Consiglio Nazionale delle Ricerche

A platform for gathering eyewitness reports from social media users in the aftermath of emergencies

S. Bellomo, S. Cresci, F. Del Vigna, M.La Polla, M. Tesconi

IIT TR-10/2015

Technical Report

luglio 2015

Istituto di Informatica e Telematica

A platform for gathering eyewitness reports from social media users in the aftermath of emergencies

Salvatore Bellomo, Stefano Cresci, Fabio Del Vigna, Mariantonietta N. La Polla, Maurizio Tesconi

Technical Report

Abstract

Social sensing is based on the idea that communities or groups of people can provide sets of information similar to that which is obtainable from sensor networks. Emergency management and situation awareness are candidate fields of application for social sensing.

Nowadays, two different approaches are present in literature: opportunistic and participatory crowdsensing, the former of which intends to detect events and/or gain situation awareness by gathering data 'on-the-fly' while the latter 'hires' volunteers in order to retrieve valuable information.

This work aims to create, implement and deploy a platform based on a decision support system for gathering eyewitness reports in the aftermath of an emergency, focusing in particular on earthquakes. These reports can be useful to improve the situation awareness. While doing so, we would like to find out if an approach combining opportunistic and participatory sensing methods is effective. Our system, in fact, detects eyewitnesses exploiting an opportunistic approach and then aims to transform these potential eyewitnesses into volunteers willing to share information.

The platform retrieves earthquake notifications from an official channel and, immediately after, gathers the messages shared on Twitter for a fixed timeslot. In doing so, we collect messages posted by potential eyewitnesses. Data mining and natural language processing techniques are applied in order to select meaningful and comprehensive datasets of tweets. We then concentrate on the filtered dataset in order to try to engage with their authors and obtain, in real time, information and enhance situation awareness.

Information retrieved by our system can be extremely useful to all the government agencies interested in mitigating the impact of earthquakes, as well as news agencies looking for new information to publish.

Results collected by our platform are promising and, despite being in its preliminary stages, a combined approach to the search for earthquake eyewitnesses seems possible.

Table of Contents

Introduction

Social Networks can be described as online platforms upon which people are able to socialise and interact with each other. The vast majority of social networking sites allow users to communicate with those who have either similar interests, hobbies and backgrounds, or who share a real-life connection. Users of social networks can create their own public profiles and post statuses and updates on the happenings of their everyday lives.

Today, websites such as Facebook and Twitter play a crucial role in modern communication as they facilitate the process of socialisation and allow users their own personal pedestals for sharing emotions, feelings and opinions. Thanks to the increase in recent years of social media usage, collecting information from social networking websites has become an important subject in social sensing literature.

Social sensing is based on the idea that communities or groups of people can provide sets of information similar to those obtainable from sensor networks. Emergency management and situation awareness are candidate fields of application for social sensing. Users can thus be considered Social Sensors as they represent a rich source of information on situations, facts and social contexts, as asserted by the Social Sensing (or the Human as a Sensor) paradigm (Zhou, Qian, & Ma, 2012).

A survey conducted in USA[1] showed that the benefit of the informal communication through Twitter lies in the early diffusion of emergency information and the potential to organise mutual help within neighbourhoods. People in a disaster zone typically post real-time information, including unofficial messages and rumours, and repost and *retweet* official messages. Two kinds of approaches have been developed to gather this information: opportunistic and participatory sensing.

Participatory sensing is a relatively new paradigm that allows people to voluntarily sense their environment using readily available sensor devices, such as smart phones, and share this information using existing cellular and Internet communication infrastructure. It harnesses the power of ordinary citizens to collect sensor data for applications spanning environmental monitoring, intelligent transportation and public health, which are often not cost-viable using dedicated sensing infrastructure. A well-known example of participatory system is Wikipedia. In such platform "digital volunteers" write collaboratively in order to share their knowledge to the Internet community.

Opportunistic sensing is a paradigm that leverages interactions among users by listening to the media in order to retrieve valuable information. The idea of social media websites is that data is shared and virtually available to the community. It is therefore possible to analyse posts and comments in order to retrieve information on emergency events. Opportunistic approaches have been used in literature in order to automatically detect occurring disasters. This is something that many projects focus on and, thanks to the effects that mass emergencies have on users of social media, the results are promising.

The micro blogging platform Twitter (more so than Facebook) provides an unprecedentedly open and accessible space for such activities. It builds on a much simpler networking structure where updates posted by users are either public or private, rather than visible and shareable only to selected circles of friends within one's social network. In fact, public Twitter messages are visible even to unregistered visitors.

Given that social media is used by a large majority of people, disaster response agencies have started to utilise it as a source of information. Twitter is especially suitable for this kind of analysis: over 500 million tweets[2] (messages or status updates limited to 140 characters) are posted online every day that usually contain keywords (hashtag) of trending topics.

The participatory approach, on the one hand, gathers specific information using digital volunteers 'hired' beforehand - a time-consuming task that varies according to the location of the disaster. On the other hand, the opportunistic approach focuses more on retrieving 'on-the-fly' information, meaning that the chance of

---

[1] www.fairfaxcounty. gov/emergency/flooding-090811-metrics.pdf
[2] http://www.internetlivestats.com/twitter-statistics/

finding valuable, spontaneous information is higher. However, this information is more likely to be fragmented and unstructured (due to limited characters on Twitter posts). It is therefore important that noise is detected and deleted in the opportunistic approach.

Our goal is to accelerate the damage assessment process by having users interact and participate in giving useful information. We aim to leverage the opportunistic approach in order to detect potential eyewitness and combine it with participatory sensing so that we can 'hire' volunteers to share their valuable information. The platform will first automatically detect potential eyewitnesses and then contact users in order to encourage them to help others of their own accord and give them the chance to reply with their own messages.

Exploitation of the information shared by people directly involved in the emergency allows us to automatically increase situational awareness and to obtain estimations of the consequences on communities and infrastructures without the typical delays of in situ assessments.

State of the Art

Approaches that exploit information available on social media for emergency management have been carried out and experimented on before. Collecting information from social media users is an interesting task that since the last decade, thanks to the increase in social media usage, has become a much more widely discussed topic in literature. For example, the U.S. Geological Survey (USGS) is investigating how the social networking site Twitter, a popular service for sending and receiving short, public text messages, can improve USGS earthquake response products and the delivery of hazard information.

Depending on the user's awareness and their involvement in the architecture of the system we are confronted by either an *opportunistic* or a *participatory* approach (Lane, Eisenman, Musolesi, Miluzzo, & Campbell, 2008).

