

EAIMS: Emergency Analysis Identification and Management System

Richard McCreddie, Craig Macdonald and Iadh Ounis
{firstname.lastname}@glasgow.ac.uk
University of Glasgow, UK

ABSTRACT

Social media has great potential as a means to enable civil protection and law enforcement agencies to more effectively tackle disasters and emergencies. However, there is currently a lack of tools that enable civil protection agencies to easily make use of social media. The Emergency Analysis Identification and Management System (EAIMS) is a prototype service that provides real-time detection of emergency events, related information finding and credibility analysis tools for use over social media during emergencies. This system exploits machine learning over data gathered from past emergencies and disasters to build effective models for identifying new events as they occur, tracking developments within those events and analyzing those developments for the purposes of enhancing the decision making processes of emergency response agencies.

1. INTRODUCTION

Social networks and social media have become some of the most prominent technologies currently in use today. Indeed, the number of different social media platforms continues to grow, accompanied by a steady increase in the number of users that either actively post or otherwise consume content shared using these platforms [5]. Furthermore, as these platforms continue to proliferate, the reach of these social networks and social media increases, particularly in rural areas where initial adoption was low.

Importantly, social media provide a wealth of real-time information, which, if used effectively can provide insights about what is happening in the world. In particular, during emergency situations such as natural disasters, information shared on social media has been shown to be valuable [19]. Indeed, experts in a variety of domains have developed tools to use social media to identify and track events [1, 12, 14, 16]. For instance, the Crisis Tracker system [16], uses a manually created set of search terms to crawl crisis-related content, which is then clustered automatically and manually annotated by volunteers. Such crisis tracking and analysis tools can collect, analyze and aggregate information from people

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '16, July 17 - 21, 2016, Pisa, Italy

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4069-4/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2911451.2911460>

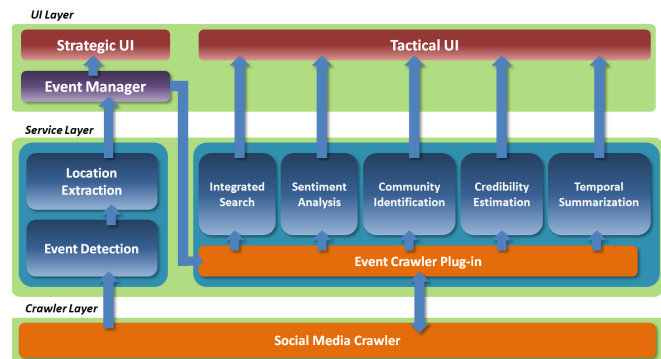


Figure 1: EAIMS architecture overview

involved in an event, enabling both the general public and emergency response agencies to monitor what is happening in real-time.

In recent years, there have been notable technological advances in fields such as machine learning for sentiment and behavioral analysis [11], real-time event detection in social streams [10], temporal event summarization [8] and credibility estimation of social content [4]. These new technologies have the potential to provide richer information about crisis events based on social media content. However, to-date, there are few crisis tracking and analysis tools that make use of them.

We present a new crisis tracking and analysis toolkit, named: EAIMS (Emergency Analysis Identification and Management System). The goal of EAIMS is to provide an integrated toolkit for automatically identifying, tracking, summarizing and managing information about emergency events derived from social media. Unlike most crisis tracking and analysis systems, EAIMS provides automatic real-time detection of new emergency events. Furthermore, once an event is detected, EAIMS provides a rich suite of automatic tracking and analysis tools, namely: automatic timeline construction, targeted sentiment analysis, user-community identification and information credibility estimation.

EAIMS is designed for use by emergency management staff stationed at their operations center. In particular, we envisage the system being used by two different groups of people, representing strategic use across different events and tactical use within a single event. First, at a strategic level, the system can be used by emergency monitoring staff responsible for identifying new crisis events and initiating first response procedures. In particular, automatic event detection within EAIMS can automatically alert users to potential crisis events that they can then verify through existing

channels. Using automatic tools to aid in crisis event identification has the potential to enable faster response times to unexpected crisis events. Second, at a tactical level, we envisage EAIMS being used by emergency analysts both during and after an event. Users can leverage the automatic tracking and analysis tools provided by EAIMS to more effectively monitor the event itself - using social media users as sensors. Furthermore, these tools can also be used to assess how the agency’s response to the emergency is perceived - facilitating more informed communication with the public.

2. EAIMS ARCHITECTURE

A broad overview of the EAIMS architecture is illustrated in Figure 1. As we can see from Figure 1, EAIMS is comprised of three main architecture layers. The lowest layer is the crawling layer, which is responsible for collecting general, as well as event-specific social media content, which will be processed by the upper layers. The second architecture layer is the service layer. This layer is comprised of a series of plug-and-play services that provide the end-user with crisis tracking and analysis capabilities. In line with the vision of EAIMS, these services are divided into strategic and tactical services. Strategic services provide functionality to users working at the strategic level, i.e. when tracking different event reports across multiple regions. Meanwhile, tactical services provide functionality when working at the level of a single event. The tactical services also include the Event Crawler Plug-in, which is responsible for performing targeted collection of content for a particular event. The third layer of EAIMS is the UI (user interface) layer. This layer is responsible for managing and visualizing the strategic and tactical information provided by the service layer to the user. In particular, this layer contains three main components. First, the event manager component is responsible for maintaining the list of candidate events that have been identified. Second, the UI layer contains the strategic UI, that visualizes the different candidate events that have been detected on a map for the user. Finally, the tactical UI is responsible for visualizing the outputs from each of the tactical services for a selected event. We discuss the motivations and implementation of the tactical and strategic services/UI below.

3. STRATEGIC SERVICES AND UI

The aim of the strategic services is to provide automatic detection and annotation of new candidate events to the user. To provide this functionality, there are two notable components that are needed: 1) automatic identification of events of interest to the user; and 2) association of those events with locations. The output of these components are visualized within EAIMS’s strategic UI, as illustrated in Figure 2, where each event is represented by a pin on the map. Clicking on a pin displays the earliest tweet identified for the event on the right hand side of the UI. Additionally, a map overlay is displayed which provides event controls. In particular, it enables the user to initialize tracking for the event¹ and to switch to the tactical view for this event. We discuss how we tackle the event detection and location identification components below.

¹This triggers targeted crawling of social media posts based on tracking terms (named entities and #hashtags) extracted from each event cluster.

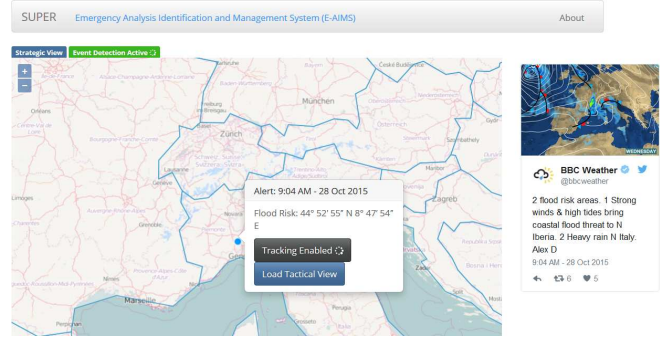


Figure 2: EAIMS: Strategic UI.

3.1 Event Detection

The aim of automatic event detection is to identify the first instance of an article/post that discusses a new emerging topic, e.g. a news article representing a new event. This is normally achieved by incremental clustering algorithms. However, such clustering approaches have been shown to be too computationally intensive to be applied to high volume streams, such as tweets. Hence subsequent approaches have focused on constant-time approximations to perform text clustering [13].

To enable EAIMS to identify events over high volume tweet streams, we use a recent distributed streaming event detection approach based on locality sensitive hashing [10]. Under this approach, when each post arrives, multiple locality sensitive hashing functions are used to generate hashes for that post. A locality sensitive hashing function is one where textually similar posts will receive similar hash codes with a high probability. Hash collisions with previous posts observed are used to determine whether a similar topic has been observed before. Clusters of similar posts are maintained, which represent different candidate events. Once a candidate event cluster reaches a predefined size (5 posts in our case), that cluster is emitted as a new event. This approach is both constant time and can be distributed over multiple machines to achieve parallelism for large events.

The above approach generates a relatively high volume of candidate events (hundreds per hour) over time, and are not specific to the interests of the user. Hence, for each cluster emitted as a candidate event, we apply an additional filtering step. In particular, we apply an SVM text classifier to filter out those events that are not about topics of interest to the user. We currently support two main use-cases, namely: flash flooding and large public demonstrations, with an associated event classifier for each.

