Visualising Topical Sentiment in Twitter Streams

Suvodeep Mazumdar[1], Vitaveska Lanfranchi[1], Amparo E. Cano[2], and Fabio Ciravegna[1]

[1] OAK Group,
University of Sheffield, Sheffield, UK
n.surname@dcs.shef.ac.uk
[2] Knowledge Media Institute,
The Open University, Milton Keynes, UK
name.surname@open.ac.uk

Abstract. Advancements in mobile technology and the proliferation of social media platforms have made it possible for individuals to stay constantly connected with friends and family. This has provided new opportunities to the emergency response domain, where the information shared by individuals in crisis can provide invaluable insight into the situation on the ground. Information shared on social media is highly dynamic, heterogeneous, large scale, geographically distributed and multilingual. Moreover, the context of such information is mostly relevant for a very short period of time and the information can be very subjective, embedded in personal feelings. This is a significant challenge for the emergency response domain, where critical decisions need to be made quickly on the basis of the users' situation awareness. We propose to address this issue using visual analytics techniques to facilitate browsing and understanding of topicality and feelings in social media. Our approach is twofold- firstly, we enrich social media posts by adding semantics to facilitate browsing and sentiment in order to gauge the emotions behind individual posts. Secondly, we combine two paradigms of data browsing - topical and temporal into a real-time dynamic visualisation of social media messages.

Keywords: social media, dynamic visualisation, visual analytics

1 Introduction

Social networking and media-sharing platforms such as Twitter, Facebook and Flickr has seen a large scale adoption by individuals and communities across the web over the past few years. Cheaper mobile phones and internet-enabled devices have made it possible for individuals to stay connected with such platforms, thereby enabling people to share their experiences at all times. This amass of data continually generated has become a major source of information for organizations, user groups and communities to understand important situations and events. Such data streams are particularly useful in the case of emergencies and crises - where highly critical and potentially life-saving decisions need to be taken in a short amount of time. Social media data have become a natural port of call to achieve Situational Awareness, i.e. accurate, complete and real-time information about an incident [33], to understand the current local and global situation and how this may evolve over time [11]. In this sense, social networks have already changed the information landscape and have become a major source of information for authorities, organisations, individuals and groups [13]. They offer an excellent opportunity for situation awareness and crisis management; they often give a timely picture of events, allowing for both early warning of incidents as well as means for early situation awareness, even before police or rescue personnel arrive on the scene. Additionally, information from social news media does not only provide facts about the physical situation, but allows the possibility to assess the state of mind of people involved, e.g. positive and negative feelings, misconceptions etc. The immense volume of real-time, user-generated content collected from social media platforms or sources has already shown serious potential in applications such as disaster detection [23], seasonal mood level changes [8], tracking inuenza rates [14], etc. Authorities increasingly use social media for monitoring events and crowds as well, during social emergencies and natural disasters [27].

The goal of our research is to boost the users' understanding of interrelations between topicality and sentiments by the timely visual analysis of large-scale social streams (aggregated from Twitter, Facebook and Flickr). To this end, we have designed and developed *VisInfluence*, a new visualisation approach that leverages the semantic value of the information, to create contextual enrichment for visually browsing information. While effectively filtering dynamic social data is a significant challenge, our focus in this work is on exploring the sentiments associated with social media. Our present work is an extension of a previous work in visualising influence in social media, where we focussed on visualising influential users[9]. Time-critical domains such as emergency response require highly precise means to filter information. This needs understanding the topics of interest to the user. Since our starting point is the user's interest, we provide a mechanism for users to identify the topics interesting to them - by either selecting topics from a tag-cloud, or manually searching for topics. While there are topics that are generic, such as 'emergency', 'earthquake' or 'volcano', highly specific topics can also be relevant based on when the user is accessing the system, such as 'Download Festival', 'Sheffield Tramlines'. Evolving topics can also be of great interest to users such as 'London Riots' or 'Aurora shooting'. Once the topics of interest are identified, the system monitors the topics and provides real-time updates of the topics within a visualisation interface.

In the following sections we present related work focused on the use of social media data for situation awareness and on visual analytics in emergency response, This is followed by a description of our approach for obtaining and visualising the data, with a focus on the new visualisation approach we designed. Details about the implementation are then presented, followed by Conclusions and Future Work.