When users consciously opt to meet an application request out of their own will, this is called participatory sensing. The public is asked to gather, analyse and share data and information with the integrated sensor capabilities of the system (usually mobile devices: camera, GPS, etc.) or, especially in the case of social media approaches, reports. A few methods of crowdsourcing systems have been experimented on before, focusing on different fields and using a variety of tools. "*Digital volunteers*", so-called for their own will of share, consciously, their experience, are a crucial element in these approaches.

An example of participatory approach is the platform created by USGS which was made available to the public and which allowed users from all over the world to share their earthquake experiences using the "Did You Feel It?" (DYFI) system[3]. By taking advantage of the vast number of Internet users, USGS is able to get a more complete description of what people experienced, as well as the effects of the earthquake and the extent of damage.

Other examples of participatory approach are: *Ushahidi platform*, a crisis-mapping platform which supports professional organisations with options for requesting citizens or digital assistants to gather, structure or share information (Heinzelman & Waters, 2010); *Mobile4D*, an application which emergency services use to request affected citizens to submit reports about their local situation (Frommberger & Schmid, 2013); *DIA-DEM* system represents another way of gathering and validating civil information where a pre-selected expert group of citizens are requested services to use a mobile application for identifying strange smells during chemical disasters, the collected responses are shared between experts and visualised on a map, so that emergency services can derive possible locations of an affected chemical factory (Winterboer, Martens, Pavlin, Groen, & Evers, 2011).

In contrast with participatory approach, in the opportunistic sensing, information are gathered by simply listening to the media (e.g. social media) and extracted from spontaneous messages or status shared by users. An example of this approach is described in EARS – Earthquake Alert Report System (Avvenuti, Cresci, Marchetti, Meletti, & Tesconi, 2014), a system that detects and evaluates how consequences of earthquakes are assessed. The system has been tested on Italian territory and employs data mining and natural language processing techniques on social media data in order to enhance situation awareness following seismic activity. Other recent efforts have focused on the assessment of the consequences of natural disasters on communities and infrastructures. In (Cresci, Tesconi, Cimino, & Dell'Orletta, 2015), authors employed machine learning techniques on Twitter data to automatically detect mentions of damage among tweets. Tweets conveying actionable information about damage have been later exploited for crisis mapping tasks (Cresci, Cimino, Dell'Orletta, & Tesconi, 2015).

Several other initiatives, exploiting an opportunistic approach, have been developed: *CrisisTracker*, for example, is a platform for exploring Twitter within different types of disaster (like tsunamis, nuclear disasters, terroristic attacks). It retrieves tweets by a keyword and location with the aim of creating 'social awareness', tweets then can be visualised on a map or a timeline (Rogstadius, Vukovic, Teixeira, Kostakos,

[3] http://earthquake.usgs.gov/earthquakes/dyfi/

Karapanos, & Laredo, 2013). An early warning system (EWS) is another example for the real-time detection of earthquakes and tornadoes in Japan based on Bayesian statistics is described in (Sakaki, Okazaki, & Matsuo, 2010).

On the whole, opportunistic crowd sensing methods proved to be valuable in the task of detection, whereas participatory techniques were important and useful in retrieving precise, additional information. As far as we know, previous works have only focused on one approach and never combined both of them. After having analysed each method, we have decided to direct our project towards improving the ways that information is collected by reaping the benefits from and integrating the principles of both participatory and opportunistic methods in order to create one reliable system.

Scenario

The National Earthquake Information Centre (NEIC) locates around 50 earthquakes each day, or 20,000 a year[4]. Despite this, the estimation is much higher than these figures as many earthquakes go undetected, especially those that hit remote areas. This statistic suggests that earthquakes are a perfect example of an emergency situation that is easy to monitor, since those who feel a tremor are likely to flock onto social networking websites to share their experiences.

Twitter is an online social networking service that enables users to post and read short 140-character messages called "tweets". The short size of these statuses is arguably what made Twitter so famous. Since messages can be no longer than 140 characters, people need to be as concise as possible when writing a tweet. It's for this reason that the hashtag system (brief keywords preceded by the symbol, "#") has been introduced to Twitter. This system provides a simple and elegant solution for tagging one's own updates as relevant to specific topics that thereby makes them (and the continuing discussions to which they may belong) discoverable and traceable by others. Hashtags enable public conversations by large groups of Twitter users without each participating user needing to subscribe to (to "follow") the update feeds of all other participants. They are also especially effective at establishing topical communities ad hoc, such as in response to breaking news stories.

When Twitter users post tweets, they have the option of displaying their current location alongside their post – this feature is called *geolocation*. This data could well play an important role in our project as it allows us to potentially identify eyewitnesses in the event of an earthquake (or generally an emergency situation). Unfortunately, the usefulness of this approach is limited; in fact, in our dataset, composed by 3 million tweets, collected from Twitter since December 2014 to April 2015, around 4% of all tweets were "*geotagged*" with explicit geographical information. Given the lack of geotagged posts, we cannot commit ourselves solely to finding earthquake witnesses via this approach (Cresci, Cimino, Dell'Orletta, & Tesconi, 2015).

So in order to find earthquake eyewitnesses on Twitter, we decided to keep the approach with geotagged tweets but separate from the majority of tweets that may be not. Thus we analyse the stream of tweets using two different approaches: **geotagged** tweets and **keywords-related** tweets.

Overall functioning of the platform

The aim of this project is to create a platform able to gather more information from spontaneous reports on social media, using tweets as source, and approach them in order to establish a direct contact with the authors of earthquake-related tweets, immediately after the detection of an earthquake.

Both the approaches, geotagged and keyword-related, can be categorised into four main stages (see Figure 1):

---

[4] http://earthquake.usgs.gov/earthquakes/eqarchives/year/eqstats.php

- ***Search Phase***: Monitoring USGS feed for detecting new earthquakes and retrieving tweets published within a time window (15 minutes in our experiments) of the occurrence of an earthquake –
- ***Filter Phase***: Filtering of collected tweets
- ***Contact Phase***: Contacting Twitter users
- ***Reply Phase***: Collecting and analysing responses

During the **Search Phase** the system monitors the USGS news feed in order to detect new occurrences of earthquakes. Starting from ***Search Phase,*** if we find an earthquake to analyse, we store all the information in a database. Next, we retrieve and store all the tweets within a period of **15** minutes (until we have filled the fixed timeslot) since the earthquake occurred. Then, we analyse them into the ***Filter Phase***, classifying the tweets that are referring to an ongoing event and the ones that are not. The filtered tweets belong to potential eyewitnesses and we want to gather more information by contacting them into the **Contact phase**. Even though there has not been any data retrieving from Twitter stream, thus, no new earthquake has been detected, **Reply Phase** will perform a check and will retrieve eventual replies by the contacted users.