3.2 Location Detection

Once a relevant event has been identified it is vital to identify a location for that event. Indeed, most response agencies are responsible only for particular geographical regions, hence it is important to only highlight events relevant to the user. A candidate event is represented by a cluster of social media posts, such as tweets. However, only about 2-3% of all tweets are geo-located [3]. Hence, there is a good chance that none of the posts within a cluster will have geolocations.

To tackle this issue, we use the Carmen location attribution approach [6]. Carmen aims to assign a location to a tweet from a database of known locations. To do so, it combines three types of information sourced from each tweet,

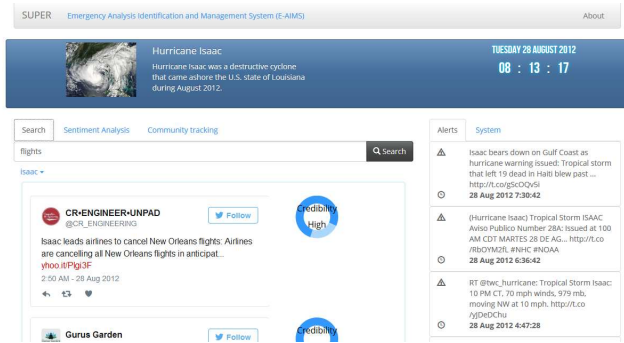


Figure 3: EAIMS: Tactical UI - Search Tab.

namely: place information encoded within the tweet; geo-coordinate information for the small number of tweets that have them; and location(s) extracted from parsing the user profile for the tweet author. An evaluation of Carmen’s performance over a tweet dataset indicates that 22% of posts can be geo-located [6], although the granularity of the identification varies (e.g. in many cases location can only be narrowed down to the country of interest). This is often sufficient to identify the location for one or more posts in our event clusters.

4. TACTICAL SERVICES AND UI

In contrast to the strategic services, which identify and visualize multiple events, the aim of the tactical services are to provide a series of tools to enable a deeper analysis of a single event. When a user initiates tracking of an event within the strategic UI, the cluster of social media posts for that event are sent to the Event Crawler Plug-in. This plug-in then formulates a new tracking query to collect content specifically for that event. To do so, it extracts all named entities and hashtags from the event cluster. These query terms are passed to the crawler layer as a query.

The tactical UI is illustrated in Figure 3. As we can see from Figure 3, event-related information is displayed along the top of the page (this can be customized by the user). Below the event information is a series of tabs that provide access to the search, sentiment analysis and community identification functionalities (the search tab is selected in Figure 3). Down the right hand side, the output of the timeline generation service is displayed, showing the most recent updates automatically extracted from the social media content stream. We discuss each of these components in more detail below.

4.1 Integrated Search

The first of the tactical services is integrated search. This service provides dynamic indexing and retrieval of social media content crawled for the event. This service enables information exploration and discovery by emergency analysis staff during, as well as after the event. To provide social search capabilities, we build on top of the real-time indexing and retrieval functionality of the Terrier search engine platform [7]. In particular, using Terrier v4.1, we deploy a real-time distributed index using a temporal document-partitioning scheme. Under this setup, multiple mini-indices are constructed, each containing documents from a particular time period. For an emergency management scenario, this partitioning scheme has two major advantages. First, since many search requests will be time-orientated, we can

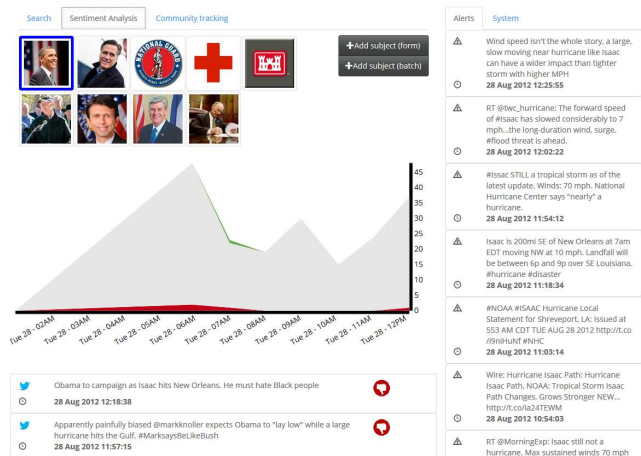


Figure 4: EAIMS: Tactical UI - Sentiment Tab.

limit the number of indices that have to be searched based on the user’s search constraints. Second, for cases where the user wants results ranked by time, we initially need to only search the most recent indices to supply the most recent results. Indices containing older content can then be searched on-demand as the user scrolls down the search result page.

