

Analyzing Crowd-Sourced Information and Social Media for Crisis Management

Item Type	Book chapter
Authors	Andrews, S.;Day, T.;Domdouzis, K.;Hirsch, L.;Lefticaru, Raluca;Orphanides, C.
Citation	Andrews S, Day T, Domdouzis K et al (2017) Analyzing Crowd-Sourced Information and Social Media for Crisis Management. In: Akhgar B, Staniforth A and Waddington D (Eds) Application of Social Media in Crisis Management. Transactions on Computational Science and Computational Intelligence. New York: Springer International Publishing.
DOI	https://doi.org/10.1007/978-3-319-52419-1_6
Publisher	Springer International Publishing
Rights	© 2017 Springer. Reproduced in accordance with the publisher's self-archiving policy. The final publication is available at Springer via https://doi.org/10.1007/978-3-319-52419-1_6
Download date	2025-08-07 16:43:45
Link to Item	http://hdl.handle.net/10454/17662

Chapter 1

Analyzing Crowd-Sourced Information and Social Media for Crisis Management

Simon Andrews, Tony Day, Konstantinos Domdouzis, Laurence Hirsch, Raluca Lefticaru, and Constantinos Orphanides

Abstract The analysis of potentially large volumes of crowd-sourced and social media data is central to meeting the requirements of the Athena project. Here, we discuss the various stages of the pipeline process we have developed, including acquisition of the data, analysis, aggregation, filtering and structuring. We highlight the challenges involved when working with unstructured, noisy data from sources such as Twitter, and describe the crisis taxonomies that have been developed to support the tasks and enable concept extraction. State of the art technology such as formal concept analysis and machine learning is used to create a range of capabilities including concept drill down, sentiment analysis, credibility assessment and assignment of priority. We present an evaluation of results obtained from a set of tweets which emerged from the Colorado wild fires of 2012.

1.1 Introduction

Social media data is rapidly changing the way emergency data is created and distributed during a crisis. The proliferation of mobile devices together with ubiquitous tools for online dissemination means that a huge pool of dynamic information is of increasing importance to emergency personnel in a crisis situation. Crowdsourcing can provide the fastest access to localised information and a number of studies have suggested that emergency response can be improved by employing local community knowledge to provide aggregated situational awareness emerging in real time [9, 18].

However, the challenges of using such data has also been highlighted and indeed it has been identified that social media data can be a source of misinformation, propaganda and rumour, both intentional and unintentional [2, 16]. The obvious risks

S. Andrews, T. Day, K. Domdouzis, L. Hirsch, R. Lefticaru, C. Orphanides
CENTRIC, Sheffield Hallam University, 153 Arundel Street, Sheffield, S1 2NU, UK,
e-mail: {S.Andrews, T.Day, K.Domdouzis, L.Hirsch, R.Lefticaru, C.Orphanides}@shu.ac.uk

associated with using an unregulated stream of information implies that assessing the reliability of crowdsourced data has emerged as a crucial task [16]. Gao [11] also points out that the level of messaging in a disaster will be so high that meaningful filtering, aggregation, categorisation and summarisation will be essential capabilities if we are to make effective use of the data. Narvaez [14] further argues that appropriate organisation of social network information is the key to providing support in terms of ground action.

The Athena project aims to harness available data from a variety of sources as a way of extracting actionable intelligence to the public and first responders. This is achieved by developing systems for the searching, acquisition, aggregation, filtering and presentation of knowledge from social media and crowd-sourced data to support crisis management. In this chapter, we describe the technology employed at each of these stages. The raw data comes from a variety of sources and is often challenging to process and analyse, as it is by its nature loosely structured and noisy. Real examples are useful in development and evaluation of the tools and also in explaining how a system works. We therefore refer to a particular data set, namely a set of tweets transmitted during the Colorado wild fires in 2012 [17], to demonstrate aspects of the system, such as development of the crisis taxonomy or delivering credibility assessment.

We include an explanation of how we use and combine various commercial and freely available open source software tools to support the tasks. We also describe a powerful visual instrument, enabled by formal concept analysis, which we have developed. The tool allows an analyst to drill down into the concept hierarchy which has been established in previous stages.

1.2 Overview

In section 1.3 we begin by describing how we obtain and structure the data. We include an explanation of the information acquisition process and the data processing pipeline that has been established. Assessing the credibility and priority of any crisis information is obviously a critical task. In section 1.4 we discuss the challenges involved and the process employed in Athena, together with informative examples from the Twitter data. In recent years the research community has been very active in the area of sentiment analysis and it is clear that establishing the sentiment, and any changes in sentiment from large volumes of crowd sourced information can be extremely useful to analysts and first responders. In section 1.5 we discuss how we track and report sentiment information and include a discussion around the supporting taxonomy. In section 1.6 we review the use of formal concept analysis in the data aggregation process and in section 1.7 we discuss filtering and searching.

1.3 Obtaining structured crisis data

The Athena Project has developed data collection and pre-processing tools for the real-time acquisition of data, such as video and voice, from the social media (both generally and from the dedicated Athena crisis pages) and from mobile/high-tech devices. The Athena Project is characterized by an Information Processing Process that must collect and provide robust, high quality data to the rest of the components. This process is based on two stages. The first stage is the acquisition and pre-processing of data. The second stage is based on data analysis and aggregation.

In Athena, information acquisition is the process of obtaining crisis data from a number of sources including social media websites, such as Facebook and Twitter. Athena will then use a number of filtering processes and crisis taxonomies in order to focus on information that can be formulated to valuable intelligence. Twitter is considered a valuable source of crisis data. A tweet can include a variety of information, such as text, image, video, links and hashtags. There is also a significant amount of metadata that characterizes each tweet. These include geo-location, author name and Twitter handle, author pre-set location, timestamp, number of retweets, number of favourites, list of hashtags, list of links and other users mentioned in the tweet. Athena also focuses on the collection of data from stories posted to public pages and the responses of the public to these stories. Specifically, the analysis of the Facebook data will concentrate on the comments made on the dedicated Athena crisis pages. As both Twitter and Facebook often link to external websites, it may be useful to follow these links and crawl the content of the page they have linked to. In some cases, this will be news reports containing standard information but in other cases, these may contain live reports posted by journalists that are not detected via the Twitter and Facebook crawlers or posts by individual bloggers.

