A Case Study for Monitoring Fires with Twitter

Robert Power
CSIRO, Australia
robert.power@csiro.au

Bella Robinson
CSIRO, Australia
bella.robinson@csiro.au

John Colton
CSIRO, Australia
john.colton@csiro.au

Mark Cameron
CSIRO, Australia
mark.cameron@csiro.au

ABSTRACT

This paper presents a user configurable monitoring system to track in near-real-time tweets describing fire events. The system targets fire related words in a user defined region of interest published on Twitter which are further processed by a text classifier to determine if they describe a known fire event of interest. The system was motivated from a case study that examined a corpus of tweets posted during active bushfires. This demonstrated that useful information is available on Twitter about fire events from people who are in the vicinity.

We present an overview of the system describing how it is initially configured by a user to focus on specific fire events in Australia, the development of a text classifier to identify tweets of interest, especially those with accompanying photos, and the monitoring system that can track multiple events at once.

Keywords

Situational Awareness, Disaster Management, Social Media, Twitter.

INTRODUCTION

The management of bushfires in Australia is the responsibility of State and Territory governments with agencies such as the New South Wales Rural Fire Service (NSW RFS) being responsible for managing fire events. They inform the community about known incidents, providing updates of progress on their web sites and in some cases through social media, such as Twitter and Facebook.

The NSW RFS is the world's largest volunteer fire service, consisting of over 2,000 rural fire brigades managing 126 fire districts coordinated from over 50 offices with the central headquarters located in Sydney. They have a total volunteer membership of approximately 74,000 and an extra 800 staff to manage the service operations. During the 2013/14 fire season, they attended over 23,000 incidents, of which more than 8,000 were bush or grass fires[1].

We undertook a requirements gathering exercise with NSW RFS to better understand how they envisage using social media to gather evidence of fires from the community. This was during the period September 2013 until February 2014 when incident controllers were exploring content published on Twitter. The most significant requirements found related to finer-grained user configurable monitoring and in summary were: enabling greater situational awareness at an incident level by being able to configure a tool to focus on multiple fire events and from then on readily monitoring the separate events; and finding 'high value' images such as those with smoke plumes.

We present a case study to explore these requirements, our user configurable fire monitoring tool that can track multiple fire events and plans for future work.

[1] http://www.rfs.nsw.gov.au/about-us/fast-facts

RELATED WORK

Social media has been recognized as an effective communication channel. However, its use as a source of situational awareness from the public has not been widely adopted (Anderson 2012). This is due to a combination of the difficulty in appropriately 'framing' social media content, sifting through the large volumes of information available and the issue of trusting content (Lindsay 2011).

Previous work has focused on disaster event detection using social media with particular success for earthquakes (Sakaki, Okazaki and Matsuo 2013; Robinson, Power and Cameron 2013; Avvenuti, Cresci, Marchetti, Meletti and Tesconi 2014). This is due to the unexpected nature of earthquakes and the characteristics of the tweets that follow from people who experience them. People are quick to report their experience as short messages, which include exclamation marks or other such punctuation, swear words and do not include URLs, user mentions or retweets (Avvenuti et al 2014). After an earthquake event has been detected, the Earthquake Alert and Report System (EARS) system processes the tweets that follow to gather information about impacts in the affected area with the aim of producing a damage assessment report. They also include measures to deter and minimize the effect of rumors or misinformation from malicious users.

A similar methodology has been applied to detect fires (Power, Robinson and Ratcliffe 2013). The authors describe a notification system to identify in near-real-time tweets describing fire events. Tweets including fire related keywords are first filtered to identify candidate messages. These tweets are then processed by a text classifier to refine the results to target actual fire events. Their system detects new events after a 'quiet' period of inactive fire related discussion.

Other tools exist to detect and monitor different disaster events or crisis management issues. Twitcident (Terpstra, de Vries, Paradies and Stronkman 2012) targets specific event types such as natural disasters or gatherings of people at riots or organized celebrations. Their tool can be customized to specific locations and incident types using message content and tweet type filters. Similarly, Tweet4act (Chowdhury, Imran, Asghar, Amer-Yahia and Castillo 2013) uses keywords from an incident specific dictionary to target tweets related to a crisis situation with text classifiers used to categorize them into groups of before, during and after the event. Other research (Imran, Elbassuoni, Castillo, Diaz and Meier 2013; Traverso, Cerutti, Stock and Jackson 2014) makes use of ontologies combined with Natural Language Processing (NLP) and machine learning techniques to categorize tweets of interest contributing to situational awareness. A similar approach (Schulz, Ristoski and Paulheim 2013) is used for real-time identification of small-scale incidents using machine learning for text classification augmented by semantic enrichment of microblog content using Linked Open Data. Their approach has been applied to detect three categories of small-scale incidents: car accidents, urban fires and shootings.