2 Related Work

In this section we present related work in two areas : Situational awareness and Visual Analytics in emergency response. Situational Awareness in times of emergency is paramount to deliver a timely and effective response [7]. To achieve effective Situational Awareness, emergency services must collate information from multiple sources and use it to build an understanding of the current situation and how this may evolve over time [11]. Leveraging data from citizens to build a form of collective intelligence [28], during emergencies or for security purposes, is becoming a vital resource for Situation Awareness [22]. During the 2007 southern California wildfires, two bulletin boards were set up to facilitate the exchange of information between citizens and authorities [25]. A later analysis of Twitter postings during the 2009 Red River flooding [29] indicated that the service was being used by citizens and communities to collate and propagate information in a concise and responsive manner. Several systems have been developed to support citizen participation during emergencies that either directly foster data from citizens through custom apps [20] or analyse public data stream to extract real-time knowledge [32, 2, 31]. Existing techniques for searching social media involve exploiting entity-based semantic features[30]; entity mentions, hashtags, URLs and metadata [18]; and entity annotations coupled with user models for personalised searches [1]. Recommendations and filtering systems are used to help users reduce information overload, i.e. recommending links that users may find interesting; using dynamic semantic models of user interests [2]; recommending posts and friends based on categories [10, 21].

Visual Analytics techniques have been proposed to represent and filter social media at different levels of specificity [16] [3] and to convey information evolution in the crisis management domain [24]. When visualising large scale social media data, visual analytics is mainly used to provide high level overviews. Lee [15] explores information regarding social media campaigns, Sakaki [24] uses twitter to understand the progression of earthquakes and Wongsuphasawat [34] explores trends in emergency medicine. While these systems manage to efficiently display the chosen information, they are limited in the amount of data displayed. Systems with a broader focus try to capture the properties of generic data, allowing users to filter the data to items of interest. TwitInfo [16] for example, uses multiple views to present a large data set and Eddi [6] allows users to explore real time data streams relating to a given keyword.

Whilst most social media visualisation approaches rely on geographical and temporal features TwitterReporter [19], some systems are starting to exploit the semantic of the data to enhance the visualisations. Examples of such systems are ThemeCrowds [5], TwitInfo [16], mediaWatch [12]. Marcus [16] uses features such as sentiment and link popularity to geographically plot the data. mediaWatch uses features such as sentiment to create news flow diagrams that analyses the evolution of keywords and sentiments over time with an innovative display. Adams et al. [4] also focuses on interactive colour-coded timeline displays. ThemeCrowds [5] cluster groups of users and their evolution over time for a particular topic.

3 Enriching social media data

The first step in our approach is enriching social media posts by adding semantics to facilitate browsing and sentiment to gauge the emotions behind individual posts. Social Media messages are typically composed of metadata (e.g. about the users, the device used to post, the location, etc.) and content. Both metadata and content can be analysed to extract information (e.g. keywords, terms, named entities, events, etc.) and to establish semantic data. In our implementation, we have used four processes to extract information from the messages - a tag processor to extract user-generated tags, a user detail extractor to extract metadata from the posts (such as user names, profiles and so on), an entity extractor and a sentiment analyser. The entity extractor and sentiment analyser are modules that use the Alchemy API[3] to extract different entities and sentiments of the messages. Our approach is based on two sub-steps: (1) Content analysis to identify related concepts that are later fed into the topical navigation widget; and (2) Sentiment analysis, to understand the mood of users and their influence on a topic. We illustrate this using the following example,

```
Alyssa Milano @Alyssa_Milano: I am so blown away by the police officers and all 1st
responders in Boston. Awesome bravery. I salute you! #BostonStrong
```

The user name (Alyssa Milano) and the user ID (@Alyssa_Milano) is initially stored along with the user-generated hashtag BostonStrong. The text of the tweet is then passed to the Alchemy API to extract relevant concepts (Tags: Awesome Bravery, police officers, responders) and sentiment (0.146867). We combine this information as shown below and store it within a local data store, to be queried later. The triples are formalised based on a visinfluence vocabulary as well as several other well-known vocabularies such as dcterms[4], sioc[5] and geonames[6]. While our visinfluence vocabulary is rather simplistic, our focus in this paper is on visual analytics instead of formalising the domain.

```
@prefix vin:  <http://dcs.shef.ac.uk/visinfluence/>
<http://dcs.shef.ac.uk/visinfluence/status/325418041229848577>
rdf:type sioc:Post ;
sioc:hasCreator <http://twitter.com/Alyssa_Milano> ;
dcterms:created "2013-4-20 02:18:01" ;
dcterms:subject "BostonStrong" ;
vin:City <vin:Boston> ;
vin:Tag "Awesome Bravery" ;
vin:Tag "police officers" ;
vin:hasSentiment 0.146867 ;
```

When all the tweets are processed, the system can search for tweets that are relevant to topics as well as created within specific temporal durations. The retrieved set of tweets is then ‘bundled’ into one single representation in the visualisation. In our current implementation, the sentiment score for a set of

[3] http://www.alchemyapi.com/
[4] http://dublincore.org/documents/dcmi-terms/
[5] http://sioc-project.org/ontology
[6] http://www.geonames.org/ontology/documentation.html

tweets is calculated as the mean of their individual scores. In the near future, we will investigate other ways of generating combined sentiment scores for a group of tweets.

4 A Visual Approach

We build our approach on the classical visual information - seeking mantra proposed by Shneiderman - Overview first, zoom and filter, then details-on-demand [26], where a visual approach toward finding information lies in presenting an overview of the data, providing means for users to filter information they are interested in, and then making the details available when required. While this approach has been embedded in most visualisations and interfaces, we believe it is most effective for large scale static data. However, in our scenario, the data is highly dynamic, and this results in situations where such an approach would not be effective. For example, a visual overview of highly dynamic data would be computationally expensive, as well as introduce a cognitive burden on the users. We propose to extend Shneiderman's approach by incorporating means to exploit dynamic social data. Our approach is in four steps: **identify topics**, **monitor**, **explore**, **details-on-demand**. We explain each of these steps using a scenario, where an emergency response personnel in the control room uses VisInfluence to understand what is currently happening.

4.1 Identify Topics

Being a time-critical application domain, users need to be able to quickly communicate what they are interested in. This calls for easy means to identify their topics of interest, and then subsequently add them to their list of topics to monitor. In our present implementation, topics can be of two types - user-generated hashtags or entities (extracted from the text). Upon initialisation, users are presented with a tag cloud that represents a summary of all the hashtags that are posted from the time the data is collected. The tag cloud is also presented with a drop-down menu showing different types of entities that have been identified while processing the posts. Additionally, the system provides users with a topical search facility to add their own search terms. In the next stages, the tweets relevant to the identified topics will only be considered, as shown in Figure 1. The final set of topics are shown to the user as a list, which can be edited as desired before proceeding to the next step of exploration. This stage is a notable change from the original metaphor of visual information seeking proposed by Shneiderman, where the author proposed that the first step would be to provide an overview of everything, and then provide filtering operations. We argue that in order to make quick and important decisions, users need to first identify the topics that are most relevant to them. This saves critical time and can help users guage situations that are of utmost importance to them.

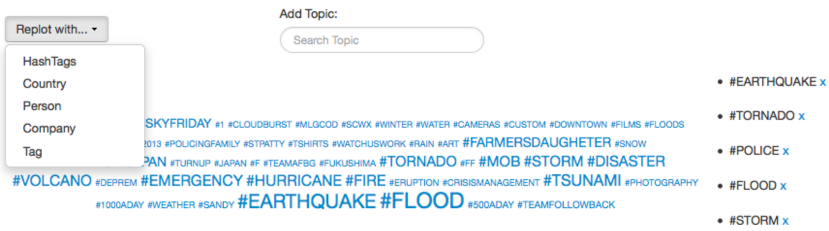

Fig. 1. A screenshot of topic selector - the user can either enter their own search topics, or select from the tagclouds generated from the hashtags or extracted entities. The topics selected are shown on the list (right).

4.2 Monitor

The set of topics identified by the user are then used by VisInfluence to render the final visualisations. The visualisation space consists of individual horizontal panels for each topic, and each panel presents the dynamic information collected about the topic. The panels are divided into several vertical time frames (600 time frames in our present implementation), each time frame indicates a duration of time. Figure 2 shows one such example, where each time frame represents 10 seconds - this means that every 10 second, a new reading will be requested from the backend, where the tweets collected over the last 10 seconds will be processed to calculate a combined sentiment score.

