

Crisees: Real-Time Monitoring of Social Media Streams to Support Crisis Management

David Maxwell¹, Stefan Raue¹, Leif Azzopardi¹, Chris Johnson¹, and Sarah Oates²

¹ School of Computing Science,
² School of Social and Political Sciences,
University of Glasgow, UK

Abstract. The *Crisees* demonstrator is a service that aggregates and collects social media streams to support Crisis Management

1 Crisis Event Manager

Social media streams provide access to unprecedented amounts of information describing ongoing events as they unfold [1]. Tapping into these real-time sources about critical events can provide authorities and agencies dealing with such emergencies and crises with valuable information. Such events include the recent London riots and Egyptian uprising. For example, Spinsanti and Ostermann analysed how *Twitter*³ could provide useful information about European Forest Fires [3]. In their analysis, they found that the location of *tweets* regarding outbreaks of fire were closely correlated with the officially recorded locations of each fire. Using social media in this way has led to the idea that citizens can act as “social sensors” [2].

Whilst using such social sensors can provide valuable operational intelligence at ground level and in real-time, there are obvious and numerous problems that need to be resolved in order to manage these new sources of information effectively. These include: collecting and processing data in real-time, filtering and aggregating the content, assessing the integrity of the material, identifying cogent information, and finally both visualising and conveying the information [3].

The aim of the *Crisees* demonstrator is to provide real-time analysis of crisis events, and is designed to help support analysts monitoring such events. *Crisees* is developed upon a flexible service-oriented architecture where information is obtained from specified social media APIs via a **collector**; this information is then **filtered**, passed to an **indexer** for storage and finally **analysed** (and subsequently re-indexed if required). This processed information is then made available to web services through the **Crisees REST API**. Multithreading is used to allow for several incoming streams to be collected and processed concurrently.

The *Crisees* demonstrator is configured to use the Twitter and *YouTube*⁴ APIs to obtain content - using specific keywords and geographical data to find location-specific and relevant matches. The user interface is shown in Figure 1. Analysts, such as emergency managers, are presented with a tabbed interface

³ <http://www.twitter.com>

⁴ <http://www.youtube.com>



Fig. 1. A snapshot of the *Crisees* user interface during the recent Scottish storms, focussing on the Kingston Bridge in Glasgow.

(see item 1 in the figure), enabling them to track different parts of the same event simultaneously - or different events entirely. Events can be paused as required (2); query terms can be added or removed through the query manager (3). Sources can be enabled or disabled for each event (4) - and historical data can be viewed if required (5). A stream from the various social media sources utilised providing either historical or real-time information is displayed down the left of the interface (6) - along with a set of icons (7), each of which denote different properties of the content, such as the source, a sentiment value, and whether there is geographic information associated with the content. For example, (8) denotes the sentiment of the tweets, where values above zero denote a positive sentiment, and those below indicate a negative sentiment. The map shows the location of where information (text, video or image) has been posted from (9). Finally, any audio and visual content collected is presented at the bottom of the interface, complete with thumbnail images - if available - for easy access (10). Currently the demonstrator is being extended to include sentiment analysis of incoming textual content and to identify the geographical location of such content (if no geo-data is available). Our future plans are to add further types of analysis and to enable user feedback so that the quality and relevance of the information can be rated.

The demonstrator is hosted at: <http://www.dcs.gla.ac.uk/access/crisees>.

References

1. A. L. Hughes and L. Palen. Twitter Adoption and Use in Mass Convergence and Emergency Events. In proc. of the 6th *ISCRAM Conference*, Sweden, May 2009.
2. M. F. Goodchild. Citizens as sensors. *GeoJournal*, 69:211-221
3. L. Spinsanti and F. Ostermann. Retrieve Volunteered Geographic Information for Forest Fire. In proc. of the 2nd *Italian Information Retrieval Workshop*, Italy, January 2011.