is the last geographic object in the location description.

Figure 5.14: Spot features. The thematic view indicates that for almost all location descriptions that contain spot features, the spot feature is preceded by a populated place.

Figure 5.15: Spatial relationships. The thematic view indicates that for many location descriptions that contain spatial relationships, the spatial relationship is both preceded and followed by a populated place.

feature class, if the selected GeoNames feature class is 'A', only the first and the third location description are depicted in the map and the thematic view. A better approach would have been to cluster location descriptions according to a similarity measure and then to depict these clusters. Furthermore, if the finer-granularity GeoNames feature codes had been used instead of the coarser-granularity GeoNames feature classes, a better understanding of the patterns of geographic objects that are evident in location descriptions could have been achieved.

The second research question that motivated the development of the this software prototype was to investigate the degree to which the spatial and temporal distributions of the location descriptions that contain different geographic objects are different to each other. By using this software prototype it was found that location descriptions that contain spot features as well as those that contain spatial relationships tend to cluster in the Port-au-Prince area, although the former are more clustered than the latter. This finding is illustrated by comparing Figures 5.14 and 5.15. No difference in temporal distributions was found, however.

On reflection, the design of this software prototype did not fully support the investigation of the degree to which the spatial and temporal distributions of the location descriptions that contain different geographic objects are different to each other. In the case of the spatial distributions of the location descriptions, multiple standard ellipses—one for each GeoNames feature class—could be implemented. These standard ellipses could be displayed simultaneously and could be combined with the ability to move forwards and backwards through time, as with the first software prototype. Indeed, had location descriptions been clustered according to a similarity measure, it would have been interesting to summarise the spatial distributions of the location descriptions using a map and standard ellipse linked to a dendrogram, which is a common tree-based representation of a clustering process. In the case of the temporal distributions of the location descriptions, a better approach would have been to use each bar in the temporal histogram to depict the number of location descriptions that contained the selected GeoNames feature class as a proportion of the number of location descriptions that were reported at that point in time. This approach would have prevented the positively skewed temporal distribution of all reports influencing the interpretation of the temporal histogram.

5.4.3 Software prototype 3: Exploring the ambiguity associated with the composition of location descriptions

The aim of the research reported in Chapter 4 was to investigate the types of geographic objects and spatial relationships that were contained in location descriptions, and to investigate how these types were organised with respect to each other. Geographic objects were identified and resolved using the GeoNames gazetteer. However, there were often several candidate entries in the GeoNames gazetteer for each geographic object. Consequently, there was often ambiguity associated with the resolution process, where *ambiguity* can be defined as the doubt associated with the classification of a phenomenon (Fisher 1999).

It was argued in Chapter 4 that there is a role for geovisualisation for investigating how well geographic objects were classified; that is, for evaluating the geoparsing system that was described in Chapter 4, Section 4.1.2. This role is outlined in this section, where a geovisualisation software prototype for exploring the ambiguity associated with the composition of location descriptions is discussed. The basis of this exploration is a comparison: a comparison of the locations of geographic objects contained in location descriptions on the one hand, with the crowdsourced point locations of those location descriptions on the other.

Research questions

The research questions that motivated the development of the third software prototype are as follows:

- For each location description, what can be said about the ambiguity associated with the resolution process?

- For each location description, what is the relationship between the ambiguity associated with the resolution process and the spatial distribution (i.e. distance and direction) of entries in the GeoNames gazetteer that were selected by the resolution process?

Scenario

Following the development of the first software prototype, the *write a scenario* stage of the design, analysis, and evaluation model was not deemed worthwhile. This issue is discussed further in Section 5.5. Consequently, having completed the *formulate working hypotheses* stage, development proceeded to the *identify and justify features* stage.

Features

This software prototype can be thought of as a third iteration of the design, analysis, and evaluation model. Consequently, many of the features of this software prototype were described in Sections 5.4.2 and 5.4.3. However, the new features of the third software prototype are justified as follows:

- *The thematic view and the map.* The thematic view and the map are an overview and a detail view of the identification and resolution processes reported in Chapter 4, Section 4.1.2. The use of overview and detail view reflects an established conceptual framework (Shneiderman 1996).

- *The thematic view.* The thematic view depicts the location descriptions themselves. The background of each geographic object contained in each location description is coloured according to the number of candidate entries in the GeoNames gazetteer. Consequently, the the background depicts the ambiguity associated with the geographic object. The location descriptions can be ordered by number of geographic objects, by length (number of characters), and by an uncertainty metric. For each location description, the uncertainty metric is the ratio of the number of geographic objects to the number of candidate entries in the GeoNames gazetteer. Consequently, the uncertainty metric is an estimate of the ambiguity associated with the location description. When a location description contains one geographic object that has one candidate entry in the GeoNames gazetteer, the uncertainty metric will be 1.0. The location description 'toussaint louverture international airport, haiti' is an example of this case: 'toussaint louverture international airport' is the geographic object and it has one candidate entry in the GeoNames gazetteer. When a location description contains one geographic object that has 75 candidate entries in the GeoNames gazetteer, the uncertainty metric will be approximately 0.01. The location description 'bellevue, haiti' is an example of this case: 'bellevue' is the geographic object and it has 75 candidate entries in the GeoNames gazetteer.

- *The map*. The map depicts information for only the selected location description: the crowdsourced point location, the candidate entries in the GeoNames gazetteer, and the the entries in the GeoNames gazetteer that were selected by the resolution process.

Design decisions

The background of each geographic object contained in each location description in the thematic view is coloured according to a continuous local scale. In other words, the minimum and maximum number of candidate entries in the GeoNames gazetteer are computed for each location description, rather than for all location descriptions, and the lightest and darkest colours in the scheme are assigned to these values. Colours for intermediate values are interpolated. Two single hue, sequential ColorBrewer schemes are used for this purpose: ColorBrewer Oranges for the selected location description and ColorBrewer Blues for the non-selected location descriptions. For consistency, the colours in each scheme are comparable (i.e. the first colour in the Oranges scheme has the same index as the first colour in the Blues scheme, as do the second, third, and fourth colours). Figure 5.17 shows a single location description from the thematic view, annotated to show the geographic objects with the most candidate entries in the GeoNames gazetteer.

Location descriptions in the thematic view can be ordered by number of geographic objects, by length (number of characters), and by an uncertainty metric. Nevertheless, the number of location descriptions in the thematic view often exceeds the size of the display. Consequently, the thematic view can be scrolled. Like panning and zooming, scrolling can be thought of as an operator that manipulates the viewpoint of the user (Roth 2012, 2013*a*) and has a low degree of interactivity (Crampton 2002).

A mouse click selects a location description in the thematic view. Alternatively, the up or down arrow keys select the location description above or below the selected location description. The selected location description is always depicted at the top of the display to reduce occlusion, which Lam (2008) regards as an in-

Figure 5.16: Overview of the third software prototype, showing the thematic view and the map. The selected location description is depicted at the top of the display, using the ColorBrewer Oranges scheme.

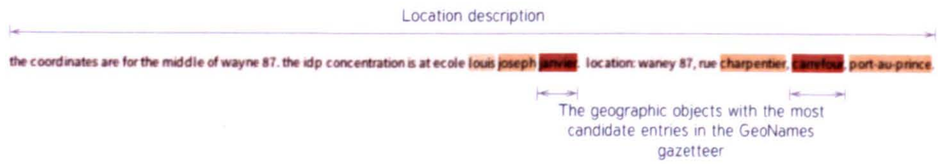

Figure 5.17: A single location description from the thematic view, annotated to show the geographic objects with the most candidate entries in the GeoNames gazetteer.

teraction cost. However, a mouse click causes the selected location description to 'jump' to the top of the display, which, again, Lam (2008) regards as an interaction cost. Nevertheless, the advantage of depicting the selected location at the top of the display—consistency—was deemed to outweigh the limitation of the 'jump'.

The map depicts information for only the selected location description. A larger circle depicts the crowdsourced point location. Smaller circles depict the candidate entries in the GeoNames gazetteer. A thin line emphasises the distance from the crowdsourced point location to the entry in the GeoNames gazetteer that was selected by the resolution process. Gleicher et al. (2011) regard this line as an explicit encoding of relationships. Figure 5.18 shows the map, annotated to show candidate entries in the GeoNames gazetteer (selected and not selected by the resolution process) and the crowdsourced point location.

Moving the cursor over a geographic object contained in the currently selected location description in the thematic view highlights the corresponding candidate entries in the GeoNames gazetteer on the map. Whilst the second and third software prototypes used colours from the same ColorBrewer scheme for selected and non-selected elements, this software prototype uses empty circles (stroke, no fill) for selected elements and full circles (no stroke, fill) for non-selected elements. This design decision was made because it was found that for identifying a small number of circles on the map, the contrast between empty and full was easier to discern than the contrast between darker and lighter colours from the same ColorBrewer scheme.

Findings

The first research question that motivated the development of this software prototype was to explore the ambiguity associated with the resolution process. Ordering location descriptions by the uncertainty metric and selecting each location description in turn confirms how the least ambiguous location descriptions often contain one or two geographic objects, each of which has one candidate entry in the GeoN-

