

Spatio-Temporal Decision Support System for Natural Crisis Management with TweetComP1

Stuart E. Middleton¹(✉), Andrea Zielinski², Öcal Necmioğlu³,
and Martin Hammitzsch⁴

¹ University of Southampton IT Innovation Centre, Gamma House,
Enterprise Road, Southampton SO16 7NS, UK
sem@it-innovation.soton.ac.uk

² System Technologies and Image Exploitation,
Fraunhofer Institute of Optronics, Karlsruhe, Germany
andrea.zielinski@iosb.fraunhofer.de

³ Kandilli Observatory and Earthquake Research Institute, Istanbul, Turkey
ocal.necmioğlu@boun.edu.tr

⁴ GFZ German Research Centre for Geosciences, Potsdam, Germany
martin.hammitzsch@gfz-potsdam.de

Abstract. This paper discusses the design of a social media crisis mapping platform for decision making in natural disasters where tweets are analysed to achieve situational awareness during earthquake and tsunami events. A qualitative end user evaluation is undertaken on our first prototype system to get feedback from practitioners working in the field of hazard detection and early warning. Participating in our evaluation is the Kandilli Observatory and Earthquake Research Institute (KOERI) and the Portuguese Institute for the Sea and Atmosphere (IPMA). We conclude that social media crisis mapping is seen as a valuable data source by control room engineers, with update rates of 10–60 s and false positive rates of 10–20 % (general public incident reports) needed. Filtering crisis maps and statistical reports by social media platform and user type is desirable as different report sources have different credibility and response times.

Keywords: Decision support · Web 2.0 · Twitter · Social media · Crisis mapping · Geoparsing · Multilinguality · Early event detection

1 Introduction

In recent years vast amounts of Web 2.0 content has become available, such as Twitter messages, YouTube videos, blogs and RSS feeds. Such ‘unconventional’ social media sensors allow rapid in situ crowd-sourced measurement by people actually experiencing a crisis event (e.g. via mobile devices). Social media can provide incident reports on tsunami wave arrival and inundation at coastal areas. This provides Tsunami simulation units and decision makers with observations on the wave’s propagation and impact. In highly populated areas such as the Mediterranean region there can be a huge

number of potential coastal areas at risk [11], so receiving timely incident reports could be of great benefit to a Tsunami warning centre.

However, social media streams also bear risks: such data has variable quality and a high noise to signal ratio. Ultimately decision makers need a situation assessment picture they can trust, is updated in real-time and offers the possibility to manually check and verify incident reports on demand. It is particularly important to find the right balance between responsive alerting and avoidance of time-wasting false positives.

We have observed during our own monitoring of recent crisis events (e.g. New York hurricane Sandy 2012, Oklahoma tornado 2013, Philippines Tsunami 2012) keyword filtered Twitter Streaming API throughput of up to 20,000 tweets per hour [17]. Given the rapid adoption rate of social media around the world this throughput is only expected to get larger. On average about 1 % of the tweets we observed during crisis periods contained a geotag, making the processing of tweet content important if we want access all available location information. Social media crisis mapping is the provision of situation assessments from real-time social media feeds. Key research challenges include intelligently handling the large volume of social media reports available, keeping false positive rates low and accurately geo-parsing map information from microblog content. Further to this is the challenge of convincing Tsunami warning control room engineers of the value social media reports offer to a situation picture for inclusion and use in the control room itself.

Our approach of combining content-based geo-parsing of social media streams, at both the region and street-level, multi-lingual event classification and real-time crisis mapping is novel. Our end-user evaluation of an integrated tsunami early warning system prototype combining social media and traditional sensor-based data sources is also a valuable contribution to practitioners in this field. Our goal is to quantify the tsunami early warning control room end-user performance targets for acceptance of social media crisis mapping into the decision support workflow, and analyse how close this technology is to meeting these targets.

2 Decision Support for Tsunami Event Natural Crises

The heart of an Early Warning System is an Operation Centre, where the seismic and/or tsunami monitoring networks are connected via reliable and robust communication channels utilizing a Decision Support System based on event source parameters, such as magnitude, location and depth. This information constitutes the essential part of the early warning process, and any errors introduced at this stage may lead to false alarms, especially considering that it would also affect the assessment of the impacted areas. In fact, an Early Warning Centre should adopt all possible methodologies to minimize the false alarm rate, since any false alarm would undermine the credibility of the Early Warning Centre, and therefore would jeopardize the community's earthquake preparedness. On the other hand, a false alarm in the case of earthquake may only correspond to the ill-determined magnitude, since a seismic network consisting of tens or hundreds of stations could not lead to a false alarm for a non-existing event.

