First author's last name (use et al. if more than one author)
NOTE: DO NOT INCLUDE NAMES IN THE INITIAL VERSION REVIEWS ARE BLIND

TRIDS: Real-time Incident Monitoring with Social Media

<*Note*: for first submission leave authors and affiliation blank for double blind refereeing>
<If your paper is accepted, you will add these back in for the second upload>

ABSTRACT

Social media is increasingly being seen as a crucial information source during emergency events, providing real-time, first-hand reporting of incidents and allowing the tracking of collective perceptions and reactions. The majority of previous work has retrospectively examined the use of social media during extreme events. In the paper we discuss the deployment of TRIDS, a real-time, social media monitoring system, as an integral part of the control room in monitoring a number of events which, although planned, share features of emergency events, i.e. involving the management of large groups of people gathered in a locality. This includes one of the world's largest musical festivals. We discuss the practicalities of embedding social media monitoring into a traditional control room setting, the effectiveness of the various data and visual analytical aspects of TRIDS and the improvements in situation awareness gained from social media monitoring in a traditional event management control room setting.

INTRODUCTION

Social media has become a ubiquitous medium of communication, even by 2008 it had overtaken pornography as the primary activity on the Internet (Tancer, 2008), more recently it has become the main means by which users receive information, with social media activity tending to lead search engine activity by 4.3 hours (Kairam, Morris, Teevan, Liebling, and Dumais, 2013). During emergency events, where there is highly dynamic, complex situation evolution often involving a large number of individuals and organisations, there is an increase demand for up-to-the-minute information. In such events people post situation-sensitive information on social media related to what they experience, witness, and/or hear from other sources (Hughes & Palen, 2009), social media is used to inform others of their whereabouts and safety, or to organise themselves. The Arab Spring uprisings, showed the importance of social media in the organisation of the protests and in dispersing information to citizens: in Egypt and Tunisia, the vast majority of people surveyed over three weeks in March said they were getting their information from social media sites (88% in Egypt and 94% in Tunisia) (Storck, 2011). An analysis of relevant messages from six crisis events identified three broad categories of information: those reporting negative consequences of an event (20%-60%); those offering or asking for donations (15%-70%); and those warning about risks or providing advice (5%-20%) (Olteanu, Castillo, Diaz, & Vieweg, 2014).

For the organisations responsible for managing emergency events social media can provide a means to both: ascertain and track the collective state, perceptions and reactions of the those involved; and gather real-time, first-hand reporting of incidents. The general social media monitoring task can be

First author's last name (use et al. if more than one author)

defined as extracting, correlating and integrating the information contained in social media messages. There are four main dimensions of information that are relevant: who (profiles of the message authors and their social network, major participants in events, etc.), where (locations referred to or associated with the message), what (events and concepts referred to or associated with in the message) and when (the time-frame to which the message refers or the message is generated). The complexity of the monitoring task lies in the fact that social media streams are: (i) high in volume, and constantly increasing, (ii) often duplicated, incomplete, imprecise and incorrect; (iii) written in informal style (i.e. short, unedited and conversational), thus much less grammatically bounded and containing extensive use of shorthand, symbols (e.g., emoticons), misspellings, slang, irony, etc.; and (iv) covering every conceivable domain. These characteristics make social media data very noisy and the monitoring task hard.

In this paper we examine the use of social media monitoring in large-scale, localised events. Although these are planned, and therefore not emergency events per se, they have a number of features in common with unplanned emergency events: they involve the same organisations (e.g. local government and emergency services); have similar management structures (e.g. gold/silver/bronze control centres); and require reliable information in real-time. The key task is the identification of messages that report critical incidents or issues that may concern public security or safety, where the relevant information may be mentioned very infrequently, even in a single message.

RELATED WORK

Numerous systems have been developed to monitor the Twitter stream, general systems are built around the concept of a dashboard, or a set of visual displays that provides a summary of social media according to temporal, spatial, and thematic aspects. Common elements in these displays include (Imran, Castillo, Diaz & Vieweg, 2014):

- Lists/timelines showing recent or important messages.
- Graphs showing the frequency of a hashtag, term, phrase, or concept over time.
- Maps showing geo-tagged messages or clusters, possibly with message/cluster descriptor.
- Tag clouds or other visual summaries of the proportion of different message values.