A recent semi-structured interview survey of barriers to the use of social media by U.S. public sector emergency managers (Hiltz, Kushma and Plotnick 2014) identified: limited personal time to use social media; lack of organizational policies and guidelines for use; and concern over trustworthiness of data.

Extracting location information from tweets is another active research area. This is important during emergency events where people often include location information. The OzCT geo-tagger (Ghahremanlou, Sherchan and Thom 2014) applies toponym resolution and recognition to detect both definite and ambiguous locations. Han, Cook and Baldwin (2014) have developed an integrated geo-location prediction framework to obtain location indicative words from tweets that may be used to infer the Twitter user's location. A similar system combines NLP, heuristics and Named Entity Recognition (NER) techniques to produce a geo-parser that is tailored for Twitter messages (Gelernter and Balaji 2013).

The benefit of our tool is that it utilizes this previous work as a pair of related manager and monitor web applications that separate the task of defining an incident from tracking it. The manager is used to define events of interest, which can be simultaneously observed using the monitor to highlight fire related tweets, identify locations using NER techniques and present images, topic clusters and a tweet volume timeline graph.

THE PROBLEM AND SUPPORTING CASE STUDY

The problem is summarized as follows:

1. Is there information on social media about bushfires described by people

experiencing them?

2. Is this information useful for incident controllers and fire responders?
3. How can this information reliably be identified and obtained from social media in near-real-time?

Answering these questions for a specific fire event in Australia defines our case study. The fire event investigated occurred in the Blue Mountains region of NSW in October 2013. This is shown in Figure 1 as the central fire icon with the subscript '4' indicating there are four active fires in close proximity.

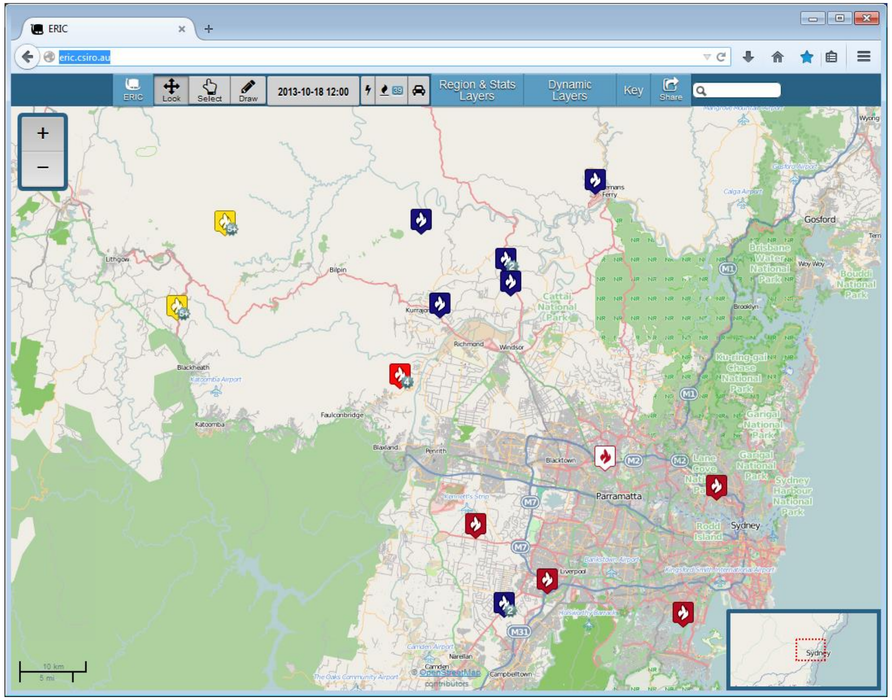

Figure 1: Fires reported by NSW RFS at noon on 18 October 2013.

GATHERING EVIDENCE

Our investigation proceeded as follows. First, we obtained all the tweets published in NSW during October 2013 and examined this content to determine if there was useful information relating to the fires in the Blue Mountains region. This included searching for tweets that contained fire related keywords ('fire' or 'smoke'); focusing on a specific geographic region of interest; manually inspecting the content found and examining photos published at this time to find high value images of the fires.

A summary of the tweets obtained from NSW from October 2013 is shown in Table 1. The elements of the tweet that are of interest are: the use of the words 'fire' or 'smoke' in the tweet text, if it is geo-tagged and if there is an associated photo. We are also interested in the tweets from users whose profile location is near the Blue Mountains region (the second column in Table 1); geo-tagged tweets in this region of interest (the third column) and both of these location tests combined (the last column).

Geo-coding the profile location was done using the Yahoo! GeoPlanet API[2] by checking the returned coordinates to be within the regions of interest (Power, Robinson, Colton and Cameron 2014). This service is rate limited, so we cache the results. Table 1 shows that for the Blue Mountains region there are a sufficient sample of tweets (66,042) of which 4% mention 'fire' or 'smoke' and are not retweets (2,935), and a similar number, again 4%, include a photo (2,961).

The potentially 'high value' tweets are those that are geo-coded in the Blue Mountains region, contain the text of interest, include a photo and are not retweets. There are 149 of these and reviewing them reveals details about the fire, notably, its intensity, location and behavior. An example is shown in Figure 2.

This image and others like it note the name of the fire and include further place names, such as towns and streets. Using the geo-coded coordinates from the tweet along with the name of the fire described, the direction the fire is moving can be inferred from the direction of the smoke. An incident controller with fire fighting experience can also infer some measure of the intensity of the fire from the colour,

[2] https://developer.yahoo.com/geo/geoplanet/

height and volume of the smoke plume.

This case study demonstrates that information is published by the general public on Twitter about fire events from people who are in the vicinity. The remaining issue to be investigated is how to readily identify this information in near-real-time. Previous research (Power et al. 2013) has found that a text classifier can be used to identify tweets from people describing fire events.

	NSW	Location	Geo-tagged	Comb
Total number of tweets	10 452 415	59 999	7 667	66 042
Tweets mentioning ‘fire’	157 613	3 625	454	3 988
Tweets mentioning ‘smoke’	18 024	363	63	412
Tweets ‘fire’ or ‘smoke’	167 651	3 853	487	4 239
Tweets ‘fire’ or ‘smoke’ (ex RT)	80 881	2 550	486	2 935
Tweets with photos	1 580 998	2 940	835	3 669
Tweets with photos (ex RT)	294 456	2 234	833	2 961
Tweets ‘fire’ or ‘smoke’ & photos	27 733	439	150	572
‘fire’ or ‘smoke’ & photos (ex RT)	7 404	174	149	306
Geo-tagged tweets	278 093	2 182	7 667	8 225
Geo-tagged tweets (ex RT)	277 202	2 177	7 659	8 214
Geo-tagged & photos (ex RT)	22 728	135	835	864
Geo, ‘fire’/‘smoke’, photos (ex RT)	839	20	149	152

Table 1: Tweet summary from NSW and the Blue Mountains during October 2013.

A Support Vector Machine (SVM) (Joachims 1998) text classifier configured using a linear kernel was developed and a comprehensive investigation of the best features to use was explored using a 10-fold cross-validation procedure. The best results achieved had an F1 score of 0.831 and accuracy of 84.54%. This text classifier was used to process the NSW October 2013 (non-retweet) tweets containing ‘fire’ or ‘smoke’ (80,881) with 37,635 (46%) being classified as fire related. The accuracy of this has not yet been examined.

DEVELOPING THE FIRE MONITOR

A fire monitor map based interactive web site was developed based on the findings above: the requirements gathering process, the Blue Mountains fire case study; methods of identifying a user’s location; and the text classifier. This is a working demonstrator for use by the NSW RFS over the 2014/15 fire season, with the aim of progressing discussions with incident controllers and fire responders and to demonstrate that useful content is available on Twitter.

Figure 2: Example tweet with a geo-coded image[3].

The monitor is first configured using the ‘Incident Manager’ interface shown in Figure 3. This allows the user to define the criteria to target a specific fire event of interest to identify the latest tweets to help inform situational awareness. These

[3] https://twitter.com/gadget23/statuses/391408065980796929

settings can be saved so that the user can select from a list of preconfigured events and switch between them to quickly review the current situation.