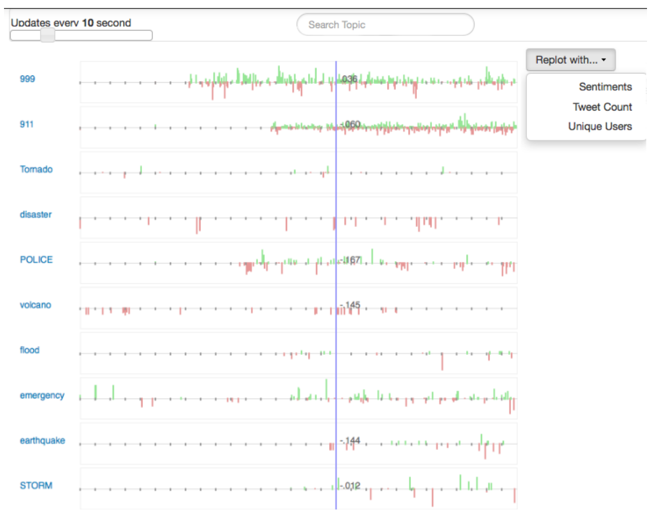

Fig. 2. A screenshot of the VisInfluence application showing how the sentiments of all the tweets associated with a topic at a given time are visualised

The backend extractors harvest tweets related to the individual topics, based on its own search criteria. The visualisation interface then queries the backend for the latest readings of the collective sentiments for all the topics, which are then rendered as a bar on the panel. The bar is color-coded to indicate a positive or a negative sentiment (positive displayed as a green bar, while negative as a red). The length of the bar indicates the extent of the general positive or negative feeling associated with the topic at the given time frame.

Users can control how often the visualisations are updated using a slider on the top of the interface, thereby choosing how quickly new data is visualised in the system. When a new time frame is initiated, all the readings for the previous timeframes are shifted toward the left, and the new readings are plotted on the far right. The oldest reading from the left most timeframe is deleted from the system. This enables users to have a historical perspective on the sentiments associated with the topic and follow how it develops as a situation evolves - however, ensuring that the users donot get overloaded by the amount of information that is presented to them.

The system also allows users to add new topics of interest at run-time, while the visualisations are being updated. The users can use the search box on the top right side of the interface (Figure 2) to enter terms they are interested in, and the topics get added to their list of topics to monitor. This is to accommodate for new topics of interest that they observe from their monitoring activities. The new topics get stacked on top of the previous topic panels - e.g. “999” and “911” panels are added by the user after they observe that these topics can be relevant as shown in Figure 2.

4.3 Explore

The users interact with the interface to explore the data to a greater extent. For example, hovering over the panels show a scale that encapsulate all the other panels, and show the readings for each panel at that time frame. In addition to the visual indication of the computed sentiments being rendered as charts, the scale shows the real values of the sentiments to enable a more fine-grained comparison. Users can also visualise the number of tweets posted as well as the number of unique users posting tweets that were relevant to the topic at the specific time frame. Users select from a drop-down menu the data they would like to visualise and the queries to the back end extractors would then be modified based on the facet of the data they would like to visualise. The tweet count and unique user count is presented in the same interface, similar to the sentiment plots where bars represent the number of tweets or users. Unlike sentiment plots, however, these plots are not color-coded, and merely represent how many tweets have contributed toward the final sentiment scores. Figure 3 shows an example where the tweet count has been selected as the facet to be visualised for the topic 911 - the figure shows the rising amount of the tweets over the latter half of the exploration session. The number of tweets had suddenly spiked, and dropped at the end of the session.

Fig. 3. Visualising tweet counts for different time frames.

We believe that this step is similar to the ‘overview’ step in Shneiderman’s information seeking paradigm, since this allows the users to have a holistic view of the data. However, our extension for this step is to facilitate an exploratory paradigm where the user can be provided with means to decide which facet of the data is of interest to them.