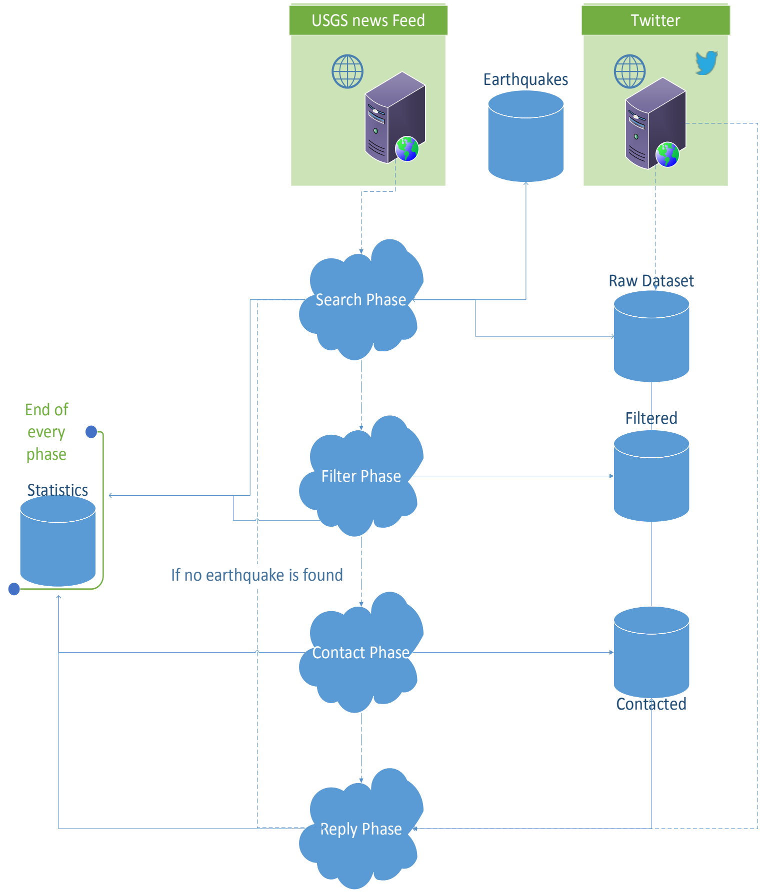

Figure 1 – Overall functioning of the platform

Search Phase

In order to collect earthquakes in near-real time, every minute, a script monitors the USGS news feed. As soon as USGS posts the detection of an earthquake online (immediately for Californian earthquakes and within 10 minutes for the rest of USA), a script will search all the tweets generated within a period of **15** minutes after the earthquake occurred.

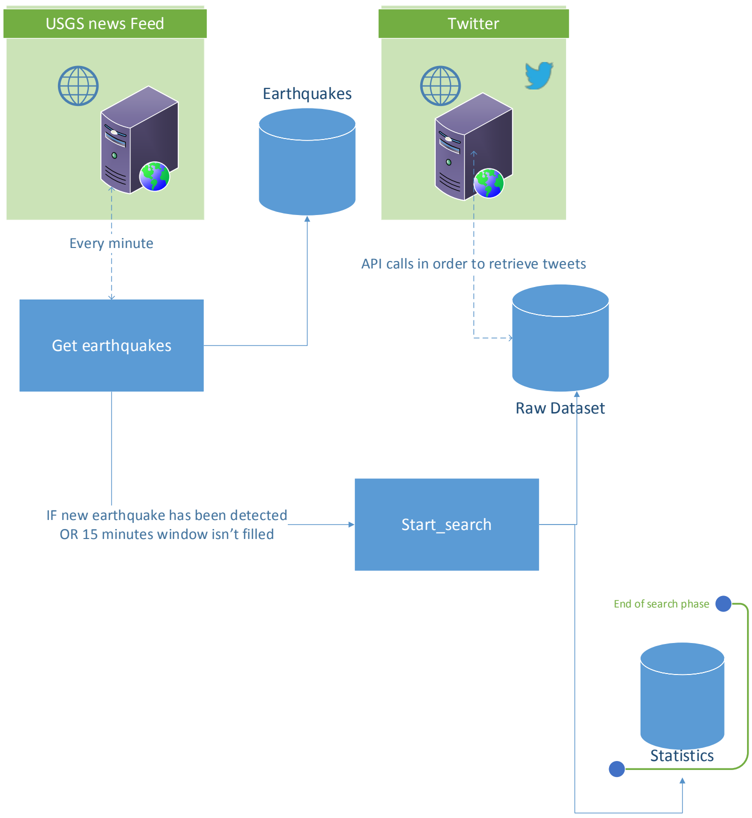

Figure 2 – Search Phase, overview

When an earthquake occurs there are many factors to consider in evaluating its consequences. One of the most important and that, which we are focussing on in particular, is the approximate distance away from an earthquake's epicentre that people are able to feel any tremors.

In order to calculate the *perceptibility radius* that is the radius in which the earthquake is detectable we used a simplified mathematical formula[5] [see Equation below]. The equation correlates the acceleration values of the seismic wave, spread through the planet's material (types of landscape including desert, plain, mountain, cliff, etc.), to a perceptibility threshold (P_{th}) as function of magnitude:

Where *g* is the acceleration of gravity, *x* is the magnitude of the earthquake and D is the so-called *hypocentral distance*.

$$\frac{10^{-1.296+0.556*x-1.582*\ln(D)}}{g} = \mathrm{Pth}$$

$$Pth = 0.005$$

[5] The project is in collaboration with the INGV (National Institute of Geophysics and Volcanology)

The distance between the hypocentre and the epicentre is the depth of the earthquake, information that we retrieve from USGS reports. Understandably the deeper the earthquake, the less likely it is that people will be able to feel any tremors.