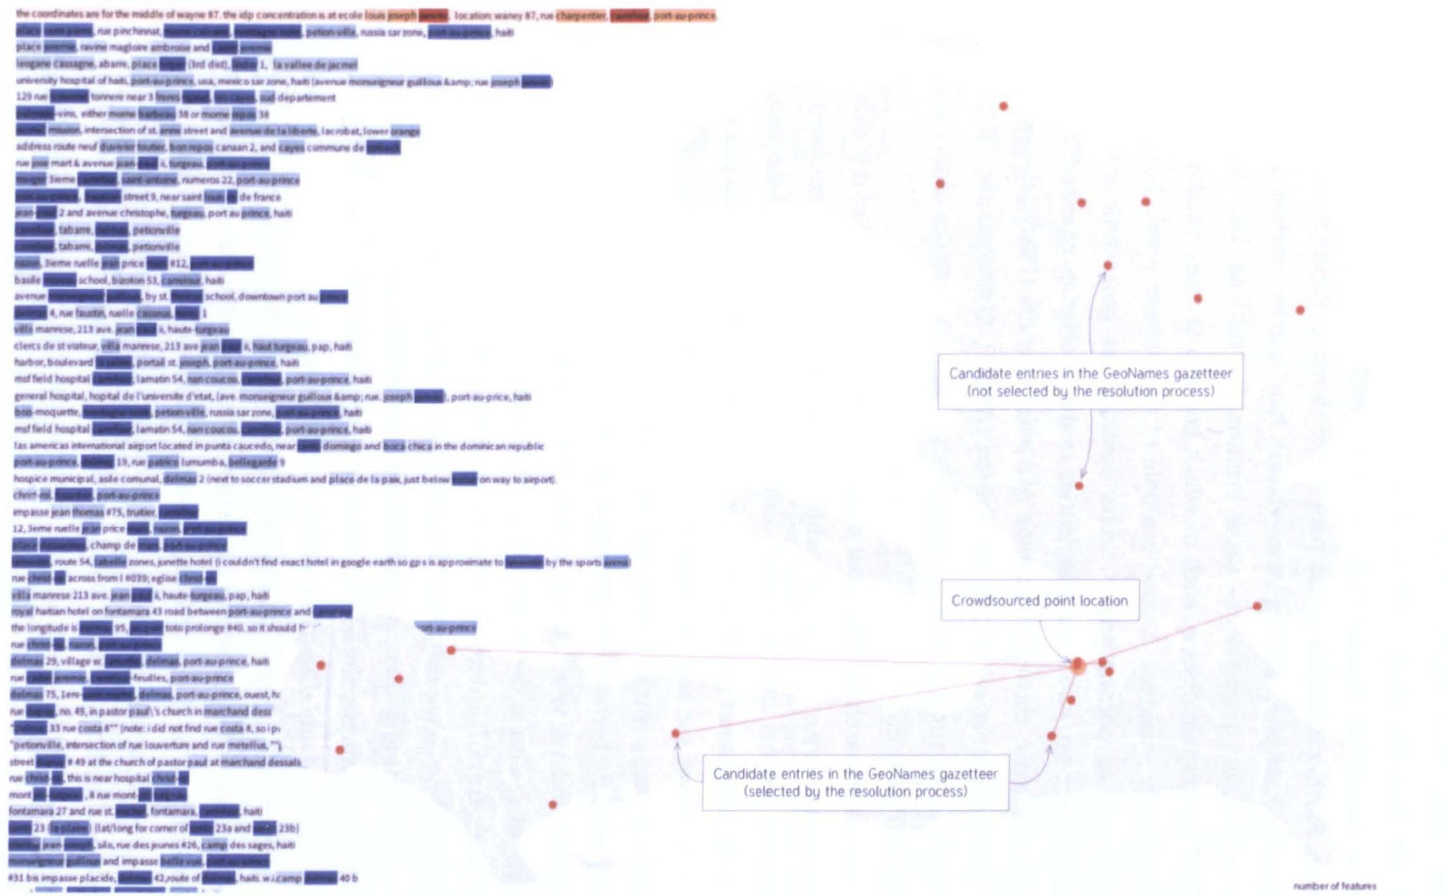

Figure 5.18: The map, annotated to show candidate entries in the GeoNames gazetteer (selected and not selected by the resolution process) and the crowdsourced point location.

ames gazetteer. A good example of this pattern would be a location description such as 'morne basile, haiti', which contains one geographic object ('morne basile').[4] There is one candidate entry in the GeoNames gazetteer for 'morne basile', which is classed as *cliffs* (CLF). Consequently, the uncertainty metric is 1.0 for this location description. Further examples of this pattern are shown in Figure 5.19.

A location description such as 'morne basile, haiti' can be contrasted with more ambiguous location descriptions—those location descriptions that often contain one or two geographic objects, each of which has two or more candidate entries in the GeoNames gazetteer. A good example of this pattern would be a location description such as 'port-au-prince', which contains one geographic object ('port-au-prince'). There are two candidate entries in the GeoNames gazetteer for 'port-au-prince': the first is classed as a *capital of a political entity* (PPLC); the second is classed as a *third-order administrative division* (ADM3). Consequently, the uncertainty metric is 0.5 for this location description. Further examples of this pattern are shown in Figure 5.20.

The second research question that motivated the development of this software prototype was to explore the relationship between the ambiguity associated with the resolution process and the spatial distribution (i.e. distance and direction) of entries in the GeoNames gazetteer that were selected by the resolution process. On reflection, however, the second research question was poorly specified: the spatial distribution of these entries has no bearing on the ambiguity associated with the resolution process, because the ambiguity associated with the resolution process is non-spatial. Indeed, whilst it might have been interesting to explore the spatial distribution of entries in the GeoNames gazetteer that shared the same name—for example, the spatial distribution of the 75 entries that shared the name 'bellevue'—doing so would not have contributed much to an understanding of the resolution process. Unsurprisingly, then, no relationship was evident between these factors. Figure 5.19, for example, shows how the least ambiguous location descriptions can have both clustered and dispersed spatial distributions.

On reflection, the design of this software prototype could have better supported the investigation of the ambiguity associated with the resolution process by serving as an interface to the geoparsing system, rather than as an interface to the results of the geoparsing system. As a reminder, the geoparsing system, which is described in Chapter 4, Section 4.1.2, consisted of identification and resolution processes. The latter consisted of three criteria—complete equality, distance, and feature codes—that resolved a match to a single entry in the gazetteer file. This software prototype could have allowed the visual inspection of the effects of the

[4]Note that in this example the geographic object 'haiti' is not identified and resolved as 'haiti' is not an entry in the Haiti subset of the GeoNames database.

(a) Map showing 'morne basile, haiti' (one geographic object, which has one candidate entry in the GeoNames gazetteer).

(b) Map showing 'petit goave' (one geographic object, which has one candidate entry in the GeoNames gazetteer).

Figure 5.19: The least ambiguous location descriptions often contain one or two geographic objects, each of which has one candidate entry in the GeoNames gazetteer.

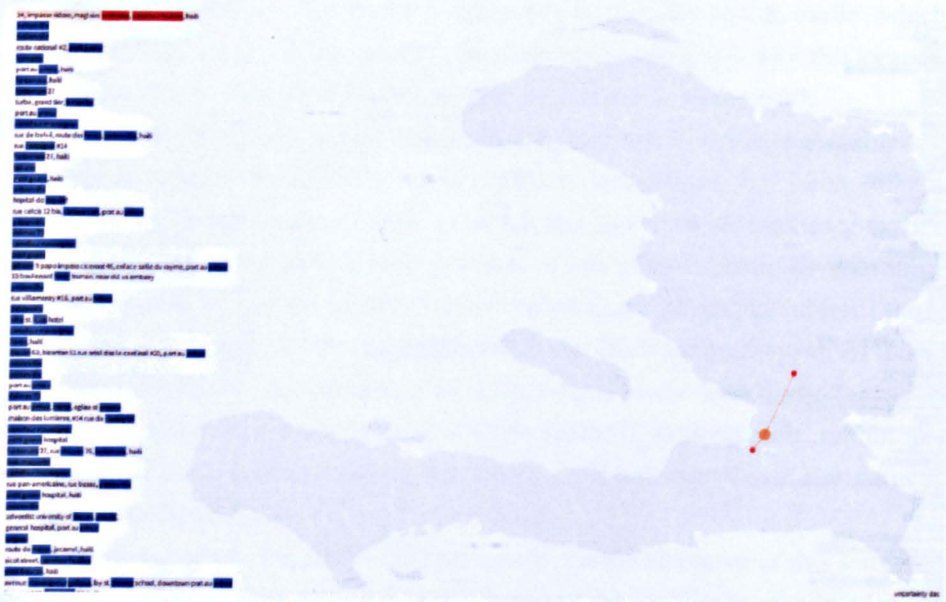

(c) Map showing '34, impasse ridore, magloire ambroise, carrefour feuilles, haiti' (two geographic objects, each of which has one candidate entry in the GeoNames gazetteer).

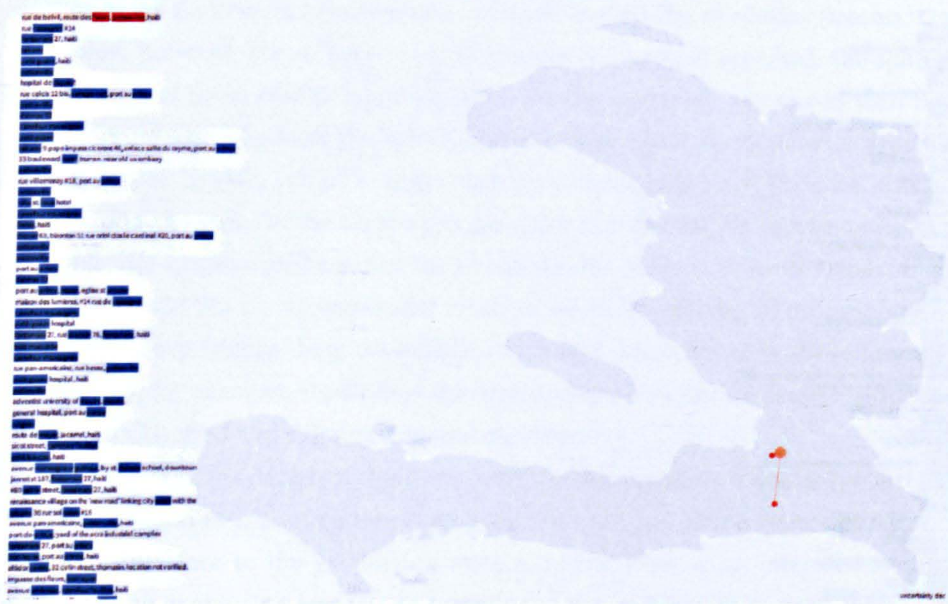

(d) Map showing 'rue de belvil, route des freres, petitionville, haiti' (two geographic objects, each of which has one candidate entry in the GeoNames gazetteer).

Figure 5.19: The least ambiguous location descriptions often contain one or two geographic objects, each of which has one candidate entry in the GeoNames gazetteer.

(a) Map showing 'port-au-prince' (one geographic object, which has two candidate entries in the GeoNames gazetteer).

(b) Map showing 'leogane, haiti' (one geographic object, which has two candidate entries in the GeoNames gazetteer).

Figure 5.20: More ambiguous location descriptions often contain one or two geographic objects, each of which has two or more candidate entries in the GeoNames gazetteer.

(c) Map showing 'jacmel, ??? route de meyer' (two geographic objects, which have four candidate entries in the GeoNames gazetteer between them).