For a real event the combination of unconventional human sensors through the use of web 2.0 system would provide a verification mechanism beyond conventional methodologies. The mapping of social media activity presented to the operators could assist the information obtained from instrumental recordings.

In the case of an earthquake, near real-time shake maps providing ground motion and shaking intensity maps, are of crucial importance for Civil Protection and Disaster and Emergency Management Authorities to assess the criticality of the event in order to better assist post-event response and recovery actions. Social-media activities, displayed jointly with the shake maps, could provide a near-real time correlation of the two datasets of different nature, which could provide the extent of the damage associated to the events.

In the case of the tsunamis, parameters such as arrival times and tsunami wave-heights obtained from pre-calculated or near real-time tsunami simulations are the key elements. Nevertheless, in the case of tsunamis, the scale of the instrumental observations is much more lower than the earthquake case, especially in a region where DART-like offshore buoys are not operated. In such an environment, Early Warning Centre's main conventional support is tide-gauge measurements at the coastal locations, from which the Centre could initiate an "early-warning" only applicable for a further distant location. In such a situation, the human sensors could play a very important role in terms of identifying the affected areas, such as through the interactive visualization of Twitter streams. This assumption, however, may not be valid in the areas affected strongly by extreme catastrophic events, where the communications lines could be destroyed or the human sensors would be ineffective in localities of major destruction. On the other hand, such inactivity could also be considered as the evidence of major destruction if some activity could be monitored in the surrounding regions favourably located. Such information would be of "supplementary" status, and its primary use would be by the Civil Protection and Disaster and Emergency Management Authorities.

The TRIDEC project [5, 6] is building a prototype tsunami early warning system for the North-Eastern Atlantic and Mediterranean (NEAM) region. Our aim is to provide decision support for National Tsunami Warning Centre (NTWC) deployments integrating 'conventional' real time sensor data (e.g. seismometer networks, tide gauges, GPS stations, GPS buoys, ocean bottom units including pressure gauges) and social media human sensor data (e.g. crisis maps, news feeds). Additionally access to support information such as bathymetric and topographic data sets, pre-computed tsunami simulations and/or on-demand simulation computations is provided to help make predictions on Tsunami wave propagation and the estimated impact times at coastal regions. Operators on duty need to quickly assess the relevance and tsunami-genic properties of earthquake events and, if necessary, make use of information logistics and dissemination facilities to send customized, user-tailored warnings to responsible authorities and the general public under immediate risk.

A 'conventional' command and control user interface [2] is shown in Fig. 1. Our crisis maps, alongside publically available social media and traditional media reports such as news reports and mobile video footage, can provide additional rapid incident reports from people experiencing both earthquakes and Tsunami wave inundation. Our hypothesis is that the combination of situation assessment from 'conventional'

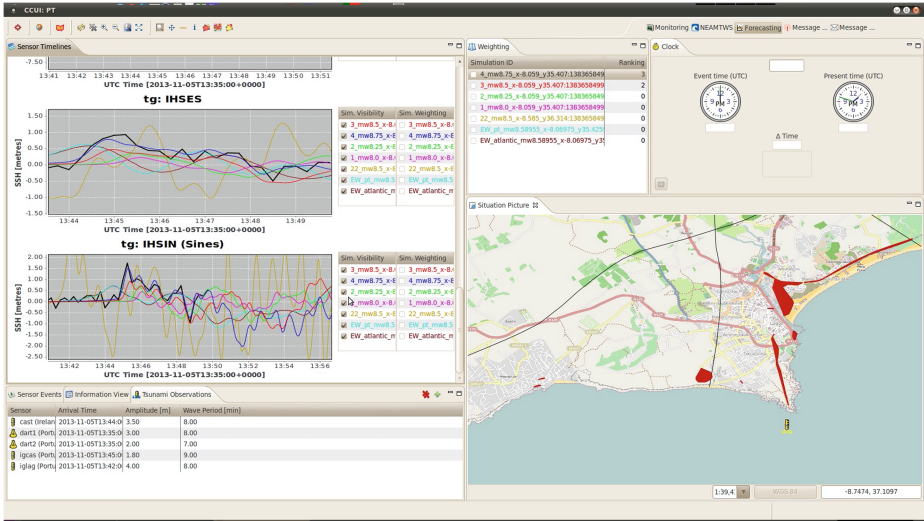


Fig. 1. TRIDEC command and control user interface (CCUI) with sensor-data (left) and social media crisis map (right) showing flooded areas (red lines and polygons) (Color figure online)

hardware-based sensor systems and ‘unconventional’ social media systems will improve the overall situation picture quality and timeliness within the control room.