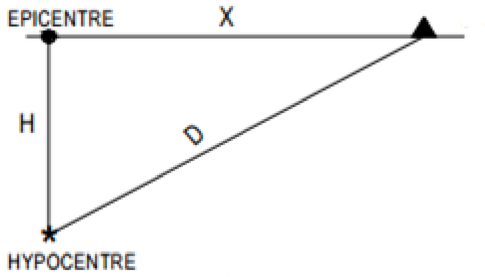

Figure 3 - The arc method

Figure 3 shows the *arc method* or *circle method*, a geometrical representation of the "epicentral" distance *X* (perceptibility radius), the "hypocentral" distance *D* and depth of focus (H) of an earthquake (Kayal, 2011). Pythagoras's theorem can be applied to this figure in order to calculate the value of X, the *perceptibility radius*.

We use the *perceptibility **radius*** (with an extra **15** km added to make up for inaccuracies), together with ***latitude/longitude*** of the earthquake to search earthquake-related, geo-tagged tweets.

As mentioned above, the percentage of geo-tagged tweets is negligible, making it more difficult to find earthquake witnesses on Twitter. For this reason, we also adopt, in this phase a keyword-based approach and retrieve tweets containing the keyword *earthquake.*

Filter Phase

Around 6,000 tweets, on average, are posted on Twitter every second, which adds up to over 350,000 tweets per minute (Zuckerberg, 2015). We monitored the stream of tweets for 6 months and we realised that, on average, we collected 20 tweets per minute. Even when there is not an earthquake, the noise is significant.

"Noise" in this case refers to the tweets that are not related to the specific emergency event despite contain the keywords searched in an attempt to gather eyewitnesses. Avvenuti et al. (Avvenuti, Cresci, Polla, Marchetti, & Tesconi, 2014) have identified two main sources of noise when completing a similar task: firstly keywords that are in fact homographs in that they have different meanings to those we are searching and keywords which refer to past events.

As shown in Figure 4, the filtering is performed by cleaning data in two steps. A pre-filtering phase, similar to that used for geotagged tweets, that applies raw rules to discard tweets that are considered noise. In this step we discard tweets posted by users in a blacklist of 124 Twitter accounts, owned by authors of tweets that periodically publish information about seismic events. It discards tweets contained text patterns that clearly do not refer to an ongoing seismic event or that contain information about earthquakes spread by a news account. It discards retweeted messages and tweets that are responding to other posts.

Tweets that are not discarded in the pre-filtering phase are then analysed with ad hoc classifier to perform a more fine-grained selection. During the offline training phase, the classifier was trained using two distinct sets of messages in order to recognise users that experienced first-hand the seismic event: tweets related and tweets not related to a current seismic event.

Our analysis of the messages reporting earthquakes has highlighted a few interesting features that help to distinguish between tweets related or related to ongoing seismic events.

On the one hand, tweets referring to an earthquake generally are very brief, present less punctuation than normal tweets and often contain slang or offensive words. We can assume that this is because people in the midst of an earthquake are likely to be frightened or apprehensive, and want to convey their fear through social media. On the other hand, tweets that refer to official news of an earthquake or that are referring to past earthquakes are normally longer, well-structured and grammatically sound. Tweets that are not related to a recent earthquake also include a higher number of mentions and URLs than spontaneous earthquake reports.

We then defined the following set of features that takes into account the results of the previous analysis:

- Character count;
- Word count;
- Punctuation count;
- URL count;
- Upper case ratio (capital letter / lower-case letter);
- Magnitude;
- Mentions;
- Exclamation marks.

The classifier was obtained using the decision tree J48, corresponding to the Java implementation of the C4.5 algorithm with a 10-fold cross validation. We gathered **5469** tweets, posted within a period of 15 minutes after the occurrence of 187 earthquakes happened between June and August 2014. We then used these tweets for the training set and manually classified them using an ad-hoc interface.

The training set had 3771 tweets classified as 'NO' and 1698 classified as 'YES'. Then, we trained the classifier in order to build the model.

TP Rate	FP Rate	Precision	Recall	F-Measure	Class
0.726	0.143	0.696	0.726	0.71	1
0.857	0.274	0.874	0.857	0.866	0
0.816	0.234	0.819	0.816	0.817	Total

Training the classifier with this set of features produced correct classifications in more than 80% of the tweets of the training set. The test phase results are reported in the confusion matrix below, where columns represent the instances in the predicted class and rows represent the instances in the actual class.

		Predicted Class	
		Yes	No
Actual class	yes	1232	466
	no	538	3233

The prediction is performed at run-time by invoking the classifier every time a message passes the pre-filtering phase. As Weka generally needs less than one second to predict the class of a new tweet, it is feasible to use the fine-grained classifier filter in our real-time system.

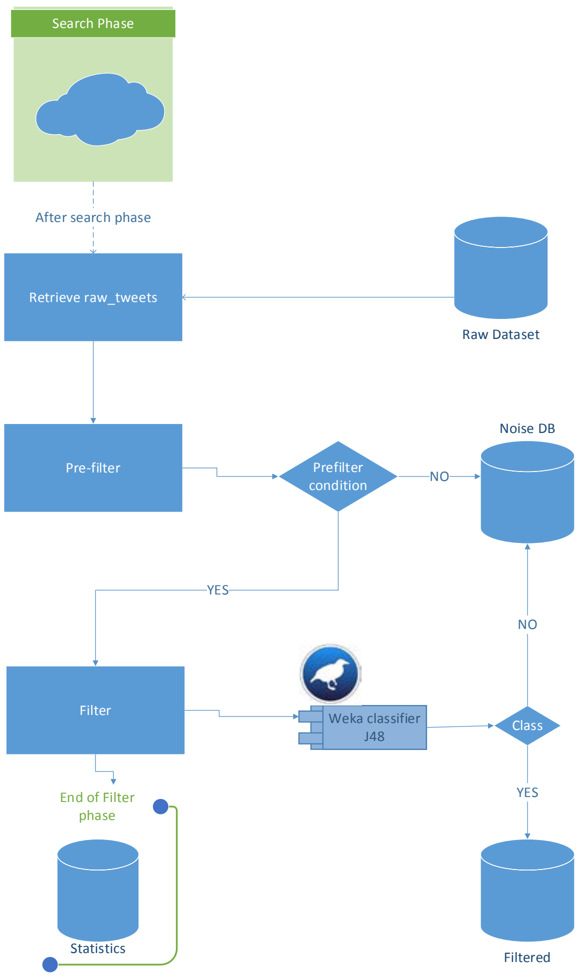

Figure 4 - Filter Phase

Contact Phase

At this stage of the process, we now have a filtered dataset of tweets and their authors. The filtered users are potential eyewitnesses and therefore valuable resources. This phase consists of the most innovative task in literature: approaching users in order to ask them to provide us with confirmation of the event that occurred along with some additional valuable information.