The visual elements are backed by data analysis processes that include:

- Filtering of social media messages (mainly tweets) matching a given set of criteria
- Natural Language Processing, particularly text classification.
- Information Extraction, particularly geo-tagging
- Clustering of similar messages, generally used to identify sub-events.

Despite the fact numerous systems have been developed to monitor social media for emergency event management (Yin, Karimi, Robinson, & Cameron, 2012; MacEachren et al., 2011; Dou, Wang, Skau, Ribarsky, & Zhou, 2012; Marcus et al., 2011; Abel, Hauff, Houben, Stronkman, & Tao, 2012; Rogstadius et al., 2013), none of the systems have been evaluated during an actual emergency event. They are also generally aimed at identifying trending incidents and therefore require those incidents to be reported in a significant number of messages. Some systems, notably ESA (Yin et al., 2012) and SensePlace2 (MacEachren et al., 2011), emphasise the processing and visualisation of spatial and temporal information, however both these facets require incidents to be reported in a significant number of messages for them to be distinct. Similarly other systems (Dou et al., 2012; Rogstadius et al., 2013) focus more on processing message text to identify topics and generate summaries for incidents, thus require considerable amount of messages and are therefore mainly relate to major incident identification. Twitcident (Abel et al., 2012) does focus on employing classification algorithms to extract messages about very specific, small-scale incidents, although it only uses geocoded messages, which account for a small proportion (around 1%) of the total number of tweets.

TRIDS has been used within live control rooms and, although it provides functionality to identify trending incidents according to their temporal, spatial, and thematic information, it's primary focuses is developing pragmatic functionality which can identify small-scale incidents in real-time.

First author's last name (use et al. if more than one author)

TRIDS

TRIDS is a system for monitoring social media that enables situation awareness in localised events. It is designed using a modular, distributed architecture comprising of three main components, namely: Data Collection; Data Processing and Analysis; and Data Visualisation.

Data Collection

The data collection component comprises of modules connecting to the various data feeds from social media providers. These data feeds can be pushed (e.g. Twitter Stream API, RSS Feeds) or pulled (e.g. Flickr/Facebook APIs), the user is provided with a feed administration interface to select the feed criteria (e.g. keywords, locations), which can be both global for all feeds or specific to a single feed. The raw messages from each feed are stored in a separate database, and the identifier for each message is added to the processing queue. This allows for asynchronous processing of messages, which is necessary to implement real-time processing policies if the message volume exceeds the systems processing capability. Messages are processed on a LIFO basis to ensure the more recent messages are immediately available. If the processing queue exceeds a specific threshold, further processing resources can be brought online (if available).

Data Processing and Analysis

The data processing and analysis pipeline is concerned with normalisation and annotation of social media messages, which can then be indexed for efficacious search/browsing. While social media is a multimedia environment, written text provides the majority of machine interpretable information. However the language use in social media offers a number of challenges to interpretation (Baldwin, Cook, Lui, MacKinlay, & Wang, 2013), including non-standard punctuation, capitalisation, spelling, vocabulary and syntax. There is a huge diversity of styles within social media both between and within different platforms. Twitter messages can vary from personal communications to official pronouncements, it is highly multilingual and an individual message can contain multiple languages. This informal language use has been shown to significantly affect the performance of text analysis methods (Gimpel et al., 2011; Ritter, Clark, Etzioni, et al., 2011). TRIDS text analysis principally concerns the extraction of key information using fast, lightweight Natural Language Processing techniques. However consideration is also given to the processing of other data types, such as locations and time, in order to ensure they can be efficiently used for real-time, large-scale visualisations.

Language identification

A language identifier is applied to the text, to determine the most likely language. The text is pre-processed to remove specific functional text, which can be language independent (such as, Twitter @mentions, #hashtags, RT prefix). Identification is performed using the language detection tools provided by Cybozu Labs (Nakatani, 2011), which employ Naive Bayes on character n-grams. This tool was shown to be highly accurate (99%) on well-formed text, however for the short, informal, idiomatic and often mixed-language text found in social media it has been shown accuracy is significantly reduced, unless identification techniques consider the foibles of social media text (Carter, Weerkamp, & Tsagkias, 2013). To mitigate this, (30) language profiles are created from social media messages (i.e. Twitter corpus), as using such domain specific profiles is shown to improve accuracy (Carter et al., 2013).

Information propagation and duplicate detection

The propagation of information is a key behaviour in social media, and particularly in Twitter, where retweeting is a fundamental and informative form of behaviour. It is crucial to understand the propagation of information to determine its provenance (where in the information originated) and importance. Within Twitter, for retweets the attribution is recorded in the metadata (with a propagation count), however it is possible for messages to be repeated without this information, both with and

First author's last name (use et al. if more than one author)

without modification. To identify such messages it is necessary to detect duplication, by normalisation of the text. The text is then normalised by removing the functional text relating to retweeting convention (e.g. prefixing of RT and @username), all non-alphanumeric characters are removed and characters are folded into ASCII (where possible) then word breaks separated by single-spaces. In addition, where possible, the links contained in the message are normalised to their final URL (to deal with the use of URL shortening). Duplicates are then identified using the integer hashcode of this normalised text string.

The repetition of links and images are also considered, as these are often the critical information being conveyed in the message. Shortened URLs are frequently used in tweets, these are resolved (if possible) to the full URL. If the link is to an image, the system also implements a fast, low-level, feature-based image-processing scheme, based upon perceptual hashing (Monga & Evans, 2006), to detect duplication via salient image feature points. In terms of latency the actual image processing is orders of magnitude less important than the time required to download the images. Therefore the real-time use of this technique is only appropriate where the volume of messages is low enough that the image download does not cause network issues.

Sentiment Analysis

Sentiment analysis is used to determine the positive or negative attitude of text. Positive and negative emoticons and terms are identified in the text, in addition preceding terms are identified as modifying or negating. These terms are provided in a predefined lexicon. Currently only an English lexicon is available, containing 58 emoticons (29+ve/29-ve), 10,973 terms (3,754+ve/7,219-ve), 60 intensity modifiers (43+ve/17ve) and 44 negations. To compute the term and modifier weights we use a training set of 1.6 million tweets (Go, Bhayani, & Huang, 2009), which is evenly distributed between positive and negative tweets, as determined by the occurrence of a positive or negative emoticon.

Autotagging

TRIDS provides a mechanism for automatically adding tags based upon string matching of the text tokens. The process supports expression matching, i.e. matching multiple tokens, and synonym matching (using a finite-state automata for efficient matching of the alternatives). The additional tagging is specified in a taxonomy such that if any of the (more specific) child tags match then the (more general) parents are also matched. For example, a location taxonomy could include continent, country and city (e.g. Africa-Botswana-Gaborone). Currently there is no online disambiguation process, however each term is, a priori, assigned two levels of ambiguity based upon the frequency of occurrence of the tagging term in a general corpus, and the frequency of occurrence within a specific domain, for example the location terms are derived from place names in OpenStreetMap and so each location name is assigned an ambiguity based on the frequency of occurrence in all the OpenStreetMap location names. These ambiguity values can be used to select the level of acceptable term ambiguity, used to determine if the message should be geo-located with the latitude/longitude of the identified location, this ensures that only unambiguous locations.

Visual Analytics

The TRIDS interface is designed to provide sense-making and analytical reasoning capabilities to the user. To solve real-world problems, users and computers must cooperate in harmoniously integrated workflows that maximise the relative analytical capabilities of humans and machines (Schreck & Keim, 2013). Therefore, the aim is to provide seamless interaction between human user and computer to gain relevant insights from the data to support the exploration and understanding of the large-scale, real-time and complex information provide by social media. The interface is web-browser based and provides a set of interconnected widgets for searching, browsing and visualisation of the information indexed within TRIDS. Each widget displays the values for an information facet associated with the messages, allowing the user to select/deselect that subset of messages relating to their current focus of interest. The currently selected messages are displayed in a list showing: (i) their provenance (which social media source), (ii) a link to the original message and to the author profile in the original

First author's last name (use et al. if more than one author)

social media platform, (iii) a link to the additional information extracted by TRIDS such as whom the user has communicated with, main topics associated to the user, etc. (iv) specific terms and concepts associated by the system (e.g. automatic tags). The user can also annotate messages (and users) using tags/terms for future retrieval.