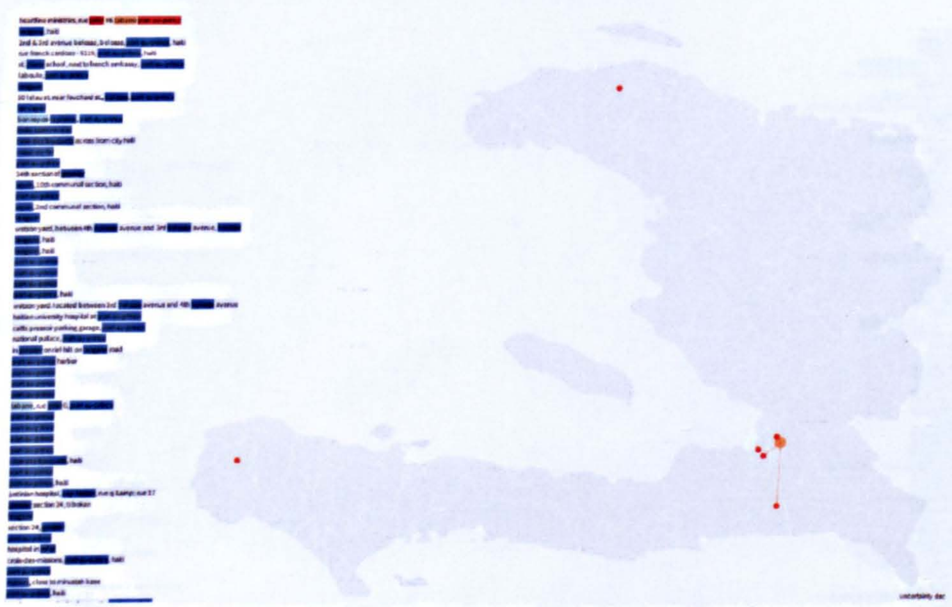

(d) Map showing 'heartline ministries, rue gelin #6 tabarre port-au-prince' (three geographic objects, which have six candidate entries in the GeoNames gazetteer between them).

Figure 5.20: More ambiguous location descriptions often contain one or two geographic objects, each of which has two or more candidate entries in the GeoNames gazetteer.

presence, absence, and order of these criteria—as well as additional criteria—on the geoparsing system, both locally (for each location description) and globally (for all location descriptions). It could also have allowed entries to be added dynamically to the gazetteer, and again, for the visual inspection of the effects of these additions.

As was noted in the introduction to this chapter, the development of each software prototype was itself an evaluation of the design, analysis, and evaluation model, which is discussed in Section 5.1.4 and elsewhere (Dillingham et al. 2012*b*). Consequently, the utility of the model is reflected on below in Section 5.5.

5.5 Reflections on the design, analysis, and evaluation model

There were two practical limitations with, and several flaws in, the design, analysis, and evaluation model. The first practical limitation was that the scenario, features, and justifications added little to the design process: rather than serving as a creative stimulus, each stage simply emphasised the existing knowledge of the designer-user. The reason for this limitation was that in scenario-based design, the scenario, features, and justifications should serve as a basis for discussion: they should be considered, deliberated, and ultimately either accepted or rejected. In the design, analysis, and evaluation model, however, the scenario and features became statements of intent: they were accepted without discussion. Consequently, the capacity for reflection was limited and the most significant benefit of scenario-based design—that it encourages "reflection in the context of design" (Carroll 2000, p.47)—was only partially realised.

How might the design, analysis, and evaluation model be modified to address this limitation? Clearly, when the designer is the user there needs to be an external source of ideas to challenge them and to enhance their capacity for reflection. This external source could take several forms. For example, it could be an existing task taxonomy: the designer-user could be challenged to implement the tasks, and then to undertake them, in the anticipation that doing so would identify strengths and limitations of the tasks with respect to the dataset, and of the capacity of the dataset for further investigation. Although the term 'task' has many meanings in the visualisation literature (Munzner 2009), there is a growing consensus as to how different tasks might be related to each other, and to how the concept of a 'task' is related to the concept of an 'interaction' (see, for example, Roth 2012, 2013*a*,*b*). Consequently, the choice of task taxonomy is unlikely to be arbitrary.

An alternative external source of ideas to challenge the designer-user and to enhance their capacity for reflection could come from their peers, consulted either individually or collectively. Lloyd & Dykes (2011) successfully used peer-review in the design process when they consulted experienced practitioners on the design of a series of geovisualisation software prototypes. However, there is the risk that peer-review emphasises the majority view, which may not be appropriate, even when the majority view is also the expert view (Kahneman 2012). Consequently, techniques that seek to enhance creativity could be adopted when consulting experienced practitioners. Goodwin et al. (2013) successfully used these techniques when designing tools for energy analysts and modellers. It would be interesting to explore whether these techniques, and the engagement of experienced practitioners, could be combined with the design, analysis, and evaluation model to encourage "reflection in the context of design" (Carroll 2000, p.47).

Sharing scenarios with the crowd (Howe 2006, 2009) and encouraging discussion, possibly using a software platform like Wikipedia (Wikipedia 2013) or OpenStreetMap (OpenStreetMap 2013), could also challenge the designer-user and enhance their capacity for reflection. Such an approach has obvious parallels with *ManyEyes* (Viegas et al. 2007), a software platform for sharing visualisations on the web, and with the open source movement more generally (Raymond 1999, Benkler & Nissenbaum 2006). Nevertheless, the willingness of the crowd to engage with a scenario is debatable. However, it would be interesting to explore crowdsourcing within the context of the design, analysis, and evaluation model.

The second practical limitation with the design, analysis, and evaluation model was that using the model was time consuming—in some cases, for example, it took longer to justify than to implement a feature—and that although the *write a scenario* stage was omitted from the development of the second and third software prototypes, using the model still proved to be antithetical to the principle of sketching, which is discussed in Section 5.3.

How might the design, analysis and evaluation model be modified to address this limitation? Clearly, it is important to identify the features that will be implemented in a geovisualisation software prototype. However, perhaps it is less important to justify these features when they are known to be effective. Again, there is a role for consulting experienced practitioners—if not directly, then through expert reviews (Tory & Moller 2005) or heuristic evaluations (Zuk et al. 2006), for example—as well as for consulting the visualisation literature to establish what is known to be effective. However, it is not necessary to have a formal documentation process to demonstrate that identification and justification took place. Instead, a more informal documentation process, such as that adopted by Lloyd & Dykes (2011) or Goodwin et al. (2013), which is characterised by sketches, field notes, and Post-It notes, would seem more appropriate and would be broadly aligned with other design disciplines, such as architecture (Kosara et al. 2008).

Practical limitations aside, however, there were several flaws in the model, or at least with the reasoning process that motivated its development. The first flaw was that it was hard to separate developing a design and documenting design decisions from undertaking analysis. During both design and analysis activities, new working hypotheses were inevitably being formulated. For example, it was not clear at exactly what stage it was found that for almost all location descriptions that contain administrative regions, the administrative region is the last geographic object in the location description. Following Chapter 4, the composition and organisation of individual location descriptions were explored, which prompted the design and development described in Section 5.4.2. In the context of the design, analysis, and evaluation model it could be said that an analysis activity led to a design activity. Consequently, it was impossible to justify the claim that a design led to an insight. Indeed, in the case of administrative regions, it could be said that an insight led to a design. Given the complex, qualitative, and unexpected nature of insight (North 2006) this flaw suggested that it might be impossible to justify the claim that a design led to an insight when the context of use and the context of design are blurred.

The second flaw with the model was that the reasoning process that motivated its development overlooked the question of *why* a design should be evaluated. Munzner (2009) argues that the answer to this question is *to reduce the threats to its validity* and identifies four nested levels at which these threats manifest themselves: because the levels are nested, the threats propagate from outer to inner levels (Figure 5.21). Scenarios, features, and justifications were intended to reduce the threats associated with characterising and abstracting the domain. However, when the designer is the user, the threats at these levels are minimal (if the designer cannot characterise and abstract their own domain, then the threat clearly lies elsewhere). Continuing with this reasoning, the main threat associated with developing a design when the context of use and the context of design are blurred is that it fails to communicate effectively with the user (Munzner 2009). Again, however, when the designer is the user, the threat at this level is minimal. Certainly, the usability of the design need not be evaluated too stringently, as the designer should know how to use their design effectively. Finally, perhaps the biggest threat to a design, which is not considered by Munzner (2009), is the threat of subjectivity: that a design was influenced by opinion rather than fact. However, the design process

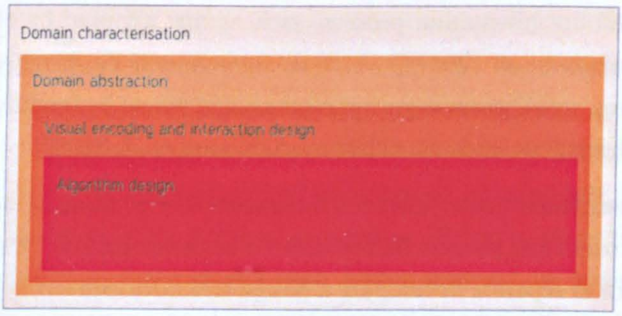

Figure 5.21: A nested model for visualisation design and validation (Munzner 2009). The single hue, sequential ColorBrewer Oranges scheme emphasises the propagation of threats from outer to inner levels.

is inherently subjective: it is a "creative process" that involves distinguishing "good choices" from "bad choices" (Sedlmair et al. 2012, p.2432) where the former and the latter are a matter of judgement, not measurement.

Further reflection on the flaws in the design, analysis, and evaluation model suggested that when the context of use and the context of design are blurred, evaluation should be more a matter of verification than validation. In other words, the question *Are we building the product right?* is more important than the question *Are we building the right product?* (Munzner 2009). In many respects, the former is easier to answer than the latter as 'best practice' guidelines are available. Aruliah et al. (2012), for example, describe several 'best practice' guidelines that range from writing good quality code to using appropriate tools. Nevertheless, the design, analysis, and evaluation model clarifies the nature of the gap between the context of use and the context of design.

5.6 Conclusion

Three geovisualisation software prototypes were developed to support the interactive visual exploration of reports from the Haiti crisis map: their goals were to explore the spatial and temporal distributions of reports; to explore the composition and organisation of location descriptions; and to explore the ambiguity associated with the composition of location descriptions. Collectively, the software prototypes represented a geovisualisation approach to crowdsourced crisis information. This approach was characterised by the overlap between the context of use and the context of design: a situation where the designer was also the user of a software artefact. This situation raised the important question of how a software artefact should be evaluated when the context of use and the context of design overlapped.

The design, analysis, and evaluation model was offered as an answer to this question: it was a synthesis of ideas from the fields of action research, human–computer interaction, and visualisation.

The design, analysis, and evaluation model was used to structure the discussion of each software prototype. However, as a precursor to these discussions, appropriate technologies, principles, and techniques were considered. From the perspective of technologies, it was found that the benefits of D3/JavaScript did not outweigh the risks and so the software prototypes were developed in Processing/Java. Principles and techniques were considered as visual encoding principles and interaction principles. Visual encoding principles included the careful use of the visual variables to establish a graphic hierarchy; the consistent use of colour value (lightness) and colour saturation both within and between the software prototypes; and removing axis ticks, labels, and lines from charts. Interaction principles included the consistent use of panning and zooming and the consistent use of brushing between the software prototypes.

The discussions of the software prototypes led to reflection on the design, analysis, and evaluation model, and to a reassessment of the research reported in Chapter 4. Taking these contributions in turn, the goal of the first software prototype was to explore the spatial and temporal distributions of reports from the Haiti crisis map. Following the design, analysis, and evaluation model, the working hypotheses, scenario, features, design decisions, and findings were reported. Whilst the findings were not especially interesting, the first software prototype provided a basis for the development of the subsequent software prototypes. Furthermore, the first software prototype demonstrated a practical limitation with the design, analysis, and evaluation model: that the most significant benefit of scenario-based design—that it encourages "reflection in the context of design" (Carroll 2000, p.47)—was only partially realised.