A challenging task is choosing the right question, that we defined *approach tweet*, to ask users. We need to convince and encourage users to give us information as their answers are valuable, but avoid annoying them at the same time. This all needs to be achieved in no more than 140 characters. The approach tweet mentions the user to contact and a text depending on the approach we are using (geotagged or keyword-related).

With the Geotagged tweets approach, we are sure that there is a relationship between the earthquake detected and the tweet retrieved. Thanks to this, approach tweets could carry more information in order to demonstrate that our request is trustworthy. This should put the potential eyewitnesses that we are contacting into a positive frame of mind and thus increase the probability of receiving a response.

Furthermore, our experiment aims to try an approach in which the user has to simply answer the question with a *YES* or a *NO*. In this case, we can analyse the reply and find out if they've been affected or not. The analysis consists of content parsing that checks to see if the response contains words that indicate a positive or a negative answer (we also check for slang words, or words written in a foreign language that can be easily translated).

The information contained in our 140 characters message includes these components:

- A mention to the user's Twitter username;

- Magnitude of the earthquake;
- Time in UTC format;
- The amount of time after which the earthquake occurred. This element changes with every message that we post. This is important in order to try to minimise the risk of being banned for posting the same message multiple times;
- The nearest city where the earthquake took place (when available);
- The state or the country where it occurred;
- A simple question that requires a straightforward answer.

Approach tweets for the keywords-related approach are different. We need to vary the way we write the approach tweets in order to avoid being banned by Twitter, since we cannot post the same tweet more than once. We wanted to take advantage of this by finding out which type of approach tweet elicited most responses from users.

We decided to ask various questions in order to test their sociological effects, i.e. how people react to questions that are worded differently. We chose two main, different ways to start our approach tweet to determine how users respond to different messages. The first way involves declaring at the start of the message that the author of the tweet is in fact a bot. The second way however avoids drawing attention to the fact that the author is a bot.

In the second part of our approach tweet, we state the fact that we detected an earthquake and then ask the user a question.

We can divide our questions into different groups:

- Type 1: “Where” - With this question we try to gather important and useful information. We also pushed the edge of the privacy in order to verify and assure that the earthquake detected by our system matched with the earthquake felt/experienced by the user.
- Type 2: “Are you alright?” - This is a safety check question that approaches users in a friendly manner.
- Type 3: “Are you okay?” - Identical to type 2.
- Type 4: “Is that right?” - Asking users to confirm of the event.
- Type 5: “Have you been affected?” - This is a more direct question asking if users have experienced anything or have been affected by the emergency situation. This sort of direct question encourages more of a direct and informative response as users may feel encouraged to elaborate on the ways they have been affected.

Following a few example of approach tweets.

- Hi @username, this is an auto-response. We have noticed you may have felt an earthquake. Are you alright?
- Hi @username, this is an auto-response. We have noticed you may have felt an earthquake. Are you OK?
- Hi @username, @socialsensing has noticed you may have been involved in an earthquake, could you tell us where are you?
- Hi @username, @socialsensing has noticed you may have felt an earthquake. Could you let us know if you've been affected?
- Hi @username, @socialsensing has noticed you may have felt an earthquake. Is that right?

Where *@username* is the user we would like to contact and *@socialsensing* is the official Social Sensing account from CNR Pisa.

Figure 5 - Example of tweets exchanged in Contact phase

The tweet we post is a reply to their first tweet, *mentioning* the user we want to contact asking them for information.

Reply phase

At the end of every iteration of the script that activate the process of the platform, even if no earthquake has been detected, a script monitors if someone has replied to our approach tweets. Since the Twitter API call retrieves the last 20 mentions for every iteration, thanks to the frequency in which we run this script we assure we don't miss any notification.

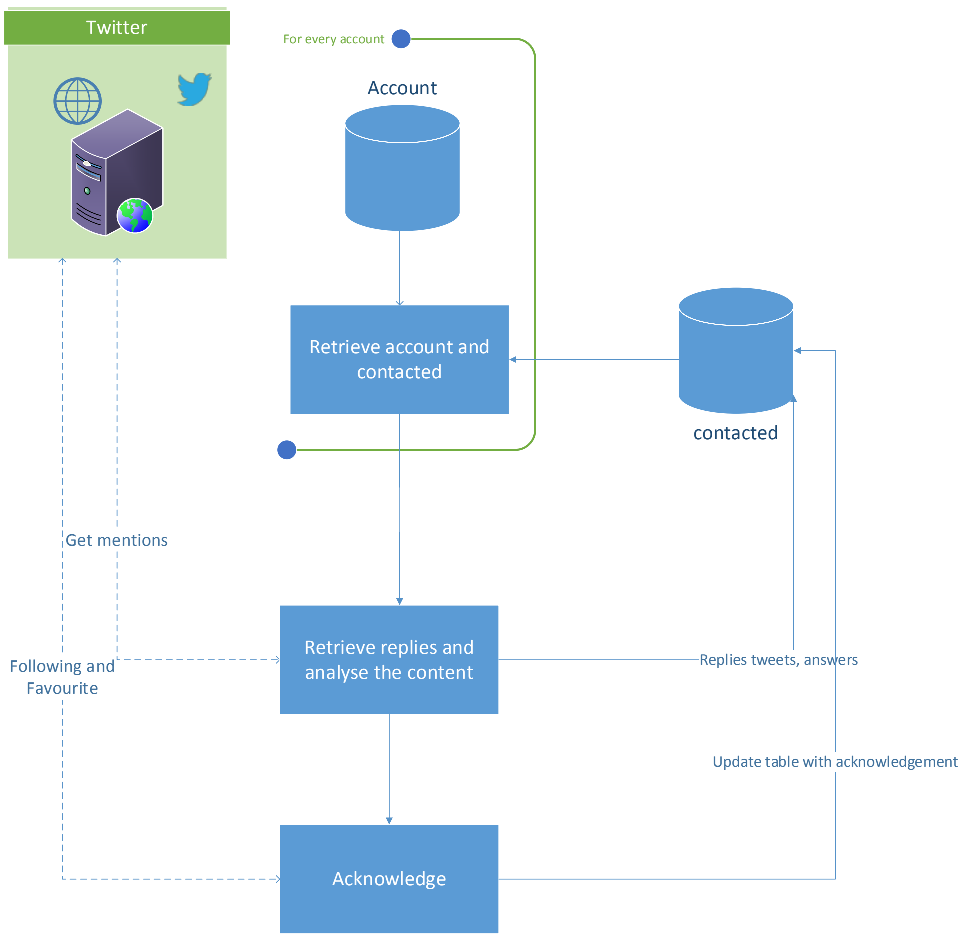

Figure 6 - Reply phase, overview

In order to retrieve all the replies, we retrieve all the account from the pool of accounts (**34** for geotagged, **24** for keyword related) and for every of them we connect to Twitter API stream using the function *GET statuses/mentions_timeline*. It returns the 20 most recent mentions (tweets containing a users's @screen_name) for the authenticating user. This method can only return up to 800 tweets[6].