The other widgets have been developed from feedback after initial trials and considering previous work surveying professionals needs (MacEachren et al., 2011), which indicates the expected feature set of a social media crisis/emergency management system is: maps (94.7%), photos/video collections (71.1%), time graphs (60.5%), keyword clouds (57.9%) and clustering tools (47.4%).

Maps

Markers on a map indicate the location of the geocoded messages, with multiple proximate messages clustered into a single marker. As the performance of the map deteriorates as the number of clusters increases each message is indexed (using a geohash) at varying levels of precision, the location precision requested by the interface is relative to the zoom level. This provides the ability for the map to show any number of data points. In addition to each marker indicating the number of messages clustered at that marker, the marker can also portray the main topic (tag or keyword) of the messages in the cluster. The map can also be used as a filter, so that information is constrained to only those messages currently within the view point.

Image/video collections

The importance of information is generically defined by its frequency within the current context, as specified by the information filters. A ranked list is used to display: text messages (as defined by the text hashing described above); URL links; and images (as defined either by the URL or, if available, the image hash).

Time graphs

The temporal information is presented in a graph, which shows the frequency of multiple tags/terms over time. The user can focus the timeline on a specific set of tags/terms, and focus on a specific timeframe. Similarly to the spatial information, the message time is indexed a varying levels of precision, with the level of precision requested being dependent on the current timeframe.

Keyword list/clouds

Lists and clouds are used as the generic widget for displaying the current information facet values. When the number of values is limited (e.g. message sentiment positive/neutral/negative) then ordered lists are used which display the frequency of messages containing the values that match the current selection. Clouds are used when more values are possible, with font size being relative to frequency, for displaying the current most frequent, terms and tags (both those assigned automatically and by the user), the most prolific message authors, and the users most often mentioned in the currently selected messages.

Clustering tools

A network representation is used to display the interconnections, where the nodes represent the message tags, terms and users, and the edges the co-occurrence frequency of that node type with the messages. The connected network of nodes aims to indicate the clusters of related information, which may indicate a specific incident.

EVENT MONITORING

The TRIDS system has been deployed in numerous event control rooms over the last 3 years to monitor some very large localised events on behalf of emergency responders and event organisers.

First author's last name (use et al. if more than one author)

The events monitored included: Glastonbury, one of the largest open air music festivals in the World (a five-day event involving around 200,000 people), the Bristol Harbour Festival (a weekend event involving 150-200,000 people), the Leeds Music Festival (a weekend event with 80,000 people) and Bristol's St Pauls Carnival (a one-day event involving 100-150,000 people). In addition TRIDS has been used during an exercise in preparation for the London Olympics and monitoring protests during the Conservative Party Conference.

The reference TRIDS user is an analyst working for emergency responders or for the organisers of the event situated in the main control room (e.g. gold, silver or bronze command in the UK). As large events have a very dynamic nature an unexpected situation can arise at any moment (e.g. a rain downpour can suddenly disrupt a concert causing the crowd to evacuate an area quickly) and stop equally suddenly, a major requirement for TRIDS was to provide the ability to quickly discover new incidents and follow up some of them without ever losing the big picture (focus without loss). This requires the coexistence of a horizon-scanning mode of operation (to discover new incidents and issues) interleaved with an investigation mode (to determine the importance and follow the evolution of identified incidents and issues).

A monitoring session is initiated with the parameters of the event defined by a set of keywords, relevant users and locations used to filter the incoming social media streams, to provide general messages of interest. Given the specifics of the event, i.e. type (e.g. flood, music festivals), location, etc., a predefined taxonomy of terms and regular expressions are used to initialise the automatic tagging process. However, in order to account for the dynamic nature of events, analysts are able to modify these initial setting, e.g. add/remove keywords, locations and people, at any time during the event.

The TRIDS analytic interface is used to implement the two main concurrent modes of operation: horizon scanning and investigation. An analyst performs horizon scanning for general term/tag patterns of interest. When an investigation lead is identified (e.g. a major car accident on an road leading to the event venue), a new browser tab is opened and the user can further focus the view by adding relevant keywords, geographic location, etc. The analyst can follow a number of different parallel, eventually determining whether to log and forward the incident/issue for consideration by other agencies.