The goal of the second software prototype was to explore the composition and organisation of location descriptions. Consequently, this software prototype was developed to extend the research reported in Chapter 4. Three patterns of geographic objects were evident in location descriptions from using the second software prototype. These patterns lent support to the finding that geographic objects tended to be paired in a finer-to-coarser granularity pattern. The goal of the third software prototype was to explore the ambiguity associated with the composition of location descriptions. Consequently, this software prototype was developed to evaluate the geoparsing system that was described in Chapter 4, Section 4.1.2. It was found that, in the majority of cases, the entries in the GeoNames gazetteer that were selected by the resolution process were both distant from, and dispersed about, the crowdsourced point location. That instances of the small distance, small dispersion pattern were not evident in the majority of cases suggests that the resolution process performed poorly.

The research overview, which is shown in Chapter 1, Figure 1.1 and at the beginning of this chapter, emphasises that there are two strands to the current research. The first strand, which relates to the nature of crowdsourced crisis information, is positioned to the left of the research overview and encompasses Chapters 2, 3, and 4. The second strand, which relates to the visualisation of crowdsourced crisis information, is positioned to the right of the research overview and encompasses Chapter 5. These strands are joined in the next chapter, which concludes the current research.

Chapter 6

Conclusion

There are two strands to the current research. The first strand, which relates to the nature of crowdsourced crisis information, is positioned to the left of the research overview and encompasses Chapters 2, 3, and 4. The second strand, which relates to the visualisation of crowdsourced crisis information, is positioned to the right of the research overview and encompasses Chapter 5. This chapter joins the two strands. It considers the contributions of the current research. It then reflects on the motivations of the current research by addressing the question *To what extent does crisis mapping guard against inaccurate information from untrustworthy sources?* Finally, it concludes with an overview of future research directions.

6.1 Research contributions

The current research contributes to the fields of GISc (Geographic Information Science) and crisis informatics; crisis mapping; and geovisualisation specifically and information visualisation more generally. This section considers the contributions of the current research to each of these fields.

6.1.1 GISc and crisis informatics

The current research contributes an understanding of, and an approach to, the geographic uncertainty associated with crowdsourced crisis information. From a crisis informatics perspective, Palen et al. (2010, p.6) call for research into the characteristics of crisis information. The authors focus on accuracy, which can be seen as a component of uncertainty (Fisher 1999, MacEachren et al. 2005). From a GISc perspective, Devillers et al. (2010, p.396) call for research into new sources of spatial data, such as volunteered geographic information (Goodchild 2007). Crowdsourced crisis information can be thought of as volunteered geographic information. Consequently, by contributing an understanding of, and an approach to, the geographic uncertainty associated with crowdsourced crisis information, the current research also contributes to the research agendas proposed by Palen et al. (2010) and Devillers et al. (2010). The understanding and the approach contributed by the current research are now considered in turn.

The *understanding* of the geographic uncertainty associated with crowdsourced crisis information contributed by the current research is twofold. First, Chapter 3 concluded that location descriptions from the Haiti crisis map do not match the types of location descriptions identified by Wieczorek et al. (2004) and Guo et al. (2008) and, consequently, that location descriptions from the Haiti crisis map do not 'fit' the conceptual model proposed by Guo et al. (2008), which was discussed in Chapter 2, Section 2.2.1. Second, Chapter 4 concluded that location descriptions from the Haiti crisis map tended not to contain vague places; tended to contain more, relatively coarse granularity types of geographic objects (i.e. more than relatively fine types of geographic objects); and tended to be paired in a finer-to-coarser granularity pattern.

The implications of the finding that location descriptions from the Haiti crisis map do not 'fit' the conceptual model proposed by Guo et al. (2008) are significant because this finding suggests that existing georeferencing methods that estimate the uncertainty associated with free-text location descriptions—namely the point-radius and the probability distribution methods (Wieczorek et al. 2004, Guo et al. 2008)—cannot be easily applied to location descriptions from the Haiti crisis map.

Consequently, this finding suggests that there is a need for new or enhanced georeferencing methods that attempt to estimate the uncertainty associated with free-text location descriptions from sources of crowdsourced crisis information. The findings about the nature of location descriptions from the Haiti crisis map can serve as a basis for developing these new or enhanced georeferencing methods. The absence of vague places and the presence of more, relatively coarse granularity types of geographic objects can be used to select appropriate gazetteers for the identification stage of a future geoparsing system; the finer-to-coarser granularity pattern can be used to enhance the resolution stage of this system and can also serve as a hypothesis to test during subsequent exploratory data analysis. Consequently, as well as contributing an understanding of the characteristics of crisis information (Palen et al. 2010, p.6) and of a new source of spatial data (Devillers et al. 2010, p.396), the current research also indirectly supports the call for more research into georeferencing (Goldberg 2011).

The *approach* to the geographic uncertainty associated with crowdsourced crisis information contributed by the current research started with classification: first of location descriptions, which is described in Chapter 3, and second of geographic objects and spatial relationships contained in location descriptions, which is described in Chapter 4. The approach finished with the development of three geovisualisation software prototypes to support the interactive visual exploration of reports from the Haiti crisis map, which are described in Chapter 5. These prototypes synthesised techniques from GISc, geovisualisation, and natural language processing. The approach is best summarised by the research overview, which is shown at the beginning of this chapter.

The approach can be thought of as a partial inversion of the idealised research sequence presented by DiBiase (1990, Figure 1). In this sequence, visual thinking, which takes place in the private realm, moves from exploration to confirmation.[1] However, the approach started with two studies that were confirmatory and then moved to (or proceeded in parallel with) the development of three geovisualisation software prototypes that were exploratory. The movement from confirmatory to exploratory was necessary because the initial assessment of the nature of the location descriptions was erroneous. Initially, the location descriptions were assumed to be the natural language equivalent of the crowdsourced point locations: this assumption was made on the basis of the available documentation (Ushahidi 2013*a*) and following visual inspection of the location descriptions. However, whilst the location descriptions were clearly related to the crowdsourced point locations, they were often very 'messy', as the findings from Chapter 3 and 4 demonstrate.

[1] According to DiBiase (1990, Figure 1), visual thinking is followed by visual communication, which takes place in the public realm and encompasses synthesis and presentation.

The movement from a confirmatory to an exploratory approach reflects the difficulties of moving from a conceptual model of free-text location descriptions to the location descriptions themselves, rather than moving from the location descriptions themselves to a conceptual model of free-text location descriptions: that is, it reflects the difficulties of deductive (top-down), rather than inductive (bottom-up), reasoning when the subject of investigation is poorly understood. In the case of the current research, the conceptual model was that proposed by Guo et al. (2008). Although this model adequately describes location descriptions stored in the Mammal Networked Information System (Wieczorek et al. 2004, Guo et al. 2008) and related to search and rescue incidents (Doherty et al. 2011), it is clearly not well-suited to location descriptions from the Haiti crisis map; a finding that became apparent when exploring the composition and organisation of location descriptions during the development of the second software prototype, which is described in Chapter 5, Section 5.4.2.

Despite the difficulties it presented during the current research, the movement from a confirmatory to an exploratory approach is common within the field of natural language processing, where the gazetteer provides an implicit conceptual model of the location descriptions: in other words, because the gazetteer is independent of the location descriptions, and is selected before being used in a geoparsing system, it can be thought of as a conceptual model of the location descriptions. Consequently, the three geovisualisation software prototypes can be thought of as a basis for identifying a set of additional tasks: namely, those tasks that relate to judging the effectiveness of a geoparsing system and, more importantly, for recalibrating that system as required. The former set of tasks are partially addressed by the third software prototype, which is described in Chapter 5, Section 5.4.2. This software prototype allows the exploration of each location description in turn, rather than all location descriptions collectively, and does not summarise each location description. However, whilst it is rather cumbersome to use, it nevertheless represents a basis for future cycles of design, implementation, and evaluation. More importantly, however, the latter set of tasks offers an exciting opportunity to link natural language processing techniques and interactive visual exploration more closely. In this way, the current research also indirectly supports the call for more research that explores the links between natural language processing and GISc (Gregory & Hardie 2011), possibly within the field of geovisual analytics (Andrienko et al. 2007), as this research would need to integrate human and computer judgements on the effectiveness of a geoparsing system in an interactive, visual environment.

6.1.2 Crisis mapping

The current research contributes three geovisualisation software prototypes that can be used to identify meaningful patterns in crisis information. This contribution is discussed in Chapter 5, Sections 5.4.1, 5.4.2, and 5.4.3. Meier (2009) and Ziemke (2012) call for research into techniques that support the visual analysis of crisis information; techniques that help validate crisis information, uncover patterns, and test hypotheses. Consequently, the current research contributes to the crisis mapping research agenda proposed by Meier (2009) and Ziemke (2012). The contributions of the three geovisualisation software prototypes are now considered in turn.

The first software prototype, which is discussed in Chapter 5, Section 5.4.1, facilitates the exploration of the spatial and temporal distributions of reports and, with its spatial and temporal views, served as a basis for the second and third software prototypes. Reflection on this software prototype highlighted three requirements. First, the requirement for a global–local variant of the temporal histogram, which was introduced in the second software prototype. Second, the requirement for fading out the standard ellipse, to support the investigation of how the spatial distribution of reports changed over time. Although the second and third software prototypes did not include a standard ellipse, this requirement should nevertheless prove useful to designers of future geovisualisation software prototypes. Third, the requirement to explore the relationships between the spatial and temporal distributions of reports, and their corresponding location descriptions. This requirement is perhaps the most important, and can be thought of as the synthesis of the three software prototypes into a single application: in other words, as a step towards the interactive, visual environment mentioned above. Again, although this requirement was not addressed by the current research, it should nevertheless prove useful to designers of future geovisualisation software prototypes.

The second software prototype, which is discussed in Chapter 5, Section 5.4.2, facilitates the exploration of the composition and organisation of location descriptions using a thematic view that depicts each location description as a row and each geographic object contained in each location description as a cell. Use of this view uncovered several instances of a finer-to-coarser granularity pattern of geographic objects contained in location descriptions and, consequently, contributes to the crisis mapping research agenda proposed by Meier (2009) and Ziemke (2012). In addition, by linking the thematic view to the spatial and temporal views, it is possible to explore the spatial and temporal distributions of these patterns: a step towards the interactive, visual environment mentioned above. Reflection on this software prototype also highlighted three requirements. First, the requirement to cluster location descriptions according to a similarity measure and then to depict these clusters. Second, the requirement to implement multiple standard ellipses—one for each GeoNames feature class—that could be displayed simultaneously. The third requirement was an extension of the first and second requirements: to link a map and standard ellipse to a dendrogram, which is a common tree-based representation of a clustering process, to investigate the degree to which the spatial distributions of different clusters of location descriptions are different to each other. Again, although these requirements were not addressed by the current research, they should nevertheless prove useful to designers of future geovisualisation software prototypes.

The third software prototype, which is discussed in Chapter 5, Section 5.4.3, facilitates the exploration of the ambiguity associated with the composition of location descriptions using a thematic view that depicts the location descriptions themselves. In addition, the location descriptions can be ordered by an uncertainty metric. Use of this view highlighted a predictable relationship between ambiguity and number of geographic objects. Nevertheless, reflection on this software prototype highlighted an important requirement: namely, that a future software prototype could serve as the interface to a geoparsing system, rather than as an interface to the results of a geoparsing system. Given the reliance of most geoparsing systems on heuristics (see, for example, Grover et al. 2010), it would seem that a system that integrates human and computer judgements—that is, a geovisual analytics system (Andrienko et al. 2007)—could make a considerable contribution to the field of natural language processing. Again, although this requirement was not addressed by the current research, it should nevertheless prove useful of future geovisualisation software prototypes.

6.1.3 Geovisualisation

The current research contributes the design, analysis, and evaluation model, which situates the activities associated with designing a software artefact—and using it to undertake analysis—within an evaluative framework: it is an iterative approach to the design, implementation, and evaluation of a software artefact that attempts to bridge the gap between the context of use and the context of design. This contribution is discussed in Chapter 5, Sections 5.1.4 and 5.5.

The design, analysis, and evaluation model was proposed before, and then was developed in parallel with, the three geovisualisation software prototypes, which are discussed in Chapter 5, Section 5.4. Reflections on the model highlighted how "reflection in the context of design" (Carroll 2000, p.47) was only partially realised, and how using the model was time consuming and, consequently, proved to be antithetical to the principle of sketching. Furthermore, reflections on the model also highlighted how it was hard to separate developing a design and documenting design decisions from undertaking analysis. Perhaps most significant criticism, however, was that the model overlooked the question of *why* a design should be evaluated. When the context of use and the context of design are blurred, the threats to the validity of a design (Munzner 2009) are small: evaluation should be more a matter of verification than validation.

Despite the criticisms of the design, analysis, and evaluation model, there is still a need to establish the validity of a design when the context of use and the context of design are blurred. Even under these conditions, it is often the case that the designer will not be the sole user of the design: at some stage, other users will come into contact with it. Consequently, the question of how to validate a design when the context of use and the context of design are blurred is important. The design, analysis, and evaluation model offered one answer to this question. The model synthesised ideas from the fields of action research, human–computer interaction, and visualisation because these fields are closely associated with geovisualisation. However, other fields can provide alternative answers. Kosara et al. (2008), for example, highlights the role criticism plays in architecture. Perhaps a better approach to validating a design, which is common in architecture, would be to solicit criticism from practitioners. From a technical perspective, it is now easier than ever to make designs available to practitioners using the web; either as standalone applications using, for example, Processing (Reas & Fry 2006) or as browser-based applications using, for example, D3 (Bostock et al. 2011). Nevertheless, it could be hard to persuade busy practitioners to offer their criticism: the norms that exist within architecture do not currently exist within geovisualisation specifically and information visualisation more generally.

Two alternative approaches exist that serve as proxies for soliciting criticism from practitioners. The first is heuristic evaluation. However, heuristics exist for information visualisation (Zuk et al. 2006) but not geovisualisation. The second is the application of design guidelines, of which many exist within information visualisation and geovisualisation. In both cases, however, the question of how to report a heuristic evaluation or the application of design guidelines has not been answered. Whilst the design, analysis, and evaluation model provides one answer, reflections on the model suggest that it is not the best answer. Consequently, it would seem that there is a need for a more discursive means of reporting a heuristic evaluation or the application of design guidelines than is offered by the model. A contribution of the current research, then, is that it highlights this need.

6.2 Crisis mapping, inaccurate information, and untrustworthy sources

This section reflects on the motivations of the current research by addressing the question *To what extent does crisis mapping guard against inaccurate information from untrustworthy sources?*

As was discussed in Chapter 1, crisis mapping can be seen as an attempt to bridge the gaps between humanitarian organisations and crisis-affected communities: by facilitating the crowdsourcing of crisis information, crisis mapping claims to offer a solution to the problem of guarding against inaccurate information from untrustworthy sources, which should allow humanitarian organisations to integrate information from crisis-affected communities into their decision-making processes. The current research, however, has demonstrated the difficulties associated with substantiating this claim. Indeed, it has demonstrated that focusing on the geographic uncertainty associated with crowdsourced crisis information, by comparing free-text location descriptions to crowdsourced point locations, is of limited utility. Given that reports from the Haiti crisis map specifically and Ushahidi-based maps more generally do not indicate either the lineage or the credibility of the information, and that comparison—which could have indicated the accuracy or the consistency of the information—is of limited utility, the claim that crisis mapping offers a solution to the problem of guarding against inaccurate information from untrustworthy sources seems impossible to substantiate.

How should humanitarian organisations respond? Clearly, they still need to do more to integrate information from crisis-affected communities into their decision-making processes (Heinzelman & Waters 2010). Furthermore, rejecting the efforts of crisis mappers seems high-handed. A possible response, then, would be for humanitarian organisations to encourage crisis mappers to address a specific set of requirements. Researchers can mediate between humanitarian organisations and crisis mappers, and can help humanitarian organisations identify a specific set of requirements: the current research, for example, contributes three geovisualisation software prototypes that could serve as a basis for discussion with humanitarian organisations, raising awareness of the potential of crowdsourced crisis information.

Another possible response would be for humanitarian organisations and crisis mappers to engage more with the literature on visual analytics: "the science of analytical reasoning facilitated by interactive visual interfaces" (Thomas & Cook 2006, p.10). This field seems to be especially relevant to the problem of guarding against inaccurate information from untrustworthy sources, as both this field and the problem rely on a synthesis of human-based and computer-based approaches to possible solutions. The current research, by demonstrating the limited utility of making geographic comparisons, could be used to argue that crisis maps need to incorporate Wikipedia-like (Wikipedia 2013) or OpenStreetMap-like (OpenStreetMap 2013) processes to indicate the lineage and the credibility of the information, and consequently that crisis maps need to incorporate visual analytics-like approaches for reasoning with the lineage and the credibility information. In the context of crisis information, lineage seems especially important: if multiple reports are corroborated by one piece of information, and that piece of information is found to be spurious, or to have come from a spurious source, then these reports should be highlighted and crisis mappers should be tasked with finding alternative pieces of corroborating information.

6.3 Future research

Three areas of future research are situated within the contexts of crisis informatics and crisis mapping, GISc, and geovisualisation. In each case, a plan for how the future research should proceed is given.

6.3.1 Crisis informatics and crisis mapping

The current research was a case study of the Haiti crisis map. Since the Haiti crisis map, however, Ushahidi has become a *de facto* standard for crisis mapping, which itself has developed into a new field of study (Meier 2009, Ziemke 2012). Consequently, future research could adopt an empirical approach to investigate how crisis mappers georeference information from the crowd of contributors.

One means of investigating how crisis mappers georeference information from the crowd of contributors would be to conduct a field experiment. A smartphone application could ask participants—surrogates for members of the crisis-affected community—to describe where they were in a study location, without using a map. This application could also record the point location of each participant. Having gathered location descriptions and their corresponding point locations, a crisis map could then be simulated. Participants—surrogates for crisis mappers—could georeference location descriptions and their spatial ability could be measured (see, for example, Hegarty et al. 2002). The distance between each point location and its corresponding crowdsourced point location would indicate the accuracy associated with the crowdsourced point location. Correlation between accuracy and sense of direction could then be investigated.

If spatial ability and accuracy were correlated, then the implications for crisis mapping would be clear: the spatial ability of a crisis mapper could be used as a surrogate for the uncertainty associated with their contributions to a crisis map. Knowing more about the uncertainty associated with a crisis map could help analysts and decision-makers in their work.

6.3.2 GISc

As an alternative to a case study of the Haiti crisis map, a larger number of secondary sources could be analysed with the aim of identifying the characteristics of Ushahidi-based maps. This analysis could contribute an improved understanding of neo-geography (Goodchild 2009) and could update Liu & Palen (2010) to reflect the transition from crisis map mashups to crisis mapping. There are, however, several problems associated with the analysis of a larger number of secondary sources. The first is the problem of information access: whilst the Crowdmap service (Crowdmap 2013) hosts thousands of Ushahidi-based maps (Crowdglobe 2012), neither the service nor the maps have an API (application programming interface). Consequently, information access would need to be negotiated. The second is the problem of information availability: little would be known about the phenomena being mapped or about the crisis mappers themselves. Consequently, the analysis of a larger number of secondary sources would require more detailed consideration.

6.3.3 Geovisualisation

The design study methodology (Sedlmair et al. 2012) is a qualitative approach to evaluation that emphasises the context of use. However, fundamental to the design study methodology is a gap between the context of use and the context of design. The design, analysis, and evaluation model, which was discussed in Chapter 5, Section 5.1.4, was an attempt to bridge this gap. Its development highlighted the similarities between action research and the design study methodology. The third area of future research could explore the similarities between the design study methodology and action research further, through a comparative literature review. The aim of this comparative literature review would be to ground the design study methodology in the action research literature, so as to understand the scientific traditions of the design study methodology.

Bibliography

alias-i (2013), 'LingPipe'. v4.1 [Accessed 18 April 2013].
URL: http://alias-i.com/lingpipe/

Alvarez, G. A. & Franconeri, S. L. (2007), 'How many objects can you track?: Evidence for a resource-limited attentive tracking mechanism', *Journal of Vision* **7**(13), 1–10.

Andrienko, G., Andrienko, N., Jankowski, P., Keim, D., Kraak, M., MacEachren, A. & Wrobel, S. (2007), 'Geovisual analytics for spatial decision support: Setting the research agenda', *International Journal of Geographical Information Science* **21**(8), 839–857.

Andrienko, N. & Andrienko, G. (2006), *Exploratory analysis of spatial and temporal data: A systematic approach*, Springer, Berlin, Germany.

Aruliah, D., Brown, C. T., Hong, N. P. C., Davis, M., Guy, R. T., Haddock, S. H., Huff, K., Mitchell, I., Plumbley, M., Waugh, B., White, E. P., Wilson, G. & Wilson, P. (2012), Best practices for scientific computing, Technical report, Cornell University, Ithaca, NY, USA. [Accessed 3 October 2012].
URL: http://arxiv.org/abs/1210.0530

Bailey, T. C. & Gatrell, A. C. (1995), *Interactive Spatial Data Analysis*, Longman Scientific and Technical, Harlow, UK.

Baskerville, R. L. & Wood-Harper, A. T. (1996), 'A critical perspective on action research as a method for information systems research', *Journal of Information Technology* **11**(3), 235–246.