Decision support is provided via the command and control user interface, for situation assessment and manual decision making, and a BPMN2 workflow engine for faster automated decision making. Control room personnel can define a decision matrix with rules to handle specific events in a simple Excel spreadsheet, which is imported and converted to Drools Expert rules for automated execution. Workflows and rules exist for the Tsunami relevance, assessment, verification and warning phases. We use rules triggered on earthquake alerts to start recording tweets and generating relevant crisis map display layers for the command and control user interface.

3 Related Work

In the last decade there has been growing interest in the exploitation of social media for crisis management. The most prominent framework for earthquake and tsunami detection, Toretter [7, 8], is based on a probabilistic state sequence model which aims to locate the centres of earthquakes by means of Kalman and particle filters. Their experiments confirm that Toretter often can send notifications for detected earthquakes quicker than traditional seismic-based sensors. The Tweet Earthquake Dispatch (TED) system managed by U.S. Geological Survey (USGS) [14] uses a short-term-average over long-term-average algorithm, which compares the frequencies of raw incoming tweets over a fixed time window against the background noise to detect seismic events. Furthermore [10] propose to filter tweets with respect to whether they might contribute to enhance situational awareness during an earthquake event using linguistic features. The systems are operated in earthquake endangered zones, i.e. Japan and the United States.

Humanitarian organizations and networks of volunteer's have setup live web-based manual crisis mapping sites [4] for recent natural disasters such as the Haiti 2010 earthquake, New York's 2012 hurricane Sandy and Oklahoma's 2013 tornado. These organizations manually filter crowd-sourced information from news reports, social media and civil protection agency alerts, and collate it into live web-based crisis maps for the general public to see. The challenge [4] for these organizations is to automate the real-time data fusion of large volumes of multi-source heterogeneous information for scalability, whilst still maintaining trust and credibility in this data.

In the context of decision support or early warning, a strong focus needs to be placed on usability and integration aspects [9]. Most central to this is the possibility to browse and visualize social media data in real-time. For instance, the Emergency Situation Awareness (ESA) system has been set up in collaboration with the Australian Crisis Coordination Centre to fulfil these aims. ESA can cope with different emergency cases, for which social media activity (Twitter, images, video) is displayed on a map and presented to the operators [13]. Twitcident [1] is a framework for filtering, searching and analysing information for different kinds of real-world incident and social media types. It is connected to an emergency broadcasting service for adaptive filtering of Twitter information and uses enhanced semantic analysis tools and faceted search facilities. SensePlace2 [3] is another framework which is able to extract deeper semantics from Social Media streams. Moreover, it uses elaborate visualization techniques to present an overview to the user [3]. In a user-study, the user-interface was designed such as to provide optimal support.

4 TweetComP1 Architecture: A Framework for Human Sensor Network Analysis and Crisis Mapping

The TRIDEC social media framework [16], see Fig. 2, is called TweetComP1. For scalable communication we use an event bus, deploying a message oriented middle-ware backbone. The TRIDEC database layer is a shared workspace for processing services. A decision support framework driven by expert-authored decision tables orchestrates the processing sub-systems. Various video demonstrations of the TRIDEC prototype in action can be found at [18].

Real-Time Acquisition of Social Media: An adaptive twitter crawler [12], which is capable of following trending keywords and adapting the filter keywords used in requests to the Twitter Streaming API. For basic crawling we have a list of earthquake and flood related keywords in multiple languages (e.g., deprem, terremoto). All tweet data and metadata is logged to a MySQL database.

Geoparsing and Geospatial Clustering: Offline processing involves the download and geocoding of map data for coastal regions at possible risk of Tsunami. Map information sources include OpenStreetMap, Google Places API and gazetteers such as Geonames and GEOnet Names (GNS). The location names, addresses and OpenGIS geometry are loaded into a MySQL database and geocoded using the Goggle Geocoder API to extract address components such as short versions of street names. Map sources provide street and building level locations whilst gazetteers provide region level locations.