Athena collects crisis data also through the Athena Mobile Application. There are three levels of reports that can be sent via the mobile application. These are the public user reports, the trusted user tier-2 reports that include reports from utilities controllers, official volunteers, and professionals from local resilience forums and other identified credible community voices, and the trusted user tier-1 reports that include reports from first responders and from the operational, tactical and strategic command of the Police. Data from the Athena Mobile Application will include some but not necessarily all of the following: Report text, Report Image, Report video, Report audio, Report geo-location, Report tag, Report timestamp and User Type.

Any information which is acquired by Athena has to be filtered through the SAS Information Retrieval Studio (IRS). The SAS IRS provides a data processing pipeline that connects each of the SAS components together. It involves five main stages. The first stage is the data import stage. This stage is related to the setting up of crawlers for Athena to crawl and by inputting relevant search terms to extract important data. The second stage is filtering. In this case, any incoming information is evaluated for its significance based on certain criteria.

The third stage is the categorisation, context and contextual extraction phase. This stage identifies keywords, categories and concepts in each of the post. The

fourth stage is related to the identification of any sentiment expressed in each of the posts while the fifth stage involves the exporting of the data.



Fig. 1.1 Information processing pipeline

SAS provides a number of crawlers which integrate with the content categorization studio and the sentiment analysis studio. The SAS Information Retrieval Studio includes a web crawler which traverses the web and follows one link to the next, a file crawler which crawls each file in a directory, a feed crawler which crawls RSS or Atom feeds, a Google crawler, a Facebook crawler and a Twitter crawler.

Each crawler is written as a simple Python file that can be modified based on the requirements of a crisis situation. In this case, new data sources could be added. For the requirements of Athena, the Twitter crawler has been re-programmed so that it returns the geo-location of tweets (if they are included as part of the tweets) and the Facebook crawler has been programmatically updated so that both the public timeline and the individual pages to be crawled.

The SAS IRS pipeline server receives documents from the crawlers. Each document passes through filtering, categorisation and sentiment analysis and then onto the export process. The export process can be adjusted in such a way so that specific documents and specific data fields are exported. Categorisation in Athena is based on two methods. In the first method, the document does not match any of the concepts. In the second method, a taxonomy is developed that identifies documents that provide crisis comments without providing crisis information. Athena uses the SAS Content Categorization Studio in order to extract categories, concepts, and contexts from the data. The SAS Content Categorization Studio is responsible for the categorization (or provision of labels) to documents, the extraction of concepts or in other words, the identification of particular characteristics in a document and the identification of relationships between concepts (contextual extraction).

Categorisation is the process of analysing a document's content and applying a category to it based on a set of pre-defined conditions. In Athena, documents can belong to a number of categories if there is match between the content of the documents and the rules from different categories. The SAS Content Categorization Studio offers automatic rule generation, rule writing and statistical categorisation. Categories are built by constructing a set of Boolean rules that is applied to the incoming documents. By connecting multiple Boolean rules, rules can be made more complex.

A crisis taxonomy has been developed in order to categorise each collected post into one or more categories. The list of possible crisis events has been categorised into seven types: Attack, Crash, Hazard, Health, Natural Disaster, Other, Public Order Incidents, Terrorism. Each of these is sub-divided to further categories. For ex-

ample, an attack can be categorized in bomb attacks, hostages, killings, knife attacks, lone wolf attacks, shootings, and suicide bombs.

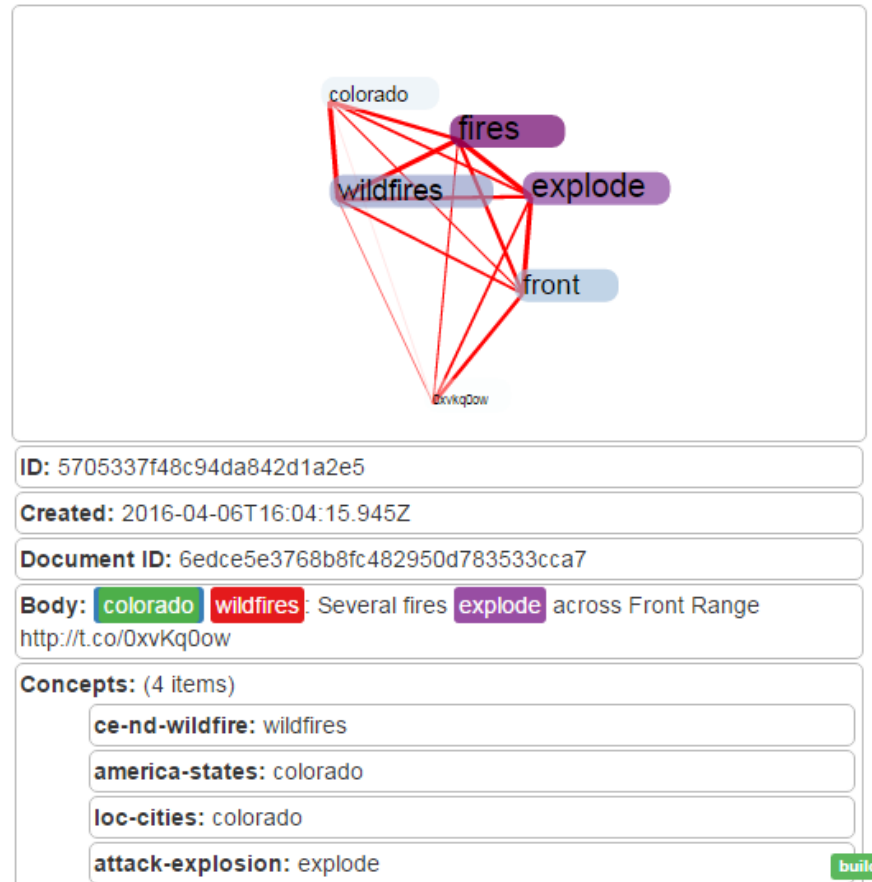


Fig. 1.2 Example of identified concepts in a tweet

Concept and contextual extraction is the process of identifying particular features included in the document. These features may be related to keywords, particular sentence constructs or even relationships between these features. An example of Concept Extraction can be shown in Fig. 1.2. In this case, the content of a document may belong to different categories if their content matches rules in a number of different categories. Categories are built based on a set of Boolean rules that are applied to the documents that have to be categorised. The complexity of Boolean rules depends on whether multiple Boolean rules are combined together in order for a number of possible scenarios to be covered. For the Athena project, a crisis taxonomy has been developed which is used to categorise each post in one or more