4.4 Details-On-Demand

The final step of our approach is similar to Shneiderman’s - where we provide means to reach the details of the data. It is extremely important to make this as seamless as possible, and make it easy for the user to understand the data being presented. Users can access details during their exploratory session by hovering their mouse (and consequently, the scale) to view the scores at the specific time frame and clicking on the plot area. The exact position of the mouse click is then calculated to identify the time frame and the topic that the user is interested in. At this stage, we provide users with views over multiple aspects of the data. The most important facet that users are presented with are the individual instances of the tweets. This assists the analyst in understanding why the sentiments around a given topic at one particular time frame generates its scores. The tweets are ranked on the basis of their sentiments - users can chose to view the tweets with the most negative scores on top, and the most positive scores at the bottom and vice versa. The second facet is the list of most influential users for the topic at the chosen time frame. The list of users is also ranked on the basis of their contribution toward the sentiment score. The user with most tweets that are negative with the greatest margin is placed on top of the list, while the user with the most positive tweets are placed at the bottom of the list.

5 Implementation and Architecture

We discuss the implementation of VisInfluence in two phases - data processing and visual representation. The first step requires streaming Tweets using Twitter API - this process introduces a few challenges - searching the public timeline for tweets would result in large number of tweets that are irrelevant. Hence, as a starting point, VisInfluence is fed a few search terms to gather an initial set of data. This set of search terms is directly related to the domain - for example, a user monitoring a music festival would prefer using terms relevant to the festival, while an emergency personnel may prefer using generic terms such as ‘emergency’, ‘disaster’ and so on. This initial set of tweets collected every 15 minutes are then fed into four extractors - a user detail extractor, tag processor, entity extractor and sentiment extractor. The entity extractor identifies the various types of entities found in the tweet - these entities can be of several types such as persons, countries, cities, tags and so on. The sentiment analyser computes the sentiments associated with each tweet. The entity extractor and sentiment analyser are essentially scripts that invoke the language processing service Alchemy API[7]. The responses from the Alchemy API are parsed and stored locally, available to be queried at later stages.

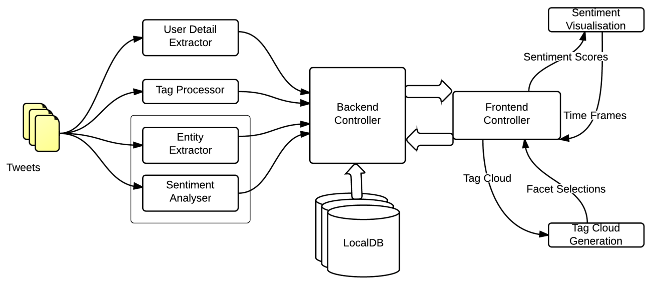

Fig. 4. VisInfluence Architecture

The other two modules, Tag Processor and User Detail Extractor extracts the hashtags and relevant user information such as user IDs, user names and locations. The extracted information is then stored in a local triplestore, via a backend controller. The user-facing part of the system, or the front end constitutes of a visualisation interface built on Processing.js[8]. The processing modules are displayed on an HTML page, along with several other Javascript-based tools such as JQueryTagCloud[9] for rendering the tag cloud and jQuery to handle AJAX queries to the backend. Each interaction with the user results in queries that are then sent to the backend controller. The queries are then forwarded to the triplestore, and the results retrieved are returned to the relevant callback functions. The responses are then parsed using Javascript, and rendered accordingly.

6 Challenges and Conclusions

The development of the VisInfluence system has followed an iterative user-centered design process, where users have been presented with incremental versions of different features. Another recently developed system, the Context and

[7] http://www.alchemyapi.com/
[8] http://processingjs.org/
[9] http://www.jondev.net/articles/jQuery_Tag_Cloud_Example

Hierarchy chain [17] has also been discussed in several focus group sessions. These discussions and focus groups have been crucial in understanding user needs and the challenges posed by their tasks and the domain itself. In the next few months, we plan to conduct an evaluation in a task-based simulated emergency aimed at understanding the applicability of our proposed paradigm for exploring social data. The participants for the evaluation will be primarily emergency response practitioners and researchers.

One of the primary challenges that the system faces is limited computational resources. Processing thousands or millions of tweets at near real-time is highly challenging and computationally expensive. This results in processing delays, that can introduce a lag in the system. While such delays are expected, a time-critical domain such as emergency response requires careful consideration of how significant and what is the possible impact of such delays. Indeed, our approach of filter-first can reduce unnecessary and irrelevant information to be discarded, a helpful addition would be efficient spam-detection techniques. Identifying spammers and spam content before as a pre-processing step would significantly reduce the amount of information to be processed, thereby reducing the possible delays.