Becker, R. A. & Cleveland, W. S. (1987), 'Brushing scatterplots', *Technometrics* **29**(2), 127–142.

Benkler, Y. & Nissenbaum, H. (2006), 'Commons-based peer production and virtue', *Journal of Political Philosophy* **14**(4), 394–419.

Bertin, J. (2011), *Semiology of Graphics: Diagrams, Networks, Maps*, ESRI Press, Redlands, CA, USA. Translated from French by William J. Berg.

Bird, S., Klein, E. & Loper, E. (2009), *Natural Language Processing with Python*, O'Reilly Media, Sebastopol, CA, USA.

Bittner, T. & Smith, B. (2001), A taxonomy of granular partitions, *in* D. R. Montello, ed., 'Spatial Information Theory: Foundations of Geographic Information Science', Vol. 2205 of *Lecture Notes in Computer Science*, Springer, Berlin, Germany, pp. 28–43.

Bostock, M. & Davies, J. (2013), 'Code as cartography', *The Cartographic Journal* **50**(2), 129–135.

Bostock, M. & Heer, J. (2009), 'Protovis: A graphical toolkit for visualization', *IEEE Transactions on Visualization and Computer Graphics* **15**(6), 1121–1128.

Bostock, M., Ogievetsky, V. & Heer, J. (2011), 'D3: Data-driven documents', *IEEE Transactions on Visualization and Computer Graphics* **17**(12), 2301–2309.

Brewer, C. A. (1997), 'Evaluation of a model for predicting simultaneous contrast on color maps', *The Professional Geographer* **49**(3), 280–294.

Card, S. K., Mackinlay, J. D. & Shneiderman, B. (1999), *Readings in Information Visualization: Using Vision to Think*, The Morgan Kaufmann series in interactive technologies, Morgan Kaufmann, San Francisco, CA, USA.

Carroll, J. (2000), 'Five reasons for scenario-based design', *Interacting with Computers* **13**(1), 43–60.

Carroll, J. M. & Rosson, M. B. (1992), 'Getting around the task-artifact cycle: how to make claims and design by scenario', *ACM Transactions on Information Systems* **10**(2), 181–212.

CIA (2013), 'CIA world factbook: Haiti'. [Accessed 2 July 2013].
URL: `https://www.cia.gov/library/publications/the-world-factbook/geos/ha.html`

Corbane, C., Lemoine, G. & Kauffmann, M. (2012), 'Relationship between the spatial distribution of sms messages reporting needs and building damage in 2010 Haiti disaster', *Natural Hazards and Earth System Sciences* **12**(2), 255–265.

Coyle, D. & Meier, P. (2009), New technologies in emergencies and conflicts: The role of information and social networks, Technical report, UN Foundation–Vodafone Foundation Partnership, Washington DC, USA and London, UK.

Crampton, J. W. (2002), 'Interactivity types in geographic visualization', *Cartography and Geographic Information Science* **29**(2), 85–98.

Crockford, D. (2008), *JavaScript: The Good Parts*, O'Reilly Media (Yahoo Press), Sebastopol, CA, USA.

Crooks, A., Croitoru, A., Stefanidis, A. & Radzikowski, J. (2013), '#Earthquake: Twitter as a distributed sensor system', *Transactions in GIS* **17**(1), 124–147.

Crooks, A. T. & Wise, S. (2013), 'GIS and agent-based models for humanitarian assistance', *Computers, Environment and Urban Systems* **41**, 100–111.

Crowdglobe (2012), Mapping the maps: A meta-level analysis of Ushahidi and Crowdmap, Technical report, Internews Center for Innovation and Learning, Washington DC, USA.

Crowdmap (2013), 'About Crowdmap'. [Accessed 16 May 2013].
URL: https://crowdmap.com/about/

Devillers, R., Stein, A., Bdard, Y., Chrisman, N., Fisher, P. & Wenzhong, S. (2010), 'Thirty years of research on spatial data quality: achievements, failures, and opportunities', *Transactions in GIS* **14**(4), 387–400.

Dewar, M. (2012), *Getting Started with D3: Creating Data-Driven Documents*, O'Reilly Media, Sebastopol, CA, USA.

DiBiase, D. (1990), 'Visualization in the earth sciences', *Earth and Mineral Sciences* **59**(2), 13–18.

Dillingham, I., Dykes, J. & Wood, J. (2011), 'Visual analytical approaches to evaluating uncertainty and bias in crowdsourced crisis information', Poster presented at the IEEE Conference on Visual Analytics Science and Technology (VAST 2011). Providence, RI, USA, 23–28 October 2011.

Dillingham, I., Dykes, J. & Wood, J. (2012*a*), Characterising locality descriptions in crowdsourced crisis information, *in* 'Proceedings of the GIS Research UK 20th Annual Conference (GISRUK 2012)'. Lancaster University, Lancaster, UK, 11–13 April 2012.

Dillingham, I., Dykes, J. & Wood, J. (2012*b*), 'A design, analysis and evaluation model to support the visualization designer-user', Poster presented at the IEEE Conference on Information Visualization (InfoVis 2012). Seattle, WA, USA, 14–19 October 2012.

Dillingham, I., Dykes, J. & Wood, J. (2012*c*), 'Exploring patterns of uncertainty in crowdsourced crisis information', Poster presented at the EuroVis Workshop on Visual Analytics (EuroVA 2012). Vienna, Austria, 4–5 June 2012.

Dillingham, I., Dykes, J. & Wood, J. (2013), Visualizing the geographies of the Haiti crisis map, *in* 'GeoViz Hamburg: Interactive Maps that Help People Think'. HafenCity University Hamburg, Hamburg, Germany, 6–8 March 2013.

Dillingham, I., Mills, B. & Dykes, J. (2011), 'Exploring road incident data with heat maps', Poster presented at the GIS Research UK 19th Annual Conference (GISRUK 2011). University of Portsmouth, Portsmouth, UK, 27–29 April 2011.

Dillingham, I., Wood, J. & Dykes, J. (2013), Types, granularities and combinations of geographic objects in the Haiti crisis map, *in* 'Proceedings of the GIS Research UK 21st Annual Conference (GISRUK 2013)'. The University of Liverpool, Liverpool, UK, 3–5 April 2013.

Disasters Emergency Committee (2013), 'Haiti earthquake facts and figures'. [Accessed 28 March 2013].
URL: `http://www.dec.org.uk/haiti-earthquake-facts-and-figures`

Doherty, P., Guo, Q., Liu, Y., Wieczorek, J. & Doke, J. (2011), 'Georeferencing incidents from locality descriptions and its applications: a case study from Yosemite National Park search and rescue', *Transactions in GIS* **15**(6), 775–793.

Dupré, L. (1998), *BUGS in Writing: A Guide to Debugging Your Prose*, Addison-Wesley, Boston, MA, USA.

Dykes, J. (2005), Facilitating interaction for geovisualization, *in* J. Dykes, A. M. MacEachren & M. Kraak, eds, 'Exploring Geovisualization', Elsevier, Oxford, UK, pp. 265–291.

Dykes, J., MacEachren, A. M. & Kraak, M.-J., eds (2005), *Exploring Geovisualization*, Elsevier, Oxford, UK.

Elwood, S., Goodchild, M. F. & Sui, D. Z. (2012), 'Researching volunteered geographic information: Spatial data, geographic research, and new social practice', *Annals of the Association of American Geographers* **102**(3), 571–590.

Fisher, P. (1999), Models of uncertainty in spatial data, *in* P. A. Longley, M. F. Goodchild, D. J. Maguire & D. W. Rhind, eds, 'Geographical Information Systems, Principles and Technical Issues', 2nd edn, Vol. 1, John Wiley and Sons, Chichester, UK, pp. 191–205.

Fisher, P. & Unwin, D. J. (2005), Re-presenting geographical information systems, *in* P. Fisher & D. J. Unwin, eds, 'Re-Presenting GIS', John Wiley and Sons, Chichester, UK, pp. 1–14.

Fleiss, J. L. (1971), 'Measuring nominal scale agreement among many raters', *Psychological Bulletin* **76**(5), 378–382.

Foley, P. & Demšar, U. (2013), 'Using geovisual analytics to compare the performance of geographically weighted discriminant analysis versus its global counterpart, linear discriminant analysis', *International Journal of Geographical Information Science* **27**(4), 633–661.

Fry, B. (2007), *Visualizing Data: Exploring and Explaining Data with the Processing Environment*, O'Reilly Media, Sebastopol, CA, USA.

GADM (2013), 'Global administrative areas'. v2.0 [Accessed 14 May 2013].
URL: `www.gadm.org`

Gamer, M., Lemon, J., Fellows, I. & Singh, P. (2012), *irr: Various Coefficients of Interrater Reliability and Agreement*. v0.84 [Accessed 23 April 2013].
URL: `http://CRAN.R-project.org/package=irr`

Gelernter, J. & Mushegian, N. (2011), 'Geo-parsing messages from microtext', *Transactions in GIS* **15**(6), 753–773.

GeoNames (2013*a*), 'GeoNames downloads'. [Accessed 17 April 2013].
URL: `http://download.geonames.org/export/dump/`

GeoNames (2013*b*), 'GeoNames feature codes'. [Accessed 17 April 2013].
URL: `http://www.geonames.org/export/codes.html`

Gerring, J. (2004), 'What is a case study and what is it good for?', *American Political Science Review* **98**(2), 341–354.

Gleicher, M., Albers, D., Walker, R., Jusufi, I., Hansen, C. D. & Roberts, J. C. (2011), 'Visual comparison for information visualization', *Information Visualization* **10**(4), 289–309.

Goldberg, D. W. (2011), 'Advances in geocoding research and practice', *Transactions in GIS* **15**(6), 727–733.

Goodchild, M. (2007), 'Citizens as sensors: the world of volunteered geography', *GeoJournal* **69**(4), 211–221.

Goodchild, M. (2009), 'NeoGeography and the nature of geographic expertise', *Journal of Location Based Services* **3**(2), 82–96.

Goodchild, M. F. & Glennon, J. A. (2010), 'Crowdsourcing geographic information for disaster response: a research frontier', *International Journal of Digital Earth* **3**(3), 231–241.

Goodchild, M. F. & Hunter, G. J. (1997), 'A simple positional accuracy measure for linear features', *International Journal of Geographical Information Science* **11**(3), 299–306.

Goodwin, S., Dykes, J., Jones, S., Dillingham, I., Dove, G., Duffy, A., Kachkaev, A., Slingsby, A. & Wood, J. (2013), 'Creative user-centered visualization design for energy analysts and modellers', *IEEE Transactions on Visualization and Computer Graphics* **19**(12), 2516–2525.

Greenberg, S. & Buxton, B. (2008), Usability evaluation considered harmful (some of the time), *in* 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems', pp. 111–120. Florence, Italy, 5–10 April 2008.