categories. The concept extraction is used to extract more specific details from the posts. The list of possible crisis events is roughly divided in to seven main types of classification. These incidents are then sub-divided into further categories as shown in Figure 1.3.




Fig. 1.3 Crisis categorisation taxonomy

Concept extraction is very helpful in the generation of useful information from social media data because of the wide variety of terms and spellings that refer to the same term. The reduction of these terms to a single term ensures that there is consistency in the analysis of the terminology. Contextual extraction is similar to concept extraction, however the data extracted in this case are more complex in nature and they are referred to as ‘facts’. Also, contextual extraction can make use of concepts within the rules it uses.

The process terminates with the output of XML versions of the scanned documents, containing metadata of the original documents as distinct XML elements, as well as all identified entities extracted from the conceptual and contextual extraction processes, as explained in the paragraphs above. Finally, through a process of data booleanization and discretization [3], the data are transformed into formal contexts, making the data accessible by the knowledge discovery and intuitive, conceptual visualization techniques [4] of Formal Concept Analysis (FCA). For an overview of FCA see section 1.6.

1.4 Assessing credibility and priority

According to the Oxford dictionary, *credibility* is “the quality of being trusted and believed in, the quality of being convincing or believable”. One objective of Athena project is to develop tools and techniques “to automatically assess information/information sources for credibility and maliciousness” and in order to achieve

this and further prioritize the information received, we are focused on two targets: detecting *credibility* and in the same time providing *situational awareness* information, in the context of a possible emergency crisis.

The overall information credibility is calculated based on knowledge of: (i) *the information provider*, when the report author is known, for example if he/she is a registered key user, an official volunteer or first responder; (ii) *the text*, on which classification techniques are applied to assess the content credibility and informativeness; and (iii) *the context*, meaning the content and corroboration of the information: more messages reporting the same incident, previously reported and already validated, will contribute to increasing the credibility and automatic validation of new incoming messages on the same topic.

In the following, as we are working with data received from different sources and which was communicated via different channels and contexts, we will present our approaches to automate the detection of credible and informative crises related messages. As detailed in Section 1.3, the information received by Athena Crisis Command & Control Intelligence Dashboard (CCCID), which has been described in [8], has different sources: the Athena mobile application, but also social media sources such as Twitter, Facebook, or RSS feeds.

For assessing the credibility of data coming from these two main streams, we have developed slightly different approaches, depending on the information source. First, the messages received from the Athena mobile app belong to the following categories: (a) *general public reports*, for example from anonymous users; (b) *trusted user tier-2 reports*, e.g. from official volunteers, professionals, and (c) *trusted user tier-1 reports*, received from first responders, operational, tactical and strategic command of Police. The messages and reports of type (b), (c) have a higher importance during the crisis, they could potentially present useful details, warnings or important actions to be taken, so they have a higher priority.

All the general public messages received from the mobile app (potentially from anonymous users) will be evaluated by human operators and they can be validated or rejected using Athena CCCID. A help message will be validated if it seems credible or genuine, otherwise it can be rejected, for example in case of people misusing the application, sending any kind of non-emergency or not relevant text, with or without malicious intentions. A semi-supervised machine learning algorithm will assist the human operators in the message classification, suggesting if the message should be validated or rejected, in case it does not seem to be a genuine message, sent from a user in an emergency situation. The automatic label proposed by the algorithm will not be definitive and can be modified by the operators. If a message is incorrectly classified by the algorithm, then a human operator can change its classification. After such a modification the labelled message will be added to the training set of the machine learning algorithm, this one being able adjust itself in the future and have a better prediction next time when receiving a similar example.

Apart from the data received from the Athena mobile app, there are also messages or documents retrieved from social media, for example from crawling Twitter or Facebook, which have been detected as possibly describing an emerging crisis situation. For this second case we propose an automatic classification, using super-

vised machine learning techniques, the methodology to be followed is detailed in Section 1.4.1.

Apart from assessing the credibility of the received messages, another important task is to automatically prioritize them. To the best of our knowledge, this is a new research area and the few reports that have emerged are concerned with crisis situations generated by floods. The proposed method is a geographical prioritization of social network messages using sensor data streams [5]. The approach is based on analysing previous data and observing the correlation between the number of Twitter messages near flood affected areas, which is much higher than in other areas [1, 5], and suggesting the combination of geographical data from the gauging stations with social media to prioritize and improve situational awareness during floods.

For the moment the Athena CCCID has a prioritization mechanism based on users tiers - different levels of trust within the Athena mobile application. However, this new approach of combining geodata with social media could be very useful when sensor data streams from the authorities are available and it should be investigated in other cases of natural disasters, e.g. earthquakes.

1.4.1 Credibility assessment of Twitter messages

There has been an increasing trend in the last decade for people to report ongoing crisis situations via Twitter, which could contribute to situational awareness and could eventually help the Law-Enforcement Agencies (LEAs), police and first responders. However, the huge amount of messages transmitted via Twitter or other social media makes it impossible for human operators to manually cull and extract relevant information in an emergency situation. Consequently, Natural Language Processing (NLP) and Machine Learning (ML) have emerged as key techniques for automating the extraction of situational awareness information broadcasted via social media [7, 19, 21].