The timeline returned is the equivalent of the one seen when you view your mentions on Twitter website (example shown in Figure 7).

[6] https://dev.twitter.com/rest/reference/get/statuses/mentions_timeline

Figure 7 - Mentions timeline for a social sensing account

In order to 'thanks' the contacted user for their feedbacks (replies to our approach tweet) in our experiment, the mentioned account will perform two Twitter API call:

- *POST friendships/create* that allows the authenticating users to follow the user specified in the ID parameter.
- *POST favorites/create* that favourites the status specified in the ID parameter as the authenticating user. This API call returns the favourite status when successful.

Experiments and Results

In order to constantly monitor the evolution and the system's response to the occurring emergencies we developed a web application that can be accessed online and that is compatible with all browsers. Using this application, we can check the developments of the contact and reply phases in real time and view reply tweets when they arrive.

The real-time nature of our application, poses some not trivial time constrains. We need to ensure that the timing of our messages is appropriate; if a user is contacted a long time after their original earthquake-related tweet is posted then there's less chance that we will receive a useful response because of loss of interest in the topic. Moreover, the tuning of the system is difficult because we deal with real time replies of users.

The platform was active from **26 February** until **01 April 2015.** In this period, we retrieved earthquakes from the USGS news feed with a magnitude greater or equal to **1.0** for the Geotagged tweets approach and a magnitude greater or equal to **2.5** for the keyword-related tweets approach. We crawled tweets while maintaining the connection between earthquake and tweets and then contacted the ones that were filtered.

We have collected **931** earthquakes with magnitude greater or equal to **2.5** (**4568** with mag >= 1). We contacted more than 10,000 users (7,000 after a geotagged tweet), approximatively **30%** (**25%** for geotagged) of which replied to us

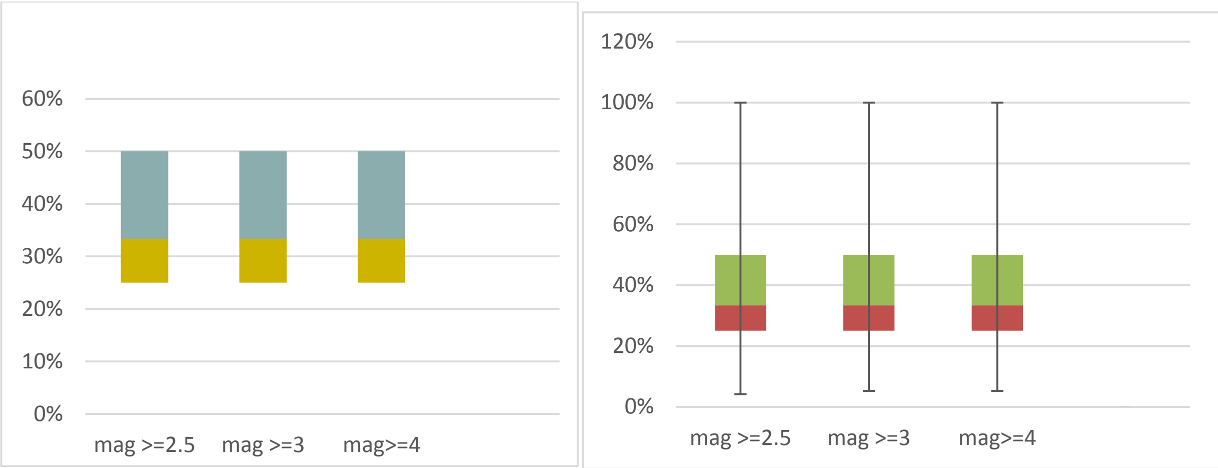

Figure 8 - Distribution of the ratio between contact/replies Geotagged tweets (left) and Keyword-based (right) approach

Figure 9 shows the ratio between replies and contact tweets. The distribution fluctuates from 20 to 50%.

We believe that the proximity to an English-speaking country is more likely to affect the probability of receiving a response from users than the magnitude of the earthquake, in fact, the ratio doesn't seem to be influenced by the magnitude of earthquake. Furthermore, since our platform is based on the Twitter paradigm, we believe that the reply rate strongly depends on the population of city that is in the vicinity of the earthquake.

In the geotagged tweets approach, our main search parameter is the radius within which we crawled tweets through the Twitter API in order to spot potential eyewitnesses. This radius is calculated, as explained above, from an estimated formula with a few extra kilometres added. Since the questions asked require a straight answer, we could test the accuracy and truthfulness of the formula by monitoring and verifying the answers.

In order to conduct geo searches, the search API first attempts to find tweets which have latitude/longitude (lat/long) within the queried geocode, and in case of not being successful, it will attempt to find tweets created by users whose profile location can be reverse geocoded into a lat/long within the queried geocode, meaning that is possible to receive tweets which do not include lat/long information. Because of this we could not show all the retrieved tweets on the map, as described below.

Figure 10 shows, as an example, the map of eyewitness replies for an earthquake, magnitude of 4.4, which occurred in Chile.

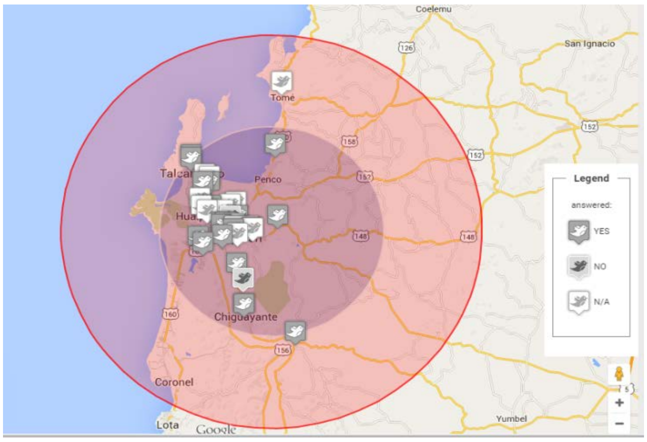

Figure 9 - Chile earthquake: map of the eyewitness replies

This earthquake occurred in the very proximity of a city. We collected 94 answers, 56% of which were positive replies (53), 13% negative (12) and 31% (29) of these responses were in a format not recognised by our parsing tool. In this case, the high number of N/A is because of the location of the earthquake. Most of these replies were written in Spanish and we couldn't analyse the content.