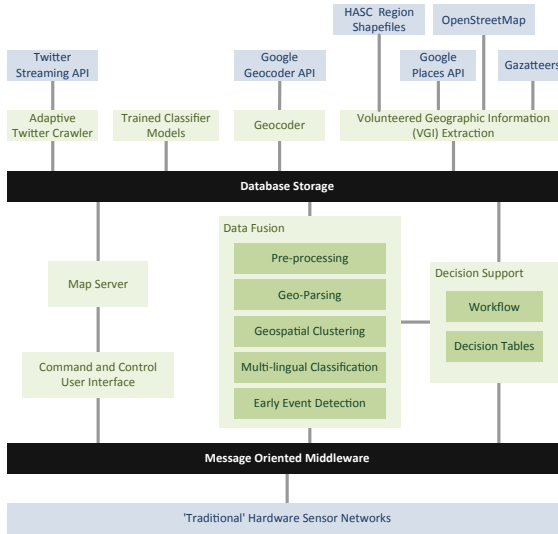


Fig. 2. Information viewpoint for the TweetComPI architecture

For online processing Tweets are geo-parsed [17] and locations matched to the offline database of coastal geographic information. Tweets are first tokenized into n-gram tokens and then named entity matching performed. A baseline historical dataset, during which no crisis occurred, is used to statistically calculate threshold levels for map reporting; this helps reduce false positives for locations with common names such as ‘Hollywood’ that are tweeted often. Hierarchical clustering is used to overlay polygons onto groups of nearby streets. Figure 3 shows an example of a crisis map display showing tweeted street level flood reports.



Fig. 3. Street level crisis map of flood inundation for a Tsunami scenario in Portugal

Multilingual Classification: It is often difficult to separate real valuable information from background noise in social media data. To handle the increasing load of noisy and off-topic information, enhanced filtering techniques are needed. Our approach is based on supervised learning [15] and works for a variety of languages. A configurable threshold level is used to filter out tweets below an acceptable relevance score.

Figures 4 and 5 shows our tweet analytics interface during an Earthquake in Turkey. In live-mode, the display is continuously refreshed. In historic mode, predefined scenarios for specific events can be selected and replayed in normal or acceleration mode. The twitter activity can be shown for any user-selected time range or can be configured for a certain geographic region.




Fig. 4. Visual analytics interface including timeline statistics for a Tsunami scenario in Turkey

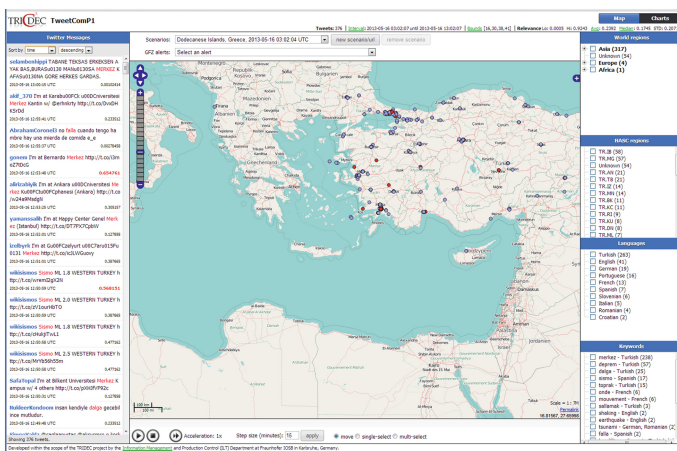


Fig. 5. Region level crisis map of earthquake reports for a Tsunami scenario in Turkey

Additionally, a visual presentation of three graphs is displayed: a timeline graph with the hourly volume of tweets, a bar graph that clusters tweets by their relevancy score, and a pie graph that shows the share of languages and geographic destination of the tweets.

Early Event Detection: With the increasing number of Twitter users all over the world it becomes possible to automatically detect novel earthquake/tsunami events from a twitter stream in a timely manner. We aggregate tweets for a certain spatial region and time-span, e.g. an interval of one minute. Geo-location information is extracted from metadata (i.e. geotag) and the unstructured twitter text (i.e. region names) and mapped to Hierarchical Administrative Subdivision Code (HASC) regions.

Emergency Alerting: Warning alerts are triggered according to specified thresholds, for example if the relevant tweets per time period in a region is exceed. The alert message format is based on the Common Alerting Protocol (CAP) standard of the Organization for the Advancement of Structured Information Standards (OASIS).