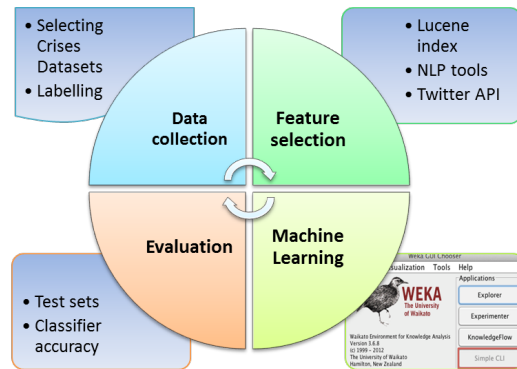
This justifies the development of appropriate classifiers for detecting credible messages and this problem has been studied by many researchers in an *off-line* (or *post-hoc*) setting, using data gathered from previous high impact situation to train the classifiers.

The majority of researchers have focused on analysing Twitter data obtained via Twitter API, which seems to be easier to access compared to Facebook data which is subject to different privacy rules. Datasets of tweets transmitted during high impact crisis, have been collected, downloaded and labelled in order to serve as training and test sets for the classifiers. An example of tweets transmitted during the Colorado fires in 2012 is given in Table 1.1, the keywords which might relate them to the crises are emphasized, however the personal data, such as name or usernames, has been hidden from ethical reasons. The table contains different tweets which were labelled in the study [17] into different categories or classes such as: related and informative; not related; not applicable; related but not informative.

Table 1.1 Sample of tweets posted during the Colorado fires (2012) and their classification

Tweet Text	Informativeness
***** and her husband had to leave their horses and evacuate immediately from the #HighParkFire. *****	Related and informative
Thousands evacuated as Colorado wildfire nears http://t.co/1jtGu6Bb	Related and informative
RT @*****: Hey Colorado Springs radio stations - Let's go ahead and remove Adele's Set Fire to the Rain from our rotation #JustSa ...	Not related
RT @*****: Lord, protect those in the path of the fires in Colorado. Please, send your Divine extinguisher for those fires!	Related but not informative

In Fig. 1.4 we present an overview of our research approach, aiming to apply ML (more precisely supervised learning) and NLP techniques in order to build a classifier to assess the credibility and information awareness of social media messages sent during crises situations. The first step, data collection, consists in obtaining the appropriate datasets and labelling them¹, this meaning annotating each record with the class to which it belongs. The collection we considered for training consists of a labelled dataset of 1200 tweets extracted from those sent during the Colorado wildfires (2012), which is presented in [17].

**Fig. 1.4** Process overview

In the machine learning process, the simple text is not enough for training; that is why for each tweet other *features* are extracted or computed, representing *attributes*

¹ Several labelled data collections, including the Colorado wildfires dataset used in this chapter, have been download from <http://crisislex.org/data-collections.html>

which might be relevant for establishing the credibility and message informativeness. The features computed for each tweet could be grouped into different scopes as in [6]:

- *Message related*: length of the text of the tweet (in characters); number of words; fraction of capital letters in the tweet; number of URLs contained on a tweet; mentions a user (e.g. @username); includes a hashtag (e.g. #HighParkFire);
- *User related*: registration age (time passed since the author registered his/her account, in days); statuses count (number of tweets at posting time); number of people following this author at posting time; number of friends (number of people this author is following at posting time); if the user has a verified account;
- *Topic*: number of tweets, average length;
- *Propagation*: number of re-tweets, degree of the root in a propagation tree etc.

The list above is not exhaustive, there are authors using larger sets, containing for example 45 features [12], in order to determine the credibility of a tweet. There are also studies showing which features are the better indicators of credibility and occur more often in data describing emergency situations [15]. For a detailed survey of existing research and tools developed for processing social media messages in mass emergency [13] can be consulted.

However, for privacy issues and other concerns regarding personal data we did not employ any user related features (such as number of friends / followers on Twitter) and we focused mainly on the message related attributes. It is worth mentioning that in the feature selection we have also considered the results provided by SAS Sentiment Analysis Studio, which provides classification into *positive*, *negative*, *neutral*, or *unclassified* as detailed in Section 1.5 and the number of identified concepts (see Fig. 1.2) and crisis related concepts which have been extracted from the text, according to the taxonomy presented in Fig. 1.3.

After processing the appropriate features for each record in the training dataset, the third stage of the process is the machine learning and here we have been experimenting with different ML algorithms, which are implemented in WEKA software [20], e.g. AdaBoost, naive Bayes, Bayesian Networks, IBk, decision trees like J48, decision table etc. Currently we are performing comparisons between these algorithms, in order to choose the ones which are more appropriate for our problem, one particular algorithm providing good performance for the current dataset is J48, an algorithm used to generate a decision tree.

The forth step consists in evaluating the classifier which has been developed against new data sets or against the same data set, but on examples which have not been used at training, the most common approach is using a 10-fold cross validation strategy. The approach we have presented is currently in the evaluation phase and after this it will be integrated in the Athena tools. Currently the accuracy of the classifier is around 80%, however it could be improved if other features, such as those user related, would be taking into account.

1.5 Sentiment analysis

Sentiment analysis is used to categorize and classify the opinions and sentiments expressed in the Colorado wildfire Twitter dataset. We have used SAS Sentiment Analysis Studio to capture the *polarity* of the text: whether the expressed opinion in a tweet, a particular sentence of the tweet, or an entity feature/aspect is positive, negative, or neutral. The following polarity classes were implemented:

- *positive* – a positive sentiment has been expressed
- *negative* – a negative sentiment has been expressed
- *neutral* – a neutral sentiment has been expressed
- *unclassified* – the sentiment expressed does not fall in any of the defined polarity classes

The polarity of each document is measured at both the overall document level (i.e. expressed towards the Colorado wildfires crisis event), and, when applicable, at the specific *feature* level (i.e. sentiment explicitly expressed towards an entity involved in the crisis event). This is possible through the creation of multi-level taxonomies to assess sentiment, as explained in the sub-section below.