Another set of issues associated with the system is the limits introduced by services - Twitter and Alchemy are both services that impose certain restrictions in terms of number of API calls. In our implementation, we have had to introduce delays to ensure that calls to such services respected these restrictions. These delays have also added to the overall delay in the system. However, such issues can be mostly addressed by paid or premium accounts with such services. In the near future we plan two major modifications in the system - improve the backend processing and provide a way to improve the filtering mechanisms. With these final set of updates, we plan to evaluate the system with real users from the emergency response domain. We aim to validate our four-step approach and understand how can such technologies be used to compliment more traditional techniques.

References

1. Fabian Abel, Ilknur Celik, Geert-Jan Houben, and Patrick Siehndel. Leveraging the semantics of tweets for adaptive faceted search on twitter. *The Semantic Web–ISWC 2011*, pages 1–17, 2011.
2. Fabian Abel, Qi Gao, Geert-Jan Houben, and Ke Tao. Analyzing Temporal Dynamics in Twitter Profiles for Personalized Recommendations in the Social Web. In *Proceedings of ACM WebSci '11, 3rd International Conference on Web Science, Koblenz, Germany*. ACM, June 2011.
3. Aniket Mulye Abhishek Pandya and Soon Tee Teoh. Enhancing online timeline visualizations with events and images. In *Electronic Imaging, Visualization and Data Analysis (VDA)*, 2011.
4. Brett Adams, Dinh Phung, and Svetha Venkatesh. Eventscapes: visualizing events over time with emotive facets. In *Proceedings of the 19th ACM international conference on Multimedia*, MM '11, pages 1477–1480, New York, NY, USA, 2011. ACM.

5. Daniel Archambault, Derek Greene, Pádraig Cunningham, and Neil Hurley. Themecrowds: multiresolution summaries of twitter usage. In *Proceedings of the 3rd international workshop on Search and mining user-generated contents*, SMUC '11, pages 77–84, New York, NY, USA, 2011. ACM.
6. Michael S. Bernstein, Bongwon Suh, Lichan Hong, Jilin Chen, Sanjay Kairam, and Ed H. Chi. Eddi: interactive topic-based browsing of social status streams. In *Proceedings of the 23nd annual ACM symposium on User interface software and technology*, UIST '10, pages 303–312, New York, NY, USA, 2010. ACM.
7. Ann Blandford and B.L. William Wong. Situation awareness in emergency medical dispatch. *International Journal of Human-Computer Studies*, 61(4):421 – 452, 2004.
8. Johan Bollen, Huina Mao, and Xiao-Jun Zeng. Twitter mood predicts the stock market. *CoRR*, abs/1010.3003, 2010.
9. A. E Cano, S Mazumdar, and F Ciravegna. Social influence analysis in microblogging platforms - a topic-sensitive based approach. *To appear in Special Issue on The Semantics of Microposts, Semantic web : interoperability, usability, applicability.*
10. Jilin Chen, Rowan Nairn, and Ed Chi. Speak little and well: recommending conversations in online social streams. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, pages 217–226. ACM, 2011.
11. Mica R. Endsley. Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64, 1995.
12. Alexander Hubmann-Haidvogel, Adrian M. P. Brasoveanu, Arno Scharl, Marta Sabou, and Stefan Gindl. Visualizing contextual and dynamic features of micropost streams. In *2nd Workshop on Making Sense of Microposts (MSM-2012*, 2012.
13. Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. What is twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 591–600, New York, NY, USA, 2010. ACM.
14. Vasileios Lampos, Tijl De Bie, and Nello Cristianini. Flu detector: tracking epidemics on twitter. In *Proceedings of the 2010 European conference on Machine learning and knowledge discovery in databases: Part III*, ECML PKDD'10, pages 599–602, Berlin, Heidelberg, 2010. Springer-Verlag.
15. Kyumin Lee, James Caverlee, Zhiyuan Cheng, and Daniel Z. Sui. Content-driven detection of campaigns in social media. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, CIKM '11, pages 551–556, New York, NY, USA, 2011. ACM.
16. Adam Marcus, Michael S. Bernstein, Osama Badar, David R. Karger, Samuel Madden, and Robert C. Miller. Twitinfo: aggregating and visualizing microblogs for event exploration. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 227–236, New York, NY, USA, 2011. ACM.
17. S Mazumdar, F Ciravegna, A L Gentile, and V Lanfranchi. Visualising context and hierarchy in social media. In *International Workshop on Intelligent Exploration of Semantic Data (IESD'2012) at EKAW 2012 2012,Galway City,Ireland. October 9 2012*, 2012.
18. Pablo N Mendes, Alexandre Passant, and Pavan Kapanipathi. Twarql: tapping into the wisdom of the crowd. In *Proceedings of the 6th International Conference on Semantic Systems*, page 45. ACM, 2010.
19. Brett Meyer, Kevin Bryan, Yamara Santos, and Beomjin Kim. Twitterreporter: Breaking news detection and visualization through the geo-tagged twitter network. In Wei Li, editor, *CATA*, pages 84–89. ISCA, 2011.