4.2 Credibility Estimation

When working with social media derived information, it is important to be able to judge the veracity/credibility of that information, if that information is to be actionable. To support this, EAIMS includes a credibility estimation service that works in parallel with the integrated search service. In particular, for each social media post ranked within a search engine result page, a credibility score is provided. To score the credibility of a post, we use the approach proposed by Castillo et al. [4], which uses a machine-learned regression model to score each post. This approach uses three main types of features to distinguish the credibility of posts, namely: text features (length of text, number of URLs/mentions/hashtags contained); information about the users who authored them (registration age, number of followers); and the general topics they refer to (such as the fraction of tweets in a topic containing URLs).

4.3 Sentiment Analysis

One of the areas where emergency management agencies see value in social media is as a means to track public opinions regarding the emergency response, i.e. as a means to judge how well the public think they are doing [9]. Prior work on sentiment analysis in social media during crises has applied well-known techniques for overall sentiment detection in posts [11]. However, we argue that sentiment analysis of the overall post might not always be suitable, as it may miss the presence of more targeted sentiments, e.g. about the people and organizations involved (which we refer to as sentiment targets). Hence, within EAIMS, we use a machine learned sentiment polarity text classification model designed to identify sentiment towards one or more named entities. In particular, we use crowdsourced sentiment labels for entity-tweet pairs to train a the classification model, using unigram and bigrams as features. The output of the sentiment analysis component is shown in Figure 4. In particular, the currently defined entities being tracked are displayed along the top of the UI, enabling the user to select which to view sen-

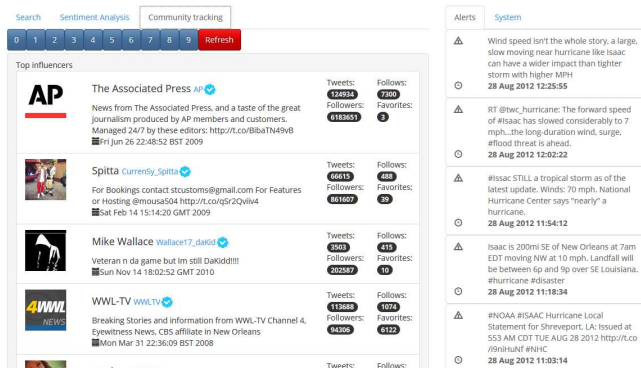


Figure 5: EAIMS : Tactical UI - Community Tab

timent for. Sentiment is then visualized in two ways. First, as a color-coded graph showing the sentiment polarity distribution for the currently selected entity over time. Second, a listing of recent posts expressing sentiment about that entity. The user can add new entities to track at any time.

4.4 Community Identification

Emergency management agencies typically maintain a social media presence as a communication channel with the general public. To make effective use of this channel, it is important to identify groups of stakeholders within the social network that can be messaged, such as local response volunteers or news agencies covering the event. To support this use-case, EAIMS provides an automatic community identification functionality. In particular, we use an approach based on topic modeling to identify communities as determined by the users' interaction patterns within the social network. User communities are visualized within the community tab of the tactical UI, as shown in Figure 5. In particular, the user can switch between any of the communities identified using the buttons displayed along the top of the UI. For the currently selected community, the top influencers are displayed, ranked by the number of social media connections they have within the network.

4.5 Timeline Generation

Finally, EAIMS provides a real-time timeline generation and notification functionality. In particular, the user can specify a series of information types to track, e.g. messages related to deaths or injuries. The timeline generation component monitors the social media content crawled for the event being tracked and extracts posts related to each of the specified information types to return to the user as updates. To achieve this, we build on recent works in the field of temporal summarization [8] to find relevant posts as they are crawled and then then apply explicit diversification with respect to the specified information types using the xQuAD framework [17]. In this way, we score a post based its relevance to the event, its coverage of information about the event, and its novelty with respect to specified information types. Posts that exceed a score threshold are displayed within the tactical UI as updates, as illustrated in Figure 3.