The darker inner circle has a radius calculated using the above mentioned formula, while the lighter outer circle has a few extra km added to the radius to make up for any inaccuracies.

We assumed that every reply proves valuable in gathering information to gain situation awareness. The overview of the replies is quite interesting.

The filtered tweets of the potential eyewitness are short and express shock, surprise or, in the worst case, fear, and often include slang and sometimes grammatical mistakes. Pictures attached to tweets show what is going on and are therefore noteworthy and important in gaining and enhancing situation-awareness.

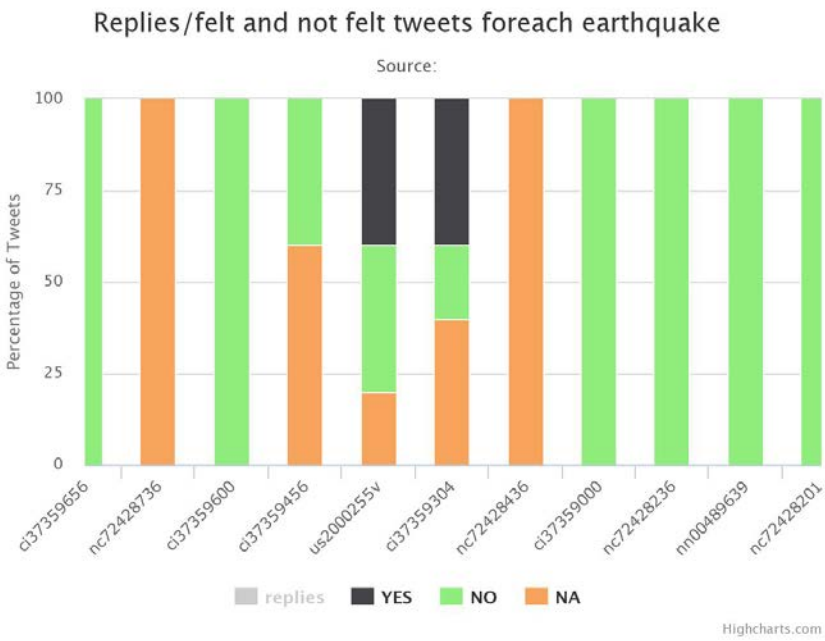

Figure 11 – Replies analysis (geotagged approach)

Geotagged tweets approach gives us the opportunity to retrieve a structured information from the replies, since the question requires a straight answer: YES or NO.

In Figure 11, we can see an example of analysis of replied tweets: Among all the replies we recognised affirmative replies (*YES*), negative replies (*NO*) and Not Available replies (NA). This last one means that we were not able to recognise the answer due to some reasons: (i) the answer is ambiguous.

It is worth noting that the timing of contacting users is really important. Our experiment (see Figure 12) shows us that users are more likely to respond if the approach tweet is posted immediately after their tweet was posted. The two boxplots indicate the distribution for approach tweets that received a reply (**Replied**) and approach tweets that didn't (**No replies**). On the axis X we have the two boxplots and on the axis Y we have the time spent in minutes. This could be due to the fact that the lifespan of a tweet is really short, a user might sign out of Twitter soon after they post their tweet, or they might simply lose interest in the topic.

Since our platform is in the early stages we want to make sure that we contact as many people as possible and as quickly as possible so that users are still discussing the event when they are contacted by us.

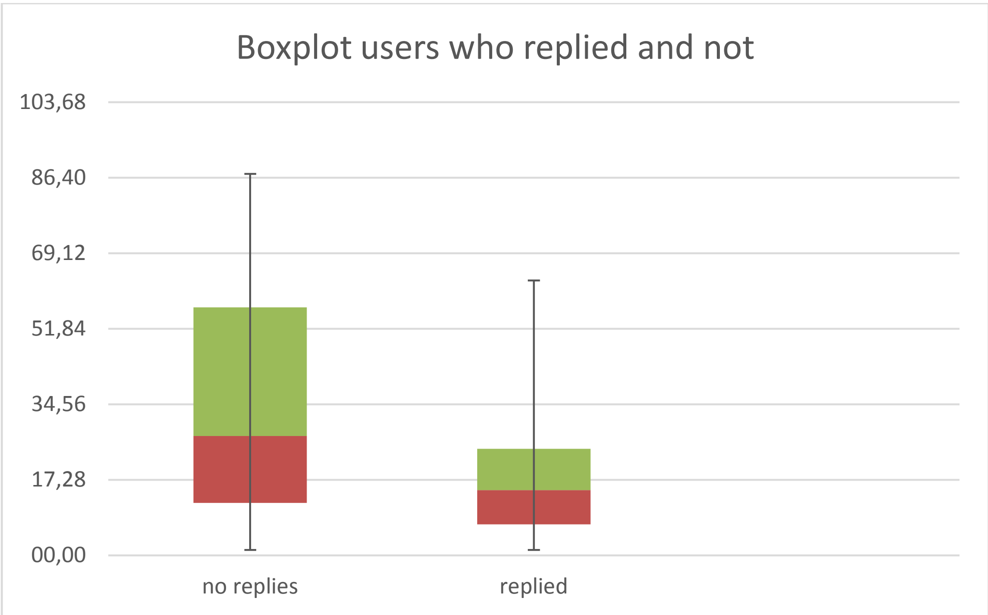

Figure 12 – Boxplot of time distribution for approach tweets that received (left) or not (right) a reply

The approach tweets posted with the first **20 minutes**, but especially in the first 15 minutes, are more likely to receive a reply. Approach tweets posted after 15-20 minutes are less likely to receive an answer.