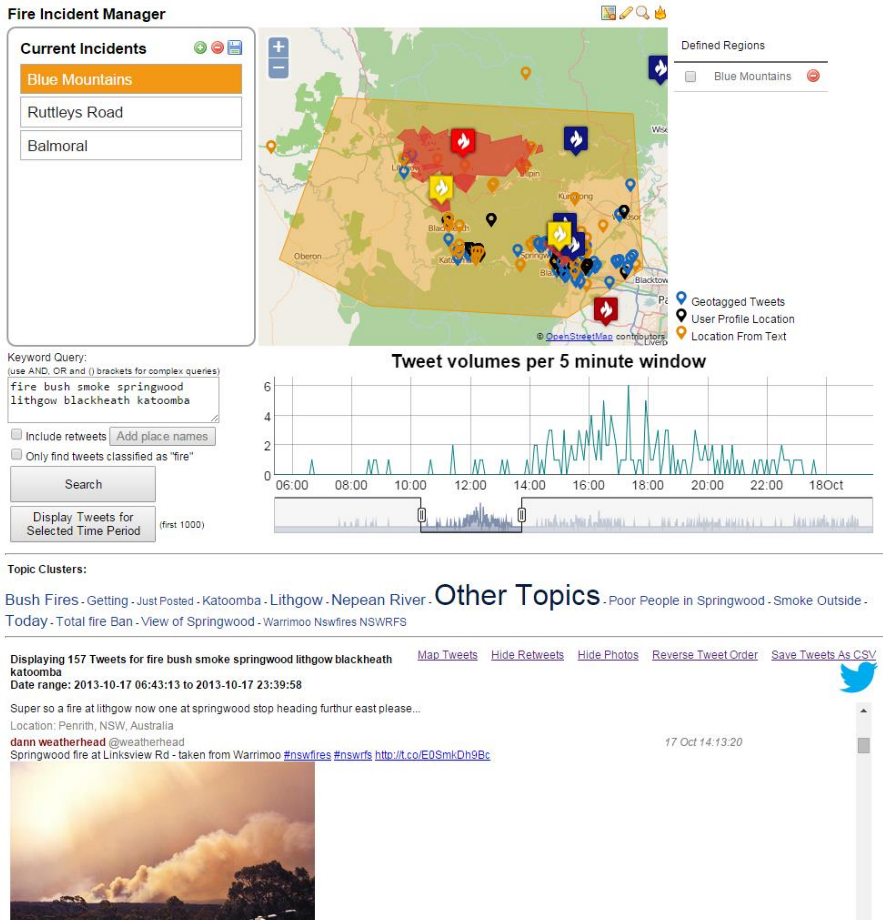

Figure 3: The Incident Manager

Initially there are no existing incidents and the map is blank. A region of interest is defined by either panning and zooming the map or by searching for a place name to zoom to, implemented using the Open Street Map Nominatim service[4]. The status of known fire events currently underway is shown on the map to help orientate the user. A 'New Incident' is then created to initiate a new incident to be monitored, appearing as a new entry in the *Current Incidents* section. The bounds of the map are used by default as a search constraint to locate tweets although a custom region can be defined by drawing a polygon.

Place names can also be defined as keywords to search for, with a list of candidate names provided by referencing the Australian gazetteer. Place names are used as search terms since people often include a location in the tweet text when reporting fire incidents. Other keywords to search for can also be entered.

The configuration information provided, the location and search keywords can be used to search for tweets. The resulting timeline chart gives an overview of the tweet volume and can be zoomed to a time interval of interest using the slider handles at the sides and corresponding tweets displayed. These are processed using the Carrot2 clustering engine[5] to provide a summary of topics discussed. The locations of users are shown as markers on the map using different colours for geo-tagged tweets, the user's profile location and locations mentioned in the tweet message. These markers can be selected to show the tweet text and photos if present. The user can also customize the content displayed by including retweets and only showing tweets considered to be fire related by the text classifier.

The example in Figure 3 shows the result of this process for the fires in the Blue Mountains occurring on 17 October 2013. The incident defined by the user has been called 'Blue Mountains' and a polygon used to restrict the region of interest. A list of keyword terms for tweet search has been defined; the first three, *fire*, *bush* and *smoke* are fire related with the remaining four being town names in the region of interest. The timeline chart has been zoomed with the topic clusters generated from the matching 157 tweets shown and the tweets are available for review. Note there is a photo showing the smoke from the Springwood fire.