Crisis Map Visualization: We can use either the MapServer or Geoserver mapping applications, driven from OpenGIS geometry data in a MySQL database. At the region level we plot tweet geotag point data, HASC regions and centroid points for mentions of local region names such as cities or suburbs. At the street level we plot roads and building outlines for any above threshold report frequencies within a rolling 3–6 h sample window updated every 30 s. The design of the interactive visualization of Twitter streams is inspired by SensePlace2 [3], aiming to present an aggregated high-level overview of the situation to the operator on duty in real-time.

5 User Evaluation of the TweetComP1

Our evaluation methodology involves each organization’s operation room staff watching a demonstration of our prototype during a Tsunami simulation off the coast of Portugal. Feedback from end-users is elicited via semi-structured questionnaires and the data used to define performance targets for end-user acceptance of social media crisis mapping in the decision support workflow. We then compare these performance targets with published [15, 17] benchmark results for our prototype.

We asked Tsunami control room engineers from the Kandilli Observatory and Earthquake Research Institute (KOERI) and the Portuguese Institute for the Sea and Atmosphere (IPMA) to participate in our study. The semi-structured questionnaires asked users to specify acceptable update and false positive rates for crisis maps, rank incident types/social media user types in order of importance, provide preferences for crisis mapping display artefacts and rank possible information in order of importance for the decision support process in the control room.

5.1 Acceptance Criteria

Our experts considered Twitter a credible information source overall, scoring it 4/5. The most trusted types of Twitter users were government units/municipality, news agencies

and humanitarian/volunteer agencies. Lowest in the ranking were the general public. The most trusted information source was tweets from agencies such as U.S. Geological Survey (USGS) and European-Mediterranean Seismological Centre (EMSC).

The type of tweeted incident considered most useful for a crisis map were major earthquake shaking/damage, infrastructure damage reports and flood inundation reports. Opinions varied for the importance of individual cries for help and Tsunami wave sightings as these may not be indicative of specific Tsunami damage. Minor earthquake shaking reports were least useful.

The minimum crisis map update delay was 10 s during a crisis, and 60 s for normal activity. A false positive rate of 10–20 % was specified for tweet users with a low credibility (i.e. general public) and a ‘minimum possible’ false positive rate for tweets originating from more trusted sources.

5.2 Performance of the TweetComP1 Prototype with Regards Acceptance Criteria

The current TweetComP1 system processes streams of tweets in real-time and provides map updates every 30 s on modest hardware (i.e. 2.5 GHz CPU, 8 GB RAM). A 10 s update rate is achievable but some optimization of the data processing is needed from twitter crawler to final crisis map.

Our false positive rate [15, 17] for large events such as major earthquakes and Tsunami is within the range 10–13 % for the major languages like English and Spanish. There is still room for improvement for other languages with more limited training material. We need to work on lowering the false negative rate for smaller events, such as minor earthquakes and floods, since the quality of any Twitter crisis map depends on having enough people tweeting about the event in the first place.

It should be noted that general public tweeters are the least trusted but they can produce incident reports sometimes within seconds of the event itself. Multiple tweeters can raise the credibility of an incident report. We need to look at ways to allow filtering of our map layers based on the user type and data source. For example viewing less trustworthy general public reports in the very early stage of a crisis and moving to more reliable reports from official organizations as the crisis develops.

6 Conclusion

We present in this paper the TRIDEC high-level architecture for crisis mapping of ‘unconventional’ social media sensors. We are able to process real-time Twitter microblog streams and extract region, street and building level location information relating to earthquake events, Tsunami wave arrival and inundation reports. We performed a qualitative evaluation of our TweetComP1 system using Tsunami control room engineers from two organisations, showing that social media crisis maps are credible as a data source for use within a control room. For control room use an update rate of 10–60 s and false positive rate of 10–20 % (general public incident reports) is needed. These targets are within the reach of our prototype system and we are looking now at future plans to move this work into the exploitation phase.

Acknowledgements. We would like to acknowledge the feedback received from control room engineers at the Kandilli Observatory and Earthquake Research Institute (KOERI) and Portuguese Institute for the Sea and Atmosphere (IPMA). The work presented in this paper is part of the research and development in the TRIDEC project (contract no. 258723), supported by the 7th Framework Program of the European Commission.