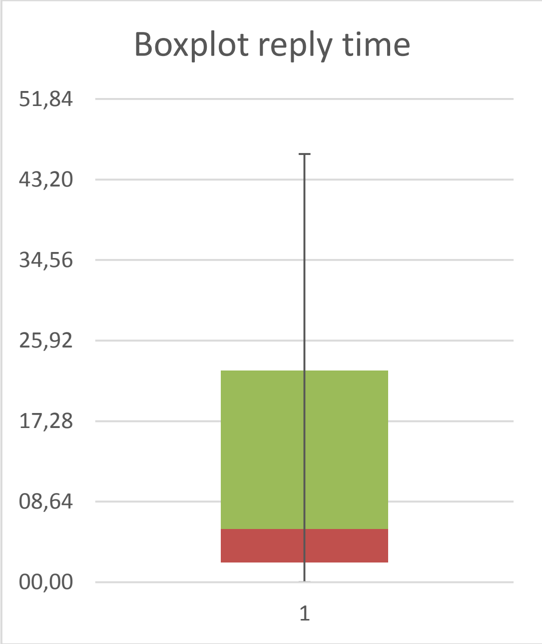

Figure 13 - Boxplot of time spent by users in replying

Figure 13 shows us the distribution of the time spent for users to reply to our approach tweets (in minutes). The majority of users reply within the first 10 minutes – this suggests that the platform could prove valuable to first responders who seek 'new' information on a disaster.

We believe that 15 minutes is a fair trade-off before the news is spread and the lifespan of the topic is used up.

Conclusion and Future Work

The development and creation of this platform has shed light on two fundamental challenges: opportunistic sensing, that is the task extraction and processing of social media data for emergency management, and participatory sensing, that is 'hiring' people in order to gather information. In this work we have discussed techniques for the detection and monitoring of emergencies, we have proposed some possible solutions, and we have discussed techniques in gathering information by contacting eyewitnesses in order to enhance situational awareness.

We have provided extensive experimental results deriving from the employment of the proposed techniques in the field of earthquake emergency management. It is true that all the results collected are preliminary and that further improvements may lead to even better performances and more valuable information. But this is due to the fact that the data collected is real time and we therefore cannot reproduce any of the experiments if we change any minor details. In fact, adding something new or removing any elements would essentially have the same effect as starting the experiments from scratch.

The experiment is in its initial stage and the first question we wanted to answer is if such a system was practical. Overall, the results we collected are promising and seem to favour the adoption of such techniques. We faced a variety of challenges and some areas we focused on definitely deserve more time and investigation, for example carrying out the online emergency monitoring and being able to detect earthquakes promptly. Techniques for extracting knowledge and gathering situation awareness from the textual and multimedia content of messages, as well as disaster intensity management, may be able to contribute to emergency management procedures.

Eliminating or reducing the delay in detection so that we are prepared to act immediately when contacting potential eyewitnesses is another crucial challenge that we are faced with, and one that resembles the issue often discussed in literature that involves finding an efficient system to detect emergency situations. One of the best systems is the one we mentioned above: EARS (Earthquake Alert and Report System) (Avvenuti, Cresci, Marchetti, Meletti, & Tesconi, 2014) that is able to detect an earthquake between 30 seconds and 4 minutes and that has an accuracy of 80%, depending on the magnitude of the earthquake.

The task of reaching the users in replying is an interesting task, having solely 140 characters in order to convince human beings to help a Twitter user who generates automatic messages in giving useful information. Our approaches have been encouraging, but improvements could be made. In hindsight, we could have improved the system by asking people to send pictures as part of their reply message.

Big earthquakes generate a lot of data traffic and our platform might be spending too much time in contacting users. Future works include the improvement of the platform with a regression linear model that contacts users ordered by criteria, maximising the probability of receiving a reply.

We believe that our experiment lays the foundations for future developments in social sensing, and is just a starting point in utilising social media to contact potential eyewitnesses in emergency situations and gain situation awareness.

References

Avvenuti, M., Cresci, S., Marchetti, A., Meletti, C., & Tesconi, M. (2014). EARS (earthquake alert and report system): a real time decision support system for earthquake crisis management. *KDD '14 Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining .* New York.

Avvenuti, M., Cresci, S., Polla, M. N., Marchetti, A., & Tesconi, M. (2014). Earthquake emergency management by social sensing. *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2014 IEEE International Conference on* (p. 587-592). Budapest: IEEE.

Cresci, S., Cimino, A., Dell'Orletta, F., & Tesconi, M. (2015). Crisis Mapping during Natural Disasters via Text Analysis of Social Media Messages. *Web Information Systems Engineering - WISE 2015.* Miami: Springer.

Cresci, S., Tesconi, M., Cimino, A., & Dell'Orletta, F. (2015). A Linguistically-driven Approach to Cross-Event Damage Assessment of Natural Disasters from Social Media Messages. *Proceedings of the 24th International Conference on World Wide Web Companion* (p. 1195-1200). Florence: International World Wide Web Conferences Steering Committee.

Frommberger, L., & Schmid, F. (2013). Mobile4D: crowdsourced disaster alerting and reporting. *ICTD '13 Proceedings of the Sixth International Conference on Information and Communications Technologies and Development: Notes - Volume 2.* New York.

Heinzelman, J., & Waters, C. (2010). *Crowdsourcing Crisis information in Disaster-Affected Haiti.* Washington: Special Report.

Kayal, J. (2011, March 30). *SEISMIC WAVES AND EARTHQUAKE LOCATION*. (Geological Survey of India) Tratto il giorno April 11, 2015 da http://escweb.wr.usgs.gov/share/mooney/SriL.II2.pdf

Lane, N. D., Eisenman, S. B., Musolesi, M., Miluzzo, E., & Campbell, A. T. (2008). Urban sensing systems: opportunistic or participatory? *HotMobile '08 Proceedings of the 9th workshop on Mobile computing systems and applications.* New York.

Rogstadius, J., Vukovic, M., Teixeira, C., Kostakos, V., Karapanos, E., & Laredo, J. (2013). CrisisTracker: Crowdsourced social media curation for disaster awareness. *IBM Journal of Research and Development , 57* (5), 4:1 - 4:13 .

Sakaki, T., Okazaki, M., & Matsuo, Y. (2010). Earthquake Shakes Twitter Users: Real-time Event Detection by Social Sensors. *WWW '10 Proceedings of the 19th international conference on World wide web.* New York.

Winterboer, A., Martens, M. A., Pavlin, G., Groen, F. C., & Evers, V. (2011). DIADEM: a system for collaborative environmental monitoring. *CSCW 'Proceedings of the ACM 2011 conference on Computer supported cooperative work.* New York.

Zhou, A., Qian, W., & Ma, H. (2012). Social media data analysis for revealing collective behaviors. *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining.* New York.

Zuckerberg, M. (2015, October 15). *Safety Check*. (Facebook srl) Tratto il giorno April 13, 2015 da http://newsroom.fb.com/news/2014/10/introducing-safety-check/