Gregory, I. N. & Hardie, A. (2011), 'Visual GISting: bringing together corpus linguistics and geographical information systems', *Literary and Linguistic Computing* **26**(3), 297–314.

Grover, C., Tobin, R., Byrne, K., Woollard, M., Reid, J., Dunn, S. & Ball, J. (2010), 'Use of the Edinburgh geoparser for georeferencing digitized historical collections', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **368**(1925), 3875–3889.

Guo, Q., Liu, Y. & Wieczorek, J. (2008), 'Georeferencing locality descriptions and computing associated uncertainty using a probabilistic approach', *International Journal of Geographical Information Science* **22**(10), 1067–1090.

Harrower, M. (2002), Visual Benchmarks: Representing Geographic Change with Map Animation, PhD thesis, Penn State University.

Harrower, M. & Brewer, C. (2003), 'ColorBrewer.org: An online tool for selecting colour schemes for maps', *The Cartographic Journal* **40**(1), 27–37.

Harrower, M. & Sheesley, B. (2005), 'Designing better map interfaces: A framework for panning and zooming', *Transactions in GIS* **9**(2), 77—89.

Heer, J. & Robertson, G. G. (2007), 'Animated transitions in statistical data graphics', *IEEE Transactions on Visualization and Computer Graphics* **13**(6), 1240–1247.

Hegarty, M., Richardson, A. E., Montello, D. R., Lovelace, K. & Subbiah, I. (2002), 'Development of a self-report measure of environmental spatial ability', *Intelligence* **30**(5), 425–447.

Heinzelman, J. & Waters, C. (2010), Crowdsourcing crisis information in disaster-affected Haiti, Special Report 252, United States Institute of Peace, Washington DC, USA.

Howe, J. (2006), 'The rise of crowdsourcing', *Wired* **14**(6), 176–183. [Accessed 16 March 2011].
URL: `http://www.wired.com/wired/archive/14.06/crowds.html`

Howe, J. (2009), *Crowdsourcing: How the Power of the Crowd is Driving the Future of Business*, Random House, London, UK.

Isenberg, P., Zuk, T., Collins, C. & Carpendale, S. (2008), Grounded evaluation of information visualizations, *in* 'Proceedings of the 2008 workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV 2008)', pp. 56–63. Florence, Italy, 5 April 2008.

Jones, C. B., Purves, R. S., Clough, P. D. & Joho, H. (2008), 'Modelling vague places with knowledge from the web', *International Journal of Geographical Information Science* **22**(10), 1045–1065.

Jurafsky, D. & Martin, J. H. (2009), *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, 2nd edn, Pearson Prentice Hall, Upper Saddle River, NJ, USA.

Kachkaev, A., Dillingham, I., Beecham, R., Goodwin, S., Ahmed, N. & Slingsby, A. (2012), Monitoring the health of computer networks with visualization: VAST 2012 mini challenge 1 award: "efficient use of visualization", *in* 'Proceedings of the IEEE Conference on Visual Analytics Science and Technology (VAST 2012)', pp. 269–270. Seattle, WA, USA, 14–19 October 2012.

Kahneman, D. (2012), *Thinking, Fast and Slow*, Penguin, London, UK.

Keim, D., Kohlhammer, J., Ellis, G. & Mansmann, F., eds (2010), *Mastering the Information Age: Solving Problems with Visual Analytics*, The Eurographics Association, Goslar, Germany.

Kitchin, R. & Tate, N. J. (2000), *Conducting Research into Human Geography: Theory, Methodology and Practice*, Pearson Education, Harlow, UK.

Koh, L. C., Slingsby, A., Dykes, J. & Kam, T. S. (2011), Developing and applying a user-centered model for the design and implementation of information visualization tools, *in* 'Proceedings of the 15th International Conference on Information Visualisation (IV 2011)', pp. 90–95. London, UK, 13–15 July 2011.

Kosara, R., Drury, F., Holmquist, L. E. & Laidlaw, D. H. (2008), 'Visualization criticism', *IEEE Computer Graphics and Applications* **28**(3), 13–15.

Kosara, R., Healey, C. G., Interrante, V., Laidlaw, D. H. & Ware, C. (2003), 'Visualization viewpoints. user studies: Why, How, and When?', *IEEE Computer Graphics and Applications* **23**(4), 20–25.

Lam, H. (2008), 'A framework of interaction costs in information visualization', *IEEE Transactions on Visualization and Computer Graphics* **14**(6), 1149–1156.

Landis, J. R. & Koch, G. G. (1977), 'The measurement of observer agreement for categorical data', *Biometrics* **33**(1), 159–174.

Leidner, J. L. (2007), Toponym Resolution in Text, PhD thesis, The University of Edinburgh. [Accessed 15 April 2013].
URL: `http://www.era.lib.ed.ac.uk/handle/1842/1849`

Liu, S. B. & Palen, L. (2010), 'The new cartographers: Crisis map mashups and the emergence of neogeographic practice', *Cartography and Geographic Information Science* **37**(1), 69–90.

Lloyd, D. & Dykes, J. (2011), 'Human-centered approaches in geovisualization design: Investigating multiple methods through a long-term case study', *IEEE Transactions on Visualization and Computer Graphics* **17**(12), 2498–2507.

Longley, P., Goodchild, M. F., Maguire, D. J. & Rhind, D. W. (2011), *Geographic Information Systems and Science*, 3rd edn, Wiley, Hoboken, NJ, USA.

MacEachren, A. M. (1994*a*), *Some Truth with Maps: A Primer on Symbolization and Design*, Association of American Geographers, Washington DC, USA.

MacEachren, A. M. (1994*b*), Visualization in modern cartography: Setting the agenda, *in* A. M. MacEachren & D. Taylor, eds, 'Visualization in Modern Cartography', Vol. 2 of *Modern Cartography*, Pergamon, Oxford, UK, pp. 1–12.

MacEachren, A. M., Jaiswal, A., Robinson, A. C., Pezanowski, S., Savelyev, A., Mitra, P., Zhang, X. & Blanford, J. (2011), SensePlace2: GeoTwitter analytics support for situational awareness, *in* 'Proceedings of the IEEE Conference on Visual Analytics Science and Technology', pp. 181–190. Providence, RI, USA, 23–28 October 2011.

MacEachren, A. M., Robinson, A., Hopper, S., Gardner, S., Murray, R., Gahegan, M. & Hetzler, E. (2005), 'Visualizing geospatial information uncertainty: What we know and what we need to know', *Cartography and Geographic Information Science* **32**(3), 139–160.

MapAction (2011), 'About MapAction'. [Accessed 15 August 2013].
URL: http://www.mapaction.org/about.html

McClendon, S. & Robinson, A. C. (2012), Leveraging geospatially-oriented social media communications in disaster response, *in* 'Proceedings of the 9th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2012)'. Vancouver, Canada, 22–25 April 2012.

Meier, P. (2009), 'Proposing the field of crisis mapping'. [Accessed 21 May 2013].
URL: http://irevolution.net/2009/08/08/proposing-crisis-mapping/

Meier, P. (2012), 'Crisis mapping in action: How open source software and global volunteer networks are changing the world, one map at a time', *Journal of Map and Geography Libraries* **8**(2), 89–100.

Meier, P. (2013), 'Using crowdsourcing to counter the spread of false rumours on social media during crises'. [Accessed 21 May 2013].
URL: http://irevolution.net/2013/03/31/crowdsourcing-critical-thinking/

Meier, P. & Munro, R. (2010), 'The unprecedented role of SMS in disaster response: Learning from Haiti', *SAIS Review* **30**(2), 91–103.

Mernik, M., Heering, J. & Sloane, A. M. (2005), 'When and how to develop domain-specific languages', *ACM Computing Surveys (CSUR)* **37**(4), 316–344.

Meyer, M., Munzner, T., DePace, A. & Pfister, H. (2010), 'MulteeSum: A tool for comparative spatial and temporal gene expression data', *IEEE Transactions on Visualization and Computer Graphics* **16**(6), 908–917.

Meyer, M., Munzner, T. & Pfister, H. (2009), 'MizBee: A multiscale synteny browser', *IEEE Transactions on Visualization and Computer Graphics* **15**(6), 897–904.

Monmonier, M. (1989), 'Geographic brushing: Enhancing exploratory analysis of the scatterplot matrix', *Geographical Analysis* **21**(1), 81–84.

Munzner, T. (2008), Process and pitfalls in writing information visualization research papers, *in* A. Kerren, J. T. Stasko, J.-D. Fekete & C. North, eds, 'Information Visualization: Human-Centered Issues and Perspectives', Vol. 4950 of *Lecture Notes in Computer Science*, Springer, Berlin, Germany, pp. 134–153.

Munzner, T. (2009), 'A nested model for visualization design and validation', *IEEE Transactions on Visualization and Computer Graphics* **15**(6), 921–928.

Natural Earth (2013), 'Natural Earth'. [Accessed 14 May 2013].
URL: `http://www.naturalearthdata.com/`

Nielsen, C. B., Jackman, S. D., Birol, I. & Jones, S. J. (2009), 'ABySS-Explorer: Visualizing genome sequence assemblies', *IEEE Transactions on Visualization and Computer Graphics* **15**(6), 881–888.

Norman, D. A. (1998), *The Design of Everyday Things*, MIT Press, Cambridge, MA, USA.

North, C. (2006), 'Toward measuring visualization insight', *IEEE Computer Graphics and Applications* **26**(3), 6–9.

Oates, B. J. (2006), *Researching Information Systems and Computing*, Sage, London, UK.

OCHA (2011), 'Libya crisis map'. [Accessed 19 October 2011].
URL: `http://libyacrisismap.net/`

OpenStreetMap (2013), 'OpenStreetMap: The free wiki world map'. [Accessed 20 May 2013].
URL: `http://www.openstreetmap.org/`

O'Sullivan, D. & Unwin, D. J. (2003), *Geographic Information Analysis*, Wiley, Hoboken, NJ, USA.

Palen, L., Anderson, K. M., Mark, G., Martin, J., Sicker, D., Palmer, M. & Grunwald, D. (2010), A vision for technology-mediated support for public participation and assistance in mass emergencies and disasters, *in* 'Proceedings of the 2010 ACM-BCS Visions of Computer Science Conference', pp. 1–12. The University of Edinburgh, Edinburgh, UK, 14–16 April 2010.

Palen, L. & Liu, S. B. (2007), Citizen communications in crisis: Anticipating a future of ICT-supported public participation, *in* 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems', pp. 727–736. San Jose, CA, 30 April–3 May 2007.

Pike, W. A., Stasko, J., Chang, R. & O'Connell, T. A. (2009), 'The science of interaction', *Information Visualization* **8**(4), 263–274.

Plewe, B. (2002), 'The nature of uncertainty in historical geographic information', *Transactions in GIS* **6**(4), 431–456.