1.5.1 Sentiment taxonomy

The Athena sentiment taxonomy utilizes a hybrid model, comprising of both statistical methods and predefined sentiment vocabularies, as well as handcrafted rules, custom-tailored to crisis events. These rules comprise of term matching, regular expressions and part-of-speech tags, along with pre-built Boolean operators expressing constraints, such as the distance and occurrence of concepts in relation to other words.

Definitions	Type	Body	Weight
1	PREDICATE_RULE	(SENT, (ORDDIST_5, (OR, "a_def(TonakkeywordPositive)"), "a_def(senPositiveWor	
2	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(TYPEFIRESERVICE)", (OR, "a_def(TonakkeywordPositive)"))	1
3	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(TYPEFIRESERVICE)", "a_def(strMediaPhrases)", "a_def(strP	1
4	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(TYPEFIRESERVICE)", "a_def(strPositiveSubject)"))	1
5	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(TYPEFIRESERVICE) _def(strMediaPhrases)", (OR, "a_def(1
6	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(strPositiveObject)", "a_def(TYPEFIRESERVICE)"))	1
7	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(strPositiveObject)", "a_def(strMediaPhrases)", "a_def(TYPR	1
8	PREDICATE_RULE	(SENT, (ORDDIST_5, "a_def(strPositiveObject)", "a_def(strMediaPhrases)", "a_def(TYPR	1

Fig. 1.5 The Athena sentiment analysis model (partial screenshot)

Figure 1.5 shows some of the custom-built predicate rules defined for the purposes of Athena, which allow for the definition of semantic relationships between

concepts (i.e. entities). The left column displays some of the concepts which are usually identified in crisis management situations, such as citizens, first responders and law enforcement agencies. Each entity can have it's own set of rules, allowing for *contextual* identification of sentiment; the fourth rule, for example, looks for instances of the fire service concept combined with words expressing positive sentiment, in the same sentence, having a maximum distance of five words between the two.

1.5.2 An example

Figure 1.6 shows a partial example of the sentiment identified in the Colorado wildfires Twitter dataset. The overall sentiment of each document (tweet) is expressed in the first column. When sentiment has been identified for specific concepts in each document, it is expressed in the third column. The fourth row contains the actual body of each tweet, which has been purposefully truncated for ethical purposes.

Inspecting the results of the first row, for example, shows how an overall positive sentiment was identified for that particular tweet, but also how the positive sentiment was expressed towards the Colorado fire service, which evidently handled the crisis successfully. In fact, out of the 1200 tweet corpus, a majority of the sentiment expressed towards governmental entities such as the fire service and the military has been positive.

sentiment	product-sentiment	feature-sentiment	tweet-body
Positive	TYPE--Positive	TYPE--FIRESERVICE--Positive	colorado fire fighte...
Unclassified			Nexus 7
Negative			Hundreds of homes de...
Positive			RT @JoVIClo: God Ble...
Negative	TYPE--Negative	TYPE--CITIZENS--Negative	My prayers go out to...
Positive	TYPE--Positive	TYPE--MILITARY--Positive	So proud of the Colo...
Positive	TYPE--Positive	TYPE--FIRESERVICE--Positive	Colorado wildfire ho...
Neutral			RT @themicininja: H...
Negative	TYPE--Negative	TYPE--FIRESERVICE--Negative	Federal firefighters...
Negative	TYPE--Negative	TYPE--CITIZENS--Negative	RT @kmitcheIIDP: Col...
Positive			shout out to colorad...
Negative	TYPE--Negative	TYPE--GOVERNMENT--Negative	#newbedon 6/26/2012 ...
Negative	TYPE--Negative	TYPE--CITIZENS--Negative	RT @kariontour: Pray...

Fig. 1.6 An example of identified sentiment in the Colorado wildfires Twitter dataset

1.6 Aggregation to reduce information overload

In a crisis situation, decision makers need a clear picture of the events occurring. It is no use being overloaded with information from hundreds or even thousands of

sources, which may well be the case if social media and citizen reporting are being used to obtain information. Thus the Athena system has a process to aggregate sources when they contain information about the same event, greatly reducing the number of information points presented to the decision maker. Furthermore, this aggregation can give an indication of the size, seriousness and credibility of the event simply by the number of sources involved (although the number of corroborating sources should not, of course, be relied upon as the only measure of these factors). In Athena this aggregation is carried out by a clustering technique called Formal Concept Analysis (FCA) [10].

A formal description of formal concepts begins with a set of objects G and a set of attributes M . A binary relation $I \subseteq G \times M$ is called the *formal context*. If $i \in G$ and $j \in M$ then ij says that object i has attribute j . For a set of objects $A \subseteq G$, a derivation operator $'$ is defined to obtain the set of attributes common to the objects in A as follows:

$$A' := \{ j \in M \mid \forall i \in A : ij \}.$$

Similarly, for a set of attributes $B \subseteq M$, the operator is defined to obtain the set of objects common to the attributes in B as follows:

$$B' := \{ i \in G \mid \forall j \in B : ij \}.$$

(A, B) is a formal concept iff $A' = B$ and $B' = A$. Thus A and B have the following properties: (i) Every object in A has every attribute in B , (ii) For every object in G that is not in A , there is an attribute in B that that object does not have, and (iii) For every attribute in M that is not in B there is an object in A that does not have that attribute.

In Athena, the information sources are the objects and the structured data extracted from them are their attributes. Thus an object might have attributes such as a location, a crisis category, a sentiment, a date and time and so on. If the aggregated information is to be presented to the decision maker via a map of the crisis area, we can define a ‘crisis concept’ as being a formal concept that contains at least one location and a least one crisis category. Thus FCA has been implemented in Athena to compute crisis concepts from the structured data obtained from the social media and citizen reporter information sources.

The Colorado wildfire Twitter data provides an example of this and Figure 1.7 shows part of a formal concept tree generated from the structured data obtained from the Tweets. Each node is a formal concept and the node on the left represent the set of 642 Tweets that have a location and a crisis category in their text. Each of the numbered nodes on the right is a ‘crisis concept’ containing a location and a crisis category. The label above each node contains the attributes of the concept and the label below gives the number of Tweets sharing those attributes. If a node is filled in, it means there are further specialised ‘sub-concepts’ that can be explored - containing fewer Tweets but a greater number of shared attributes.

Figure 1.8 shows one of the filled in nodes expanded to reveal its sub-concepts. Each sub-concept inherits the attributes of the ‘parent’ concept but has additional attributes potentially containing more valuable information about the crisis. In the

example, there is a group of Tweets that mention a shooting attack and two groups mentioning a fire or wildfire. There are also groups of Tweets containing negative sentiment.

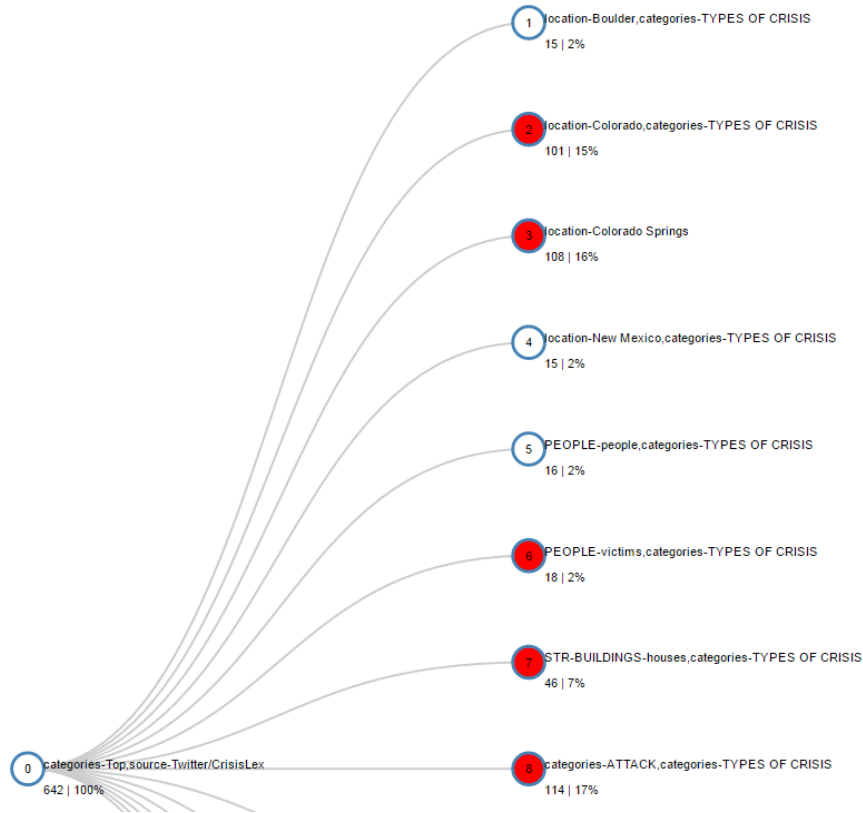


Fig. 1.7 Crisis concepts computed from the Colorado tweets

The analyst/decision maker also has the ability to trace back Tweets of interest to the original source. The crisis concepts, instead of displaying the number of Tweets can display the source URLs. The analyst can select a URL to link back to the original text.

Thus, the application of FCA to aggregate source information into ‘crisis concepts’ facilitates the decision maker by reducing information overload and focusing on crisis information along with its location. In Athena the usability and interpretation is further improved by displaying crisis concepts in a map-based interface with additional filtering and search facilities.




Fig. 1.8 Sub-concepts showing additional crisis information

1.7 Filtering and searching

This section discusses the considerations and approaches to providing command and control operators with the necessary data they need in order to carry out their roles in a crisis situation when using the Athena Crisis Command and Control Intelligence Dashboard (CCCID).

There are various design considerations that should be made when providing filtering and searching capability and what could be arbitrarily applied depending on the role of the individual user or need. Two distinct phases of the crisis may also be recognised, which are during the time of the crisis, where information is more important in a relative context to the current situation, and the post-crisis phase where in depth analysis and visualisation has its importance.

1.7.1 Design considerations

With a system such as the Athena CCCID being widely deployed, at any level, there is likely to be a large amount of data generated at any given time. This is simply due to the link between scale and distribution of crises, i.e. small crises occur often whilst large crises occur infrequently. As a result, the data flows in Athena could be reasonably steady. The Athena consortium has found the following considerations to be of importance, particularly in the scope of crisis management.

Geographical location

Clearly, it is of great importance that the geographical location of any report coming in from the general public, is available and as accurate as possible. Much of this also depends on the scale of the Athena deployment. Where Athena is deployed, for instance, on a national scale, it would be feasible to have various groups or teams of

operators monitoring and responding to particular geographical regions, just as call handlers and first responders would in any emergency response. With appropriate arbitrary filtering on geographical boundaries in place, a national deployment would result in similar behaviour to a regional or even international deployment.

In such a scenario, it would be important for operators to be effectively registered or operational for only a subset of the reports and response in Athena - whilst there may be more strategic overviews at a higher level. Consideration should also be made in that there may be local and hyper-local responders using the mobile app or even the Athena CCCID (voluntary organisations, neighbourhood watch, private sector security) who may require a geographically restricted view of activity.

Validity, credibility and priority

It is occasionally the case that misinformation, or even disinformation, may be propagated throughout a system such as Athena for various reasons (hoax, malicious, political) which is why the validity and credibility of incoming reports should be assessed. This may be a manual process though could be made automated in various ways with varying results.

In Athena, the focus with regards to validity and credibility is primarily on the source of the information, or user tier, which ranges from a full access operator level user, right down to an unregistered member of the general public. Both, validity and credibility assessment or recommendations can be achieved based on whether other information has already been validated or rejected by the CCCID operators. The benefit of this comparative approach is the reduced risk of invading individual privacy.