Case study

This section focuses on one of the main deployments of TRIDS, involving the monitoring the Twitter stream concerning the 2013 Glastonbury music festival. In order to provide 24-hour coverage over the course of 5 days a team of six people working eight-hour shifts was employed. The shifts overlapped to allow for handover, so that pertinent information could be given to the incoming analyst. The monitoring activity took place within the control room, where all the agencies involved in the management of the event were located. The initial monitoring brief was very wide; to report on any ”interesting” events in the social media, especially were they might involve public safety and security issues. The only specific brief was to identify the occurrence of ”secret gigs”, where potentially renowned artists can perform unannounced, this can lead to crowd movements which can cause numerous issues if the venue becomes crowded beyond its safety capacity.

The monitors logged the ”interesting” messages on a time-line, this log was used to create a report that was input into the periodic (8-12 hourly) multiagency team briefing sessions. In addition, if the analyst determines an incident/issue is urgent they can report immediately to the relevant agency. On a number of occasions other agency members would also ask the monitoring team to focus attention on specific events, rumours or potential issues.

At the beginning of the festival the system was initialised to identify: messages from the location of the festival site, and surrounding area; messages sent by a number of Twitter users related to the festival, e.g. the organisers, performers, local media; and messages which contained one or more terms pertinent to the festival. As the festival progressed some addition users and terms were added to the filter, in addition a number of terms were added to a blacklist, which removed messages containing

First author's last name (use et al. if more than one author)

unrelated or ambiguous terms from the system.

Over the course of the festival 1,951,499 tweets were indexed by the system, the majority of these (84.1%) on the three main festival days (Friday-Sunday). The peak hour was 22:00-23:00 on the Friday, when 74,645 tweets (20.74 tweets per second) where received, the peak minute was at 01:41 on the Sunday night, when 2,948 tweets (49.13 tweets per second) where received. The vast majority (93.7%) of tweets were identified as English, unsurprising due to the filtering terms used. 3.83% of the tweets were geo-coded, although only 0.15% were geocoded within the festival site. In terms of sentiment, three main negative issues were identified: a resentment that the main event organiser had shown support for Badger culling (there was some concern that this may lead to protest action but none occurred); a negative response to a controversial performer (despite which his actual performance passed without incident); and the set of messages bemoaning the fact the message author was not at the festival.

In terms of actual incidents logged over the course of the festival four broad categories were used, crowd management (with a sub category for "secret gig" rumours), complaints related to the onsite festival, traffic related, and other. In total 330 (0.017%) messages were logged, which related to 217 separate incidents. The split was crowd management 30% (of which 84% related to "secret gigs"), complaints 13%, traffic 27%, and other 30%. In terms of crowd management not relating to secret gigs, the main concern was potential protests over the Badger culling (which did not occur). There were 13 rumours of "secret gigs", 6 of which actual occurred all of which were managed without incident. The majority of the complaints concerned facilities, particularly the provision of water and the state of toilets, however none of these were deemed to required specific attention. Of the traffic issues 74% occurred on the first day and all were minor. The other issues mainly concerned thefts and drug related health concerns.

The two main other issues which required immediate attention were the unexpected arrival of a member of Royalty, which cause concern over crowd management and also media coverage, and the posting of the sites wireless passwords. Both of these incidents were identified by trending images, although in general incidents were identified by the automatic tagging of key terms relating to the types of categories, supplemented with additional focus on terms/tags and users for individual investigations. The timeline, while providing indication of the trending term/tags, was not useful in identifying specific incidents/issue. The map was also deemed not useful for this task, and in general we have found that the location information suffers from the low proportion of geocoded data, which means that the vast majority of information is not represented on the map. Although in previous deployments, e.g. during a political party conference, the identification of place names was useful in identifying the occurrence of location specific protests. It is worth noting that the other agencies did not report any incidents that were missed by the social media monitoring.

DISCUSSION AND FUTURE WORK

A number of observations arise from the experience of the deployment of TRIDS. They are: verifying reliability of information in social media, connectivity problems, impact on governance and the resources used for monitoring.