Processing (2013), 'Processing'. v1.5.1–2.0 [Accessed 3 July 2013].
URL: http://processing.org/

Raymond, E. S. (1999), *The Cathedral and the Bazaar: Musings on Linux and Open Source by and Accidental Revolutionary*, O'Reilly Media, Sebastopol, CA, USA.

Reas, C. & Fry, B. (2006), 'Processing: programming for the media arts', *AI & Society* **20**(4), 526–538.

Richter, D., Vasardani, M., Stirling, L., Richter, K.-F. & Winter, S. (2013), Zooming in–zooming out hierarchies in place descriptions, *in* J. M. Krisp, ed., 'Progress in Location-Based Services', Lecture Notes in Geoinformation and Cartography, Springer, Berlin, Germany, pp. 339–355.

Richter, D., Winter, S., Richter, K.-F. & Stirling, L. (2012), How people describe their place: identifying predominant types of place descriptions, *in* 'Proceedings of the 1st ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information (GEOCROWD 12)', pp. 30–37. Redondo Beach, CA, USA, 6–9 November 2012.

Richter, D., Winter, S., Richter, K.-F. & Stirling, L. (2013), 'Granularity of locations referred to by place descriptions', *Computers, Environment and Urban Systems* **41**, 88–99.

Rosson, M. B. & Carroll, J. M. (2002), Scenario-Based design, *in* J. A. Jacko & A. Sears, eds, 'The Human-Computer Interaction Handbook: Fundamentals, Evolving Technologies and Emerging Applications', Human Factors and Ergonomics, Lawrence Erlbaum Associates, Mahwah, NJ, USA, pp. 1032–1050.

Roth, R. E. (2009), 'A qualitative approach to understanding the role of geographic information uncertainty during decision making', *Cartography and Geographic Information Science* **36**(4), 315–330.

Roth, R. E. (2012), 'Cartographic interaction primitives: Framework and synthesis', *Cartographic Journal* **49**(4), 376–395.

Roth, R. E. (2013*a*), 'An empirically-derived taxonomy of interaction primitives for interactive cartography and geovisualization', *IEEE Transactions on Visualization and Computer Graphics* **19**(12), 2356–2365.

Roth, R. E. (2013*b*), 'Interactive maps: What we know and what we need to know', *Journal of Spatial Information Science* (6), 59–115. [Accessed 17 June 2013].
URL: http://www.josis.org/index.php/josis/article/view/105

Saraiya, P., North, C. & Duca, K. (2005), 'An insight-based methodology for evaluating bioinformatics visualizations', *IEEE Transactions on Visualization and Computer Graphics* **11**(4), 443–456.

Sedlmair, M., Meyer, M. & Munzner, T. (2012), 'Design study methodology: Reflections from the trenches and the stacks', *IEEE Transactions on Visualization and Computer Graphics* **18**(12), 2431–2440.

Shneiderman, B. (1996), The eyes have it: a task by data type taxonomy for information visualizations, *in* 'Proceedings of the IEEE Symposium on Visual Languages', pp. 336–343. Boulder, CO, USA, 3–6 September 1996.

Sieber, R. (2006), 'Public participation geographic information systems: A literature review and framework', *Annals of the Association of American Geographers* **96**(3), 491–507.

Slingsby, A., Dykes, J. & Wood, J. (2009), 'Configuring hierarchical layouts to address research questions', *IEEE Transactions on Visualization and Computer Graphics* **15**(6), 977–984.

Slocum, T. A., McMaster, R. B., Kessler, F. C. & Howard, H. H. (2009), *Thematic Cartography and Geovisualization*, Prentice Hall Series in Geographic Information Science, 3rd edn, Pearson Prentice Hall, Upper Saddle River, NJ, USA.

Starbird, K. (2012), Crowd computation: organizing information during mass disruption events, *in* 'Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work', pp. 339–342. Seattle, WA, USA, 11–15 February 2012.

Starbird, K. & Palen, L. (2011), "Voluntweeters": self-organizing by digital volunteers in times of crisis, *in* 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems', pp. 1071–1080. Vancouver, BC, Canada, 7–12 May 2011.

Susman, G. I. & Evered, R. D. (1978), 'An assessment of the scientific merits of action research', *Administrative Science Quarterly* **23**(4), 582–603.

Sutherlin, G. (2013), 'A voice in the crowd: Broader implications for crowdsourcing translation during crisis', *Journal of Information Science* **39**(3), 397–409.

Tapia, A. H., Bajpai, K., Jansen, J., Yen, J. & Giles, L. (2011), Seeking the trustworthy tweet: Can microblogged data fit the information needs of disaster response and humanitarian relief organizations?, *in* 'Proceedings of the 8th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2011)'. Lisbon, Portugal, 8–11 May 2011.

The Apache Software Foundation (2013), 'Apache Commons Lang'. v3.1 [Accessed 17 April 2013].
URL: http://commons.apache.org/proper/commons-lang/

Thomas, J. J. & Cook, K. A. (2006), 'A visual analytics agenda', *IEEE Computer Graphics and Applications* **26**(1), 10–13.

Thomas, J. T. & Cook, K. A., eds (2005), *Illuminating the Path: The Research and Development Agenda for Visual Analytics*, IEEE.

Thomson, J., Hetzler, B., MacEachren, A., Gahegan, M. & Pavel, M. (2005), A typology for visualizing uncertainty, *in* 'Conference on Visualization and Data Analysis', pp. 16–20. San Jose, CA, USA, 17–18 January 2005.

Tory, M. & Moller, T. (2005), 'Evaluating visualizations: do expert reviews work?', *IEEE Computer Graphics and Applications* **25**(5), 8–11.

Tufte, E. R. (2001), *The Visual Display of Quantitative Information*, 2nd edn, Graphics Press, Cheshire, CT, USA.

Unwin, A., Theus, M. & Hrdle, W. (2008), Exploratory graphics of a financial dataset, *in* C. Chen, W. Hrdle & A. Unwin, eds, 'Handbook of Data Visualization', Springer Handbooks of Computational Statistics, Springer, Berlin, Germany, pp. 831–852.

Ushahidi (2010), 'Haiti crisis map'. [Accessed 19 October 2011].
URL: http://haiti.ushahidi.com/

Ushahidi (2013*a*), *How To Map SMS Messages*. [Accessed 17 July 2013].
URL: http://community.ushahidi.com/uploads/documents/c_How-To-Approve-SMS.pdf

Ushahidi (2013*b*), 'Ushahidi: About us'. [Accessed 16 May 2013].
URL: http://www.ushahidi.com/about-us

van Wijk, J. J. (2006), 'Bridging the gaps', *IEEE Computer Graphics and Applications* **26**(6), 6–9.

Veregin, H. (1999), Data quality parameters, *in* P. A. Longley, M. F. Goodchild, D. J. Maguire & D. W. Rhind, eds, 'Geographical Information Systems, Principles and Technical Issues', 2nd edn, Vol. 1, John Wiley and Sons, Chichester, UK, pp. 177–189.

Viegas, F., Wattenberg, M., van Ham, F., Kriss, J. & McKeon, M. (2007), 'ManyEyes: a site for visualization at internet scale', *IEEE Transactions on Visualization and Computer Graphics* **13**(6), 1121–1128.

Vieweg, S., Hughes, A. L., Starbird, K. & Palen, L. (2010), Microblogging during two natural hazards events: what Twitter may contribute to situational awareness, *in* 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems', pp. 1079–1088. Atlanta, GA, USA, 10–15 April 2010.

Weaver, C. (2010), 'Cross-filtered views for multidimensional visual analysis', *IEEE Transactions on Visualization and Computer Graphics* **16**(2), 192–204.

Weaver, C., Fyfe, D., Robinson, A., Holdsworth, D., Peuquet, D. & MacEachren, A. M. (2007), 'Visual exploration and analysis of historic hotel visits', *Information Visualization* **6**(1), 89–103.

Wickham, H., Cook, D., Hofmann, H. & Buja, A. (2010), 'Graphical inference for infovis', *IEEE Transactions on Visualization and Computer Graphics* **16**(6), 973–979.

Wieczorek, J., Guo, Q. & Hijmans, R. (2004), 'The point-radius method for georeferencing locality descriptions and calculating associated uncertainty', *International Journal of Geographical Information Science* **18**(8), 745–767.

Wikipedia (2013), 'Wikipedia, the free encyclopaedia'. [Accessed 29 July 2013].
URL: http://en.wikipedia.org/

Wills, G. J. (1996), Selection: 524,288 ways to say "this is interesting", *in* 'Proceedings of the IEEE Symposium on Information Visualization (InfoVis 1996)', pp. 54–60. San Francisco, CA, USA, 28–29 October 1996.

Wood, J., Badawood, D., Dykes, J. & Slingsby, A. (2011), 'BallotMaps: Detecting name bias in alphabetically ordered ballot papers', *IEEE Transactions on Visualization and Computer Graphics* **17**(6), 2384–2391.

Wood, J. & Dillingham, I. (2012), 'giCentre GeoMap'. v1.1 [Accessed 29 April 2013].
URL: http://www.gicentre.org/geomap/

Wood, J. & Dykes, J. (2008), 'Spatially ordered treemaps', *IEEE Transactions on Visualization and Computer Graphics* **14**(6), 1348–1355.

Wood, J., Dykes, J., Slingsby, A. & Clarke, K. (2007), 'Interactive visual exploration of a large spatio-temporal dataset: Reflections on a geovisualization mashup', *IEEE Transactions on Visualization and Computer Graphics* **13**(6), 1176–1183.

Wood, J., Slingsby, A. & Dykes, J. (2013), 'giCentre Utilities'. v3.3 [Accessed 29 April 2013].
URL: http://www.gicentre.org/utils/

Yi, J. S., Kang, Y. a., Stasko, J. T. & Jacko, J. A. (2007), 'Toward a deeper understanding of the role of interaction in information visualization', *IEEE Transactions on Visualization and Computer Graphics* **13**(6), 1224–1231.

Zang, N., Rosson, M. B. & Nasser, V. (2008), Mashups: Who? What? Why?, *in* 'Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Extended Abstracts', pp. 3171–3176. Florence, Italy, 5–10 April 2008.

Ziemke, J. (2012), 'Crisis mapping: The construction of a new interdisciplinary field?', *Journal of Map and Geography Libraries* **8**(2), 101–117.

Zuk, T., Schlesier, L., Neumann, P., Hancock, M. S. & Carpendale, S. (2006), Heuristics for information visualization evaluation, *in* 'Proceedings of the 2006 workshop on Beyond Time and Errors: Novel Evaluation Methods for Information Visualization (BELIV 2006)', pp. 1–6. Venice, Italy, 23–26 May 2006.

Appendix A

Digital Appendix

The digital appendix can be found on the attached DVD. It contains the database schema and source code for the geoparsing system described in Chapter 4, and the binaries and source code for the four geovisualisation software prototypes described in Chapter 5.