---

[4] http://wiki.openstreetmap.org/wiki/Nominatim
[5] http://project.carrot2.org/

Three incidents have been defined in this way, the user provided names of which are listed in the *Current Incidents* section. Once configured, the ‘Incident Monitor’ interface, shown in Figure 4, is used to provide an overview of the tweets corresponding to these incidents. This screenshot shows three tweet streams for these events and on the left are more details about the ‘Blue Mountains’ incident. This is the ‘active’ incident on the monitor page, indicated by the selecting this tweet stream column. Further summarizing information about this active event is shown on the left: the map, the timeline chart, cluster summary, and the tweets contributing to a user selected cluster topic.

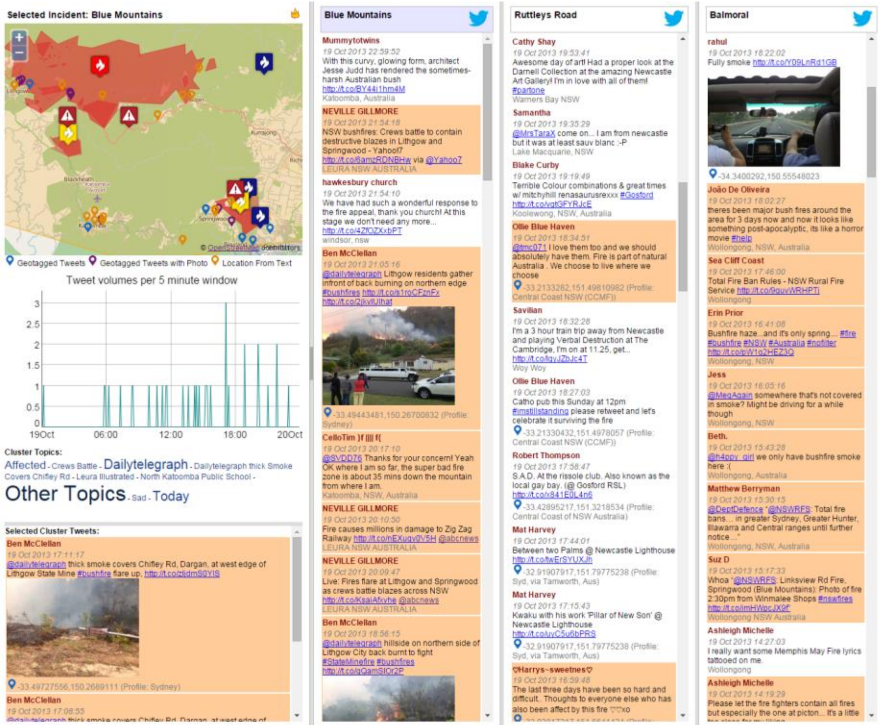

Figure 4: The Incident Monitor.

CONCLUSIONS AND FUTURE WORK

We have undertaken a requirements gathering process with NSW RFS incident controllers to discover if there is content of interest from Twitter. A case study was undertaken to explore a collection of Australian tweets during a period of severe active fires. We found there is useful information published by the general public on Twitter in the vicinity of fire events. We have used a text classifier to identify tweets of interest and specifically those that have associated images.

We have developed a fire monitoring web site that consists of an Incident Manager interface to allow a user to configure the tool to focus on an area of interest and to provide further filter conditions. These configuration details are saved and a companion Incident Monitor interface allows a user to simply monitor the previously configured and saved events of interest. This monitor interface purposefully has limited user interaction features since it is targeted for fire responders in the field or incident controllers during emergency events who are interested to obtain up to date near-real-time situational awareness from the community, but don’t have the time to explore and manipulate a user interface.

The key features of the tool are the ability to define up to four incidents to simultaneously monitor; easily switching between incidents; highlighting messages classified as fire related; showing images; integrating authoritative emergency warnings; showing tweet locations on a map; tweet summaries using a tweet timeline graph and cluster topics; and the separation of defining an incident from monitoring it.

This tool is also useable for other event types, such as earthquakes, cyclone tracking, flood events and crisis management incidents. This will be verified by exploring case studies from historical examples and testing on current tweets.