5. CONCLUSIONS

In this paper, we presented the Emergency Analysis Identification and Management System (EAIMS), a prototype service that provides real-time detection, monitoring and

analysis tools for use over social media during emergencies. It combines state-of-the-art event detection, sentiment analysis, timeline generation, credibility analysis and community identification technologies to enable emergency monitoring staff to better leverage social media during crises. This prototype system is currently being deployed for flooding and public security use-cases within the FP7 SUPER Project.

6. ACKNOWLEDGMENTS

This work has been carried out in the scope of the EC co-funded SUPER Project (FP7-606853).

7. REFERENCES

- [1] F. Abel, C. Hauff, G.-J. Houben, R. Stronkman and K. Tao. Twitcident: Fighting Fire with Information from Social Web Streams. In *Proc. of WWW*, 2012.
- [2] F. Aisopos, G. Papadakis, and T. Varvarigou. Sentiment analysis of social media content using n-gram graphs. In *Proc. of WSM*, 2011.
- [3] S. H. Burton, K. W. Tanner, C. G. Giraud-Carrier and J. H. West. "Right time, right place" health communication on Twitter: Value and accuracy of location information. *Journal of Medical Internet Research*, 14(6), 2012.
- [4] C. Castillo, M. Mendoza and B. Poblete. Information credibility on Twitter In *Proc. of WWW*, 2011.
- [5] A. Datta, M. Dikaiakos, S. Haridi, and L. Iftode. Infrastructures for online social networking services [guest editorial]. *IEEE Internet Computing*, 16(3):10–12, 2012.
- [6] M. Dredze, M. J. Paul, S. Bergsma and H. Tran. Carmen: A Twitter Geolocation System with Applications to Public Health. In *Proc. of HIAI*, 2013.
- [7] C. Macdonald, R. McCreadie, R. L. Santos, and I. Ounis. From puppy to maturity: Experiences in developing terrier. *Proc. of OSIR*, 2012.
- [8] R. McCreadie, C. Macdonald, and I. Ounis. Incremental update summarization: Adaptive sentence selection based on prevalence and novelty. In *Proc. of CIKM*, 2014.
- [9] R. McCreadie *et al.* SUPER: Towards the use of Social Sensors for Security Assessments and Proactive Management of Emergencies. In *Proc. of SWDM*, 2015.
- [10] R. McCreadie, C. Macdonald, I. Ounis, M. Osborne, and S. Petrovic. Scalable distributed event detection for Twitter. In *Proc. of IEEE Big Data*, 2013.
- [11] S. Narr, M. Hülphenhaus, and S. Albrayrak. Language-independent Twitter sentiment analysis. In *Proc. of KDML*, 2012.
- [12] M. Osborne *et al.* Real-Time Detection, Tracking, and Monitoring of Automatically Discovered Events in Social Media. In *Proc. of ACL*, 2014.
- [13] S. Petrovic, M. Osborne and V. Lavrenko. Streaming first story detection with application to Twitter. In *Proc. of NAAACL*, 2010.
- [14] H. Purohit and A. P. Sheth. Twitris v3: From Citizen Sensing to Analysis, Coordination and Action. In *Proc. of ICWSM*, 2013.
- [15] J.-C. Puras and C. A. Iglesias. Disasters2.0. Application of Web 2.0 technologies in emergency situations. In *Proc. of ISCRAM*, 2009.
- [16] J. Rogstadius, M. Vukovic, C. A. Teixeira, V. Kostakos, E. Karapanos, J. A. Laredo. CrisisTracker: Crowdsourced social media curation for disaster awareness. *IBM Journal of Research and Development*, 57(5), 2013.
- [17] R. L. T. Santos, C. Macdonald, and I. Ounis. Search Result Diversification. *Foundations and Trends in Information Retrieval*, 2015.
- [18] H. Sayyadi, M. Hurst and A. Maykov. Event Detection and Tracking in Social Streams. In *Proc. of ICWSM*, 2009.
- [19] J. Sutton, L. Palen, and I. Shklovski. Backchannels on the front lines: Emergent uses of social media in the 2007 southern California wildfires. In *Proc. of ISCRAM*, 2008.