20. Ory Okolloh. Ushahidi, or'testimony': Web 2.0 tools for crowdsourcing crisis information. *Participatory Learning and Action*, 59(1):65–70, 2009.
21. Alexander Pak and Patrick Paroubek. Twitter as a corpus for sentiment analysis and opinion mining. In *Proceedings of LREC*, volume 2010, 2010.
22. Leysia Palen and Sophia B. Liu. Citizen communications in crisis: anticipating a future of ict-supported public participation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 727–736, New York, NY, USA, 2007. ACM.
23. Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 851–860, New York, NY, USA, 2010. ACM.
24. Takeshi Sakaki, Makoto Okazaki, and Yutaka Matsuo. Earthquake shakes twitter users: real-time event detection by social sensors. In *Proceedings of the 19th international conference on World wide web*, WWW '10, pages 851–860, New York, NY, USA, 2010. ACM.
25. Irina Shklovski, Leysia Palen, and Jeannette Sutton. Finding community through information and communication technology in disaster response. In *Proceedings of the 2008 ACM conference on Computer supported cooperative work*, CSCW '08, pages 127–136, New York, NY, USA, 2008. ACM.
26. Ben Shneiderman. The eyes have it: A task by data type taxonomy for information visualizations. In *Proceedings of the 1996 IEEE Symposium on Visual Languages*, VL '96, pages 336–, Washington, DC, USA, 1996. IEEE Computer Society.
27. C.L. Slagh. *Managing Chaos, 140 Characters at a Time: How the Usage of Social Media in the 2010 Haiti Crisis Enhanced Disaster Relief*. BiblioBazaar, 2011.
28. Vassilios Solachidis, Phivos Mylonas, Andreas Geyer-schulz, Bettina Hoser, Sam Chapman, Fabio Ciravegna, Vita Lanfranchi, Ansgar Scherp, Costis Contopoulos, Ioanna Gkika, and et al. Advances in data analysis, data handling and business intelligence. *Knowledge Creation Diffusion Utilization*, pages 765–774, 2010.
29. Kate Starbird, Leysia Palen, Amanda L. Hughes, and Sarah Vieweg. Chatter on the red: what hazards threat reveals about the social life of microblogged information. In Kori Inkpen Quinn, Carl Gutwin, and John C. Tang, editors, *CSCW*, pages 241–250. ACM, 2010.
30. Ke Tao, Fabian Abel, Claudia Hauff, and Geert-Jan Houben. What makes a tweet relevant for a topic? *Making Sense of Microposts (# MSM2012)*, page 9, 2012.
31. Teun Terpstra, R Stronkman, A de Vries, and GL Paradies. Towards a realtime twitter analysis during crises for operational crisis management. In *Proceedings of the 9th International ISCRAM Conference. Vancouver, Canada*, 2012.
32. Lanfranchi V. Realtime situation awareness via social media streams. In *Demo Session at ISCRAM 2012, Vancouver, Canada*.
33. L. Winerman. Social networking: Crisis communication. *Nature*, 457(7228):376–8, 2009.
34. Krist Wongsuphasawat, John Alexis Guerra Gómez, Catherine Plaisant, Taowei David Wang, Meirav Taieb-Maimon, and Ben Shneiderman. Lifeflow: visualizing an overview of event sequences. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 1747–1756, New York, NY, USA, 2011. ACM.