References

1. Abel, F., Hauff, C., Houben, G.J., Stronkman, R., Tao, K.: Twitcident: fighting fire with information from social web streams. In: Proceedings of the 21st International Conference Companion on World Wide Web, pp. 305–308. ACM, April 2012
2. Hammitzsch, M., Lendholt, M., Esbri, M.Á.: User interface prototype for geospatial early warning systems—a tsunami showcase. *Nat. Hazards Earth Syst. Sci.* **12**, 555–573 (2012)
3. MacEachren, A.M., Jaiswal, A., Robinson, A.C., Pezanowski, S., Savelyev, A., Mitra, P., Blanford, J.: Senseplace2: geotwitter analytics support for situational awareness. In: 2011 IEEE Conference on Visual Analytics Science and Technology (VAST), pp. 181–190. IEEE (2011)
4. Meier, P.: New information technologies and their impact on the humanitarian sector. *Int. Rev. Red Cross* **93**(844) (2011)
5. Middleton, S.E., Sabeur, Z.A., Löwe, P., Hammitzsch, M., Tavakoli, S., Poslad, S.: Multi-disciplinary approaches to intelligently sharing large-volumes of real-time sensor data during natural disasters. *Data Sci. J.* **12**, 109–113 (2013). Special issue of the Proceedings of the 1st WDS Conference in Kyoto 2011
6. Moßgraber, J., Middleton, S., Hammitzsch, M., Poslad, S.: A distributed architecture for tsunami early warning and collaborative decision-support in crises. *Geophys. Res. Abstr.* **14**, EGU2012-8040-2 (2012)
7. Sakaki, T., Okazaki, M., Matsuo, Y.: Tweet analysis for real-time event detection and earthquake reporting system development. *IEEE Trans. Knowl. Data Eng.* **25**(4), 919–931 (2013)
8. Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake shakes twitter users: real-time event detection by social sensors. In: Proceedings of the 19th International Conference on World Wide Web, pp. 851–860. ACM, April 2010
9. Steinmetz, T., Raape, U., Teßmann, S., Strobl, C., Friedemann, M., Kukofka, T., Dech, S.: Tsunami early warning and decision support. *Nat. Hazards Earth Syst. Sci.* **10**(9), 1839–1850 (2010)
10. Vieweg, S., Hughes, A.L., Starbird, K., Palen, L.: Microblogging during two natural hazards events: what twitter may contribute to situational awareness. In: Proceedings of the 28th International Conference on Human Factors in Computing Systems, pp. 1079–1088. ACM, April 2010
11. Wächter, J., Babeyko, Y., Fleischer, J., Häner, R., Hammitzsch, M., Kloth, A., Lendholt, M.: Development of tsunami early warning systems and future challenges. *Nat. Hazards Earth Syst. Sci.* **12**, 1923–1935 (2012)
12. Wang, X., Tokarchuk, L., Cuadrado, F., Poslad, S.: Exploiting hashtags for adaptive microblog crawling. In: IEEE/ACM, International Conference on Advances in Social Networks Analysis and Mining, ASONAM 2013, Niagara Falls, Canada (2013)
13. Yin, J., Karimi, S., Robinson, B., Cameron, M.: ESA: emergency situation awareness via microbloggers. In: Proceedings of the 21st ACM International Conference on Information and Knowledge Management, pp. 2701–2703. ACM, October 2012

14. Young, J.C., Wald, D.J., Earle, P.S., Shanley, L.A.: Transforming Earthquake Detection and Science Through Citizen Seismology. Woodrow Wilson International Center for Scholars, Washington, DC (2013)
15. Zielinski, A.: Detecting natural disaster events on twitter across languages. In: IIMSS, pp. 291–301 (2013)
16. Zielinski, A., Tokarchuk, L., Middleton, S.E., Wang, X.: Social media text mining and network analysis for decision support in natural crisis management. In: ISCRAM 2013, 12–15 May 2013, Baden-Baden, Germany (2013)
17. Middleton, S.E., Middleton, L., Modafferi, S.: Real-time crisis mapping of natural disasters using social media. *IEEE Intell. Syst.* **29**(2), 9–17 (2014)
18. TRIDEC video demonstrations: YouTube. <http://www.youtube.com/user/tridecproject>

Decision Support Systems III – Impact of Decision Support
Systems for Global Environments

Euro Working Group Workshops, EWG–DSS 2013,
Thessaloniki, Greece, May 29–31, 2013, and Rome, Italy,
July 1–4, 2013, Revised Selected and Extended Papers
Dargam, F.; Hernández, J.E.; Zaraté, P.; Liu, S.; Ribeiro, R.;
Delibasić, B.; Papathanasiou, J. (Eds.)

2014, XVIII, 181 p. 65 illus., Softcover

ISBN: 978-3-319-11363-0