Reliability of information

If event managers or emergency responders are going to devote resources on the basis of a single or small set of social media messages, it is important that they have a notion of the reliability of the information. Misinformation may be intentional, e.g. desire to incite a response, or unintentional, e.g. due to the use of irony or sarcasm, or a genuine misinterpretation on the part of the message's author. Yuan, Guan, Huh, Lee, et al. (2013) noted that, in case of an earthquake disaster, 70% of messages relating to people being trapped were unreliable, mainly due to searchers unwillingness to give up hope on their relatives or desire to recover the body as soon as possible. In our implementation the two approaches to verification have been to inform local staff on the ground of the incident so they can investigate, or contact the originator of the message over social media for further

First author's last name (use et al. if more than one author)

confirmation. The TRIDS system supports assessing the reliability of a message by providing the ability to view the profile of the user, or review the last (200) message they posted. One of the basic criteria for determining whether to consider a message is whether the report is first-hand, to this end messages can now be automatically tagged as local to the event, dependent on geocoding or textual pattern, although the accuracy of the process has not been evaluated. Further work should examine more robust methods to determine whether incident reports are from first-hand observers and methods to automatically assess message reliability (and importance). That said, it is worth noting that while the issue of reliability of information is amplified in social media, it is also present in traditional communication methods, e.g. for emergency call centres in the UK it is estimated only 25% of 999 (emergency) calls are actual emergencies[1].

Connectivity problems

Social media monitoring relies on individuals having the ability to use their mobile devices to send messages, which requires connectivity and power. Mobile networks can suffer slowdowns or even outages during emergency events, as the network infrastructure may be adversely affected by the emergency and there will be increased communication demand. In the UK the emergency services were criticised for their reliance on mobile devices during the 2007 London bombings (Resilience, 2006), and in the aftermath investment was made into the use of alternative radio networks. However, as is apparent from numerous examples, social media and mobile devices have been widely used during emergency events. The increasing coverage and bandwidth of 3G and 4G networks combined with the desire for emergency services to utilise data traffic means they are moving towards incorporating (or fully adopting) standard mobile networks for communication[2]. From our experience, where Glastonbury takes place in a rural setting, in 2013 there were numerous messages complaining about the lack of signal and inability of the charging stations to cope with demand. By 2014 the festival site had temporary 4G cell towers and a system of battery pack rental that all but alleviated any issues, similar provisions could be used during emergency events.

Impact on governance

The monitoring of social media can be controversial, organisations are concerned about two main, and to a degree conflicting, issues: that they may be seen as undertaking covert surveillance and therefore adversely impact their relationship with the public; and that there may be legally liable for the outcome of acting on such non-verified information also, for some organisations, there can even be a legal responsibility for not acting on information once they become aware of it. To alleviate the criticism of surveillance organisations need to be open about their use of social media, however as organisations increasingly use social media and provide an "official" means of communication via social media, e.g. have a Facebook page, or Twitter account, there is an increase expectation that messages posted, particularly ones directly referring to a given organisation, will solicit some response.

Monitoring resources

One of the main issues in real-time social media monitoring is the human resources required, even with support systems, such as TRIDS, constant human attention is required. One solution is to harness volunteers as monitors, this "crowdsourcing" has been shown to be applicable to large-scale disasters, where it is possible to mobilise national and international concern for those affected (Starbird, 2011). However these volunteers were looking at messages directly related to the emergencies, volunteers may well be less inclined to wade through largely unrelated social media

---

[1] http://www.dailymail.co.uk/news/article-2078069

[2] https://www.gov.uk/government/publications/the-emergency-services-mobilecommunications-programme

First author's last name (use et al. if more than one author)

messages. Developing processes which automatically filter messages would alleviate the monitoring burden, although there would have to be a clear understanding of the level of recall of such filters, i.e. the number of incident related messages filtered out. Future work could also examine the use of collaborative monitoring techniques, to minimise any duplication of effort when multiple monitors are employed.

CONCLUSIONS

In this paper we have presented a system for event monitoring via social media that has been adopted by emergency responders and event organisers to monitor events involving hundreds of thousands of people to complement traditional control room techniques. The system has been key in successfully identifying incidents that required immediate intervention and that would have not been identified at all, or as promptly, with traditional means. The system continues to be developed given the experiences gained and has become an integral part of the control room in a number of annual large-scale events.

REFERENCES

Abel, F., Hauff, C., Houben, G.-J., Stronkman, R., & Tao, K. (2012). Twitcident: fighting fire with information from social web streams. In Proceedings of the 21st international conference companion on World Wide Web (pp. 305–308).