Priority is slightly different in that it enables the system or operators to carry out validity and credibility assessment in the most efficient manner possible. The use of a triage approach to priority would focus the manual workload of the operators only on the most important tasks first.

Due to the risk of misinformation and disinformation enter the system, avoiding such abuse could be achieved through legislation supporting for individual registration for use or the automatic sharing of identity information by telecommunications companies when interacting with the system.

Information source

Identifying the source of incoming information could be employed to aid the operators and automation processes with validation, credibility and priority. An example of this may be an individual interacting directly with the Athena mobile app which is an ambiguous step, whereas information coming in from social media is more likely to be misinterpreted. This classification should be available to the operator.

Free-text searchability

Any incoming bodies unstructured text should be indexed appropriately to enable free-text search of the system. This should also work in combination with all other filtering options, giving the operator the power to find information within information. The difference between free-text search and all other types of filtering is that there is often no reason as to why it would be arbitrarily applied (in the same way that an operator whose role is validation, may need to see only reports that have not been validated yet).

Information aggregation

Due to the potentially steady, though sizable streams of information entering a system like Athena from uncontrolled sources, it makes sense and may even be vital in many and all situations that information is properly aggregated. Athena does this on-the-fly using a formal concept analysis algorithm. Though it doesn't stop at the actual aggregation process as operators need to be able to choose between accessing individual reports or aggregated reports to ensure information is not missed or even to validate the aggregation process. The reduction in information flow here can be highly beneficial, for example, there may be fifteen reports representing the same instance with slight variations, showing the operator one report instead of fifteen will greatly improve their efficiency.

1.7.2 During crisis

As a crisis is unfolding, operators are unlikely to need functionality that provides them with full exploratory access to the Athena data set. Data needs to be geographically and temporally relevant to the individual operator and delivered in a way that enables them to carry out necessary responsibilities. Considering the types of filtering and searching capabilities discussed, some are more important during a crisis.

First of all, is location. If an operator doesn't have any information regarding the location of those reporting the crisis, it makes their job extremely difficult or even impossible. Though this is less of a challenge when more information is coming in as aggregation can help with this, i.e. a report regarding a vehicle crash in location A, could potentially be aggregated with other reports at a similar time regarding a car fire.

Such aggregation, or even association between multiple incoming reports with various information is difficult without considering the time information came in. During a crisis, it is more important for operators to see time in a relative way, such as "4 minutes ago". Using time in this way and allowing operators to find data based on the last number of seconds, minutes or hours can increase the speed in which they filter through information.

The three points of validity, credibility and priority are all quite similar, but serve their individual purposes and complement one another. These classifiers can be used by operators to change their perspective of the incoming information, i.e. see the current certain and uncertain state of the situation.

Finally, it is important due to the potential for information overload, the provide operators with the least data as possible. As Athena employs aggregation processes to identify, combine and corroborate incoming information, it enables operators to see and react to a more generalised picture of the unfolding crisis.

1.7.3 Post crisis

The needs of the accessing data post-crisis are distinct from at the time of the crisis. At this time the view is both highly generalised whilst requiring greater detail and interrogation capabilities. It is at this stage where data visualisation is likely to provide most value. That said, location and time are still key players.

The element that is more vital with regards to searchability at this stage is with crisis media. Having the ability to quickly and easily visualise and traverse all media captured during the crisis is important for gathering a picture of the entire crisis and its response. This can easily be timelined, as could the entire incident or particular geographical regions.

Timeline visualisations can provide an insight into both when media was captured, likely revealing when key events occurred. They can also provide access to the pace at which new information came into the system at key stages during the crisis (i.e. in an earthquake situation these key events could be initial incident and aftershocks as well as building collapses) to provide an overview of the crisis.

Auditing all data, updates and activities in the Athena system, whether internally or externally invoked is extremely important for analysis the crisis as well as the response to the crisis. Every individual piece of data that enters the system should be recorded in a write-only audit. Every modification to the data or its state should also be record. The end goal here is ensuring that no piece of data goes missing, i.e. removing data from the CCCID may be possible, but the audit will not miss anything. Visualising and analysis such data can provide valuable insights and lessons into the crisis but also the use of the mobile app and CCCID.

1.8 Results and evaluation

Overall we can report some success in the processes we have developed when applied to the social media data from a real disaster. Our results also point to a number of promising lines for future development.

The main objective for the automated aggregation of information sources was to reduce information overload. In the example using the Colorado Tweets, the original

1200 Tweets were firstly reduced to 642 by only considering those that contained a location and a reference to a crisis category. The FCA aggregation reduced this to 76 groups of Tweets, each group containing Tweets with the same location and crisis category. This represents a significant reduction in information points (88%) and, in this case, a final number of information points that would be manageable by an end user. However, if the starting point was a far greater number of information sources, additional steps may have to be taken to counter information overload, such as setting a minimum number of information sources per crisis concept. Better resolution of location names would also increase the level of aggregation and improve the quality of the crisis concepts. In the Colorado Tweets, there were nine 'versions' of Colorado Springs: Help Colorado Springs, Fame Colorado Springs, Colorado Springs, COLORADO SPRINGS, Davis Colorado Springs, Northwestern Colorado Springs, Technician Colorado Springs, The Colorado Springs and Hey Colorado Springs. Further work is required to improve the location recognition and extraction software to resolve identical locations and exclude erroneous locations.

The Colorado Tweets example also illustrated the ability of a crisis taxonomy in providing a useful 'drill-down' feature for the analyst. Figure 1.8 shows crisis concepts with increasingly specialised crisis categories, from TYPES OF CRISIS to NATURAL DISASTER to ND-WILDFIRE. The higher levels of aggregation are achieved at the more general levels in the taxonomy, but the analyst is able to then 'drill-down' into ever more specialised sub-concepts that have fewer information sources but potentially more specific information about a specific crisis event. The most specialised concept may contain only a few information sources and at this point the analyst is able to trace back to the original sources to examine their text.