There are a number of other areas of further research work. We plan to review the performance of the text classifier by extending the features tested to include new measures using NLP techniques such as Word Sense Disambiguation and Part of Speech tagging. Twitter specific NLP packages may also improve performance. We have developed an initial image classifier to automatically identify images that are considered to be of smoke or fire however this is still a work in progress and requires further development to improve the results obtained to date.

REFERENCES

1. Anderson, M. (2012). Integrating social media into traditional management command and control structures: the square peg into the round hole. In Australian and New Zealand Disaster and Emergency Management Conference, Brisbane Australia, May 2012, pages 18–34.
2. Avvenuti, M., Cresci, S., Marchetti, A., Meletti, C. and Tesconi, M. (2014). EARS (earthquake alert and report system): a real time decision support system for earthquake crisis management. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '14). ACM, New York, NY, USA, pp 1749-1758.
3. Chowdhury, S.R., Imran, M., Asghar, M.R., Amer-Yahia, S. and Castillo, C. (2013). Tweet4act: Using incident-specific profiles for classifying crisis-related messages. In The 10th International Conference on Information Systems for Crisis Response and Management (ISCRAM), Baden-Baden, Germany. May 2013.
4. Gelernter, J. and Balaji, S. (2013). An algorithm for local geoparsing of microtext. GeoInformatica 17(4):635-667. Springer US.
5. Ghahremanlou, L., Sherchan, W. and Thom, J. A. (2014). Geotagging Twitter Messages in Crisis Management. The Computer Journal. 2014.
6. Han, B., Cook, P. and Baldwin, T. (2014). Text-based Twitter User Geolocation Prediction. Journal of Artifical International Research. Jan 2014. 49 (1): 451-500.AI Access Foundation, USA.
7. Hiltz, S. R., Kushma, J. and Plotnick, L. (2014). Use of Social Media by US Public Sector Emergency Managers: Barriers and Wish Lists. In 11th International Conference on Information Systems for Crisis Response and Management (ISCRAM), Pennsylvania, USA. May 2014.
8. Imran, M., Elbassuoni, S., Castillo, C., Diaz, F. and Meier, P. (2013). Practical extraction of disaster-relevant information from social media. Social Web for Disaster Management (SWDM) Workshop 2013. WWW 2013 Companion, Rio de Janeiro, Brazil, May 2013.
9. Joachims, T. (1998). Text categorization with suport vector machines: Learning with many relevant features. In Proceedings of the 10th European Conference on Machine Learning, ECML 1998, pages 137–142, London, UK. Springer-Verlag.
10. Lindsay, B. (2011). Social media and disasters: Current uses, future options, and policy considerations. Technical report, Analyst in American National Government, Melbourne Australia, September. 2011.
11. Power, R., Robinson, B. and Ratcliffe, D. (2013). Finding Fires with Twitter. Proceedings of the Australasian Language Technology Association Workshop 2013 (ALTA 2013) December, 2013, Brisbane, Australia. pp 80-89.
12. Power, R., Robinson, B., Colton, J., Cameron, M. (2014) Emergency Situation Awareness: Twitter Case Studies. In 1st International Conference, ISCRAM-med, Toulouse, France, October 15-17, 2014. pp. 218-231.
13. Robinson, B., Power, R., Cameron, M. (2013). A Sensitive Twitter Earthquake Detector. Social Web for Disaster Management (SWDM) Workshop 2013. WWW 2013 Companion, May 13–17, 2013, Rio de Janeiro, Brazil. pp 999-1002.
14. Sakaki, T., Okazaki, M. and Matsuo, Y. (2013) Tweet analysis for real-time event detection and earthquake reporting system development. IEEE Transactions on Knowledge and Data Engineering, 25(4):919–931.
15. Schulz, A.; Ristoski, P. and Paulheim, H. (2013), I See a Car Crash: Real-Time Detection of Small Scale Incidents in Microblogs. In ESWC Satellite Events, volume 7955 of Lecture Notes in Computer Science, page 22-33. Springer, 2013.
16. Terpstra, T., de Vries, A., Paradies, G. and Stronkman, R. (2012) Towards a realtime Twitter analysis during crises for operational crisis management. In: Proceedings of 9th International Conference on Information Systems for Crisis Response and Management (ISCRAM), Vancouver, Canada, 2012.
17. Traverso, S., Cerutti, V., Stock, K. and Jackson, M. (2014) EDIT: A methodology for the Treatment of Non-authoritative Data. In 1st International Conference ISCRAM-med. Toulouse, France October 2014. pp 32-45.