Baldwin, T., Cook, P., Lui, M., MacKinlay, A., & Wang, L. (2013). How noisy social media text, how diffrnt social media sources? In Proceedings of the 6th international joint conference on natural language processing.

Carter, S., Weerkamp, W., & Tsagkias, M. (2013). Microblog language identification: Overcoming the limitations of short, unedited and idiomatic text. Language Resources and Evaluation, 47(1), 195–215.

Dou, W., Wang, X., Skau, D., Ribarsky, W., & Zhou, M. X. (2012). Leadline: Interactive visual analysis of text data through event identification and exploration. In Visual analytics science and technology (vast), 2012 IEEE conference on (pp. 93–102).

Gimpel, K., Schneider, N., O'Connor, B., Das, D., Mills, D., Eisenstein, J. & Smith, N. A. (2011). Part-of-speech tagging for twitter: Annotation, features, and experiments. In Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies: short papers-volume 2 (pp. 42–47).

Go, A., Bhayani, R., & Huang, L. (2009). Twitter sentiment classification using distant supervision. CS224N Project Report, Stanford, 1–12.

Hughes, A. L., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. International Journal of Emergency Management, 6(3), 248–260.

Imran, M., Castillo, C., Diaz, F., & Vieweg, S. (2014). Processing social media messages in mass emergency: A survey. arXiv preprint arXiv:1407.7071.

Kairam, S. R., Morris, M. R., Teevan, J., Liebling, D., & Dumais, S. (2013). Towards supporting search over trending events with social media. In Proceedings of ICWSM 2013, the 7th international AAAI conference on weblogs and social media.

MacEachren, A. M., Jaiswal, A., Robinson, A. C., Pezanowski, S., Savelyev, A., Mitra, P., & Blanford, J. (2011). Senseplace2: Geotwitter analytics support for situational awareness. In Visual analytics science and technology (vast), 2011 IEEE conference on (pp. 181–190).

Marcus, A., Bernstein, M. S., Badar, O., Karger, D. R., Madden, S., & Miller, R. C. (2011). Twitinfo: aggregating and visualizing microblogs for event exploration. In Proceedings of the SIGCHI conference on human factors in computing systems (pp. 227–236).

Monga, V., & Evans, B. L. (2006). Perceptual image hashing via feature points: performance

First author's last name (use et al. if more than one author)

evaluation and tradeoffs. Image Processing, IEEE Transactions on, 15(11), 3452–3465.

Nakatani, S. (2011). Language detection library for java (Tech. Rep.). Tech. rep., Cybozu Labs, Inc.

Olteanu, A., Castillo, C., Diaz, F., & Vieweg, S. (2014). Crisislex: A lexicon for collecting and filtering microblogged communications in crises. Proceedings of ICWSM.

Resilience, L. (2006). Looking back, moving forward: The multi-agency debrief–lessons identified and progress since the terrorist events on the 7 July 2005. In London: London regional resilience forum, government office for London.

Ritter, A., Clark, S., Etzioni, O., et al. (2011). Named entity recognition in tweets: an experimental study. In Proceedings of the conference on empirical methods in natural language processing (pp. 1524–1534).

Rogstadius, J., Vukovic, M., Teixeira, C., Kostakos, V., Karapanos, E., & Laredo, J. (2013). Crisistracker: Crowdsourced social media curation for disaster awareness. IBM Journal of Research and Development, 57(5), 4–1.

Schreck, T., & Keim, D. (2013). Visual analysis of social media data. Computer, 46(5), 68–75.

Starbird, K. (2011). Digital volunteerism during disaster: Crowdsourcing information processing. In Conference on human factors in computing systems.

Storck, M. (2011). The role of social media in political mobilisation: a case study of the January 2011 Egyptian uprising. University of St Andrews, Scotland.

Tancer, B. (2008). Click: What millions of people are doing online and why it matters. Hachette Digital, Inc.

Yin, J., Karimi, S., Robinson, B., & Cameron, M. (2012). ESA: emergency situation awareness via microbloggers. In Proceedings of the 21st ACM international conference on information and knowledge management (pp. 2701–2703).

Yuan, W., Guan, D., Huh, E.-N., Lee, S., et al. (2013). Harness human sensor networks for situational awareness in disaster reliefs: A survey. IETE Technical Review (Medknow Publications & Media Pvt. Ltd.), 30(3).