The sentiment analysis component identified sentiment correctly in more than 87% of the tweet corpus, while the remaining 13% showed how the analysis would benefit from further sentiment specialisation. For example, in some instances, the sentiment analysis component identified negative sentiments expressed towards citizens. Upon manual inspection, these tweets were actually expressing emotions such as sadness, anger and frustration towards the crisis event – this shows how further insight could be gained by extending the sentiment model to incorporate specialised sentiments such as anger and frustration, in the analysis, rather than only using the positive/negative scale.

1.9 Conclusion

Crowdsourcing and other online participatory practices are becoming increasingly important to emergency personnel. The benefits of harnessing social media data and the faster, localised information from technology enabling mass participation are significant. However, the risks and challenges of using such large pools of dynamic, unregulated material cannot be ignored. Here, we have reviewed the benefits and potential dangers of exploiting this information. We have discussed various tools such as sentiment analysis, credibility assessment and priority assessment which

aim to enhance the usefulness and reliability of the data and contribute to emergency assessment and response.

In this chapter we have given an overview of the processes and systems used in Athena as a means obtaining, analysing, filtering and presenting social media and crowd-sourced information. The result of this work is a set of powerful new capabilities which can be used in multiple ways to support the overall goals of the Athena project.

Acknowledgements This work was funded through the FP7 European project ATHENA (FP7-SEC-2012.6.1-30).

References

1. de Albuquerque, J.P., Herfort, B., Brenning, A., Zipf, A.: A geographic approach for combining social media and authoritative data towards identifying useful information for disaster management. *International Journal of Geographical Information Science* **29**(4), 667–689 (2015)
2. Alexander, D.E.: Social media in disaster risk reduction and crisis management. *Science and Engineering Ethics* **20**(3), 717–733 (2014)
3. Andrews, S., Orphanides, C.: Knowledge discovery through creating formal contexts. *International Journal of Space-Based and Situated Computing* **2**(2), 123–138 (2012)
4. Andrews, S., Orphanides, C.: Discovering knowledge in data using formal concept analysis. *International Journal of Distributed Systems and Technologies* **4**(2), 31–50 (2013)
5. de Assis, L.F.G., Herfort, B., Steiger, E., Horita, F.E.A., de Albuquerque, J.P.: Geographical prioritization of social network messages in near real-time using sensor data streams: an application to floods. In: XVI Brazilian Symposium on GeoInformatics, Campos do Jordão, São Paulo, Brazil, November 29 - December 2, 2015., pp. 26–37 (2015)
6. Castillo, C., Mendoza, M., Poblete, B.: Information credibility on Twitter. In: Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad, India, March 28 - April 1, 2011, pp. 675–684 (2011)
7. Cresci, S., Tesconi, M., Cimino, A., Dell’Oretta, F.: A linguistically-driven approach to cross-event damage assessment of natural disasters from social media messages. In: Proceedings of the 24th international conference companion on World Wide Web. ACM (2015)
8. Domdouzis, K., Andrews, S., Gibson, H., Akhgar, B., Hirsch, L.: Service-oriented design of a command and control intelligence dashboard for crisis management. In: Proceedings of the 7th IEEE/ACM International Conference on Utility and Cloud Computing, UCC 2014, London, United Kingdom, December 8–11, 2014, pp. 702–707 (2014)
9. Dubey, R., Luo, Z., Xu, M., Wamba, S.F.: Developing an integration framework for crowd-sourcing and internet of things with applications for disaster response. In: 2015 IEEE International Conference on Data Science and Data Intensive Systems, pp. 520–524 (2015)
10. Ganter, B., Wille, R.: Formal concept analysis - mathematical foundations. Springer (1999)
11. Gao, H., Barbier, G., Goolsby, R.: Harnessing the crowdsourcing power of social media for disaster relief. *IEEE Intelligent Systems* **26**(3), 10–14 (2011)
12. Gupta, A., Kumaraguru, P., Castillo, C., Meier, P.: TweetCred: A real-time web-based system for assessing credibility of content on Twitter. *CoRR* **abs/1405.5490** (2014)
13. Imran, M., Castillo, C., Diaz, F., Vieweg, S.: Processing social media messages in mass emergency: A survey. *ACM Comput. Surv.* **47**(4), 67 (2015)
14. Narvaez, R.W.: Crowdsourcing for disaster preparedness: Realities and opportunities. Master’s thesis, Graduate Institute of International and Development Studies, Geneva, Geneva (2012)

15. O'Donovan, J., Kang, B., Meyer, G., Höllerer, T., Adali, S.: Credibility in context: An analysis of feature distributions in Twitter. In: 2012 International Conference on Privacy, Security, Risk and Trust, PASSAT 2012, and 2012 International Conference on Social Computing, SocialCom 2012, Amsterdam, Netherlands, September 3-5, 2012, pp. 293–301 (2012)
16. Oh, O., Agrawal, M., Rao, H.R.: Community intelligence and social media services: A rumor theoretic analysis of tweets during social crises. *MIS Quarterly* **37**(2), 407–426 (2013)
17. Olteanu, A., Vieweg, S., Castillo, C.: What to expect when the unexpected happens: Social media communications across crises. In: Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW 2015, Vancouver, BC, Canada, March 14 - 18, 2015, pp. 994–1009 (2015)
18. Sievers, J.A.: Embracing crowdsourcing: A strategy for state and local governments approaching “whole community” emergency planning. *State and Local Government Review* **47**(1), 57–67 (2015)
19. Verma, S., Vieweg, S., Corvey, W.J., Palen, L., Martin, J.H., Palmer, M., Schram, A., Anderson, K.M.: Natural language processing to the rescue? extracting “situational awareness” tweets during mass emergency. In: Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Catalonia, Spain, July 17-21, 2011 (2011)
20. Witten, I.H., Frank, E., Hall, M.A.: *Data Mining: Practical Machine Learning Tools and Techniques*, 3rd edn. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (2011)
21. Yin, J., Lampert, A., Cameron, M.A., Robinson, B., Power, R.: Using social media to enhance emergency situation awareness. *IEEE Intelligent Systems* **27**(6), 52–59 (2012)