

Pulling Information from Social Media in the Aftermath of Unpredictable Disasters

Marco Avvenuti[‡], Fabio Del Vigna^{†‡}, Stefano Cresci^{*†}, Andrea Marchetti[†], Maurizio Tesconi[†]

^{*}Bell Labs, Alcatel-Lucent, Paris, France

[†]Institute of Informatics and Telematics, IIT-CNR, Pisa, Italy

{fabio.delvigna, stefano.cresci, andrea.marchetti, maurizio.tesconi}@iit.cnr.it

[‡]Department of Information Engineering, University of Pisa, Pisa, Italy
marco.avvenuti@unipi.it

Abstract—Social media have become a primary communication channel among people and are continuously overwhelmed by huge volumes of User Generated Content. This is especially true in the aftermath of unpredictable disasters, when users report facts, descriptions and photos of the unfolding event. This material contains actionable information that can greatly help rescuers to achieve a better response to crises, but its volume and variety render manual processing unfeasible. This paper reports the experience we gained from developing and using a web-enabled system for the online detection and monitoring of unpredictable events such as earthquakes and floods. The system captures selected message streams from Twitter and offers decision support functionalities for acquiring situational awareness from textual content and for quantifying the impact of disasters. The software architecture of the system is described and the approaches adopted for messages filtering, emergency detection and emergency monitoring are discussed. For each module, the results of real-world experiments are reported. The modular design makes the system easy configurable and allowed us to conduct experiments on different crises, including Emilia earthquake in 2012 and Genoa flood in 2014. Finally, some possible functionalities relying on the analysis of multimedia information are introduced.

Keywords—Web disaster management; social media mining; human safety; crisis informatics; Twitter

I. INTRODUCTION

Nowadays social media represent a powerful way to investigate preferences, tastes and activities of groups of users. Now more than ever, people continuously share comments and multimedia content about their lives, interests, feelings and opinions [1]. To this regard, platforms such as Twitter, Weibo and Instagram are privileged channels of information diffusion because of their large user base, interactive nature and ease of use while on the move. Furthermore, social media encourage citizens to participate in the process of citizen journalism, which has been proven to be much faster than traditional media in spreading news [2]. Social media users can thus be considered as sensors able to convey valuable information about situations and facts, as asserted by the social sensing (or Human as a Sensor) paradigm [3].

The amount of information shared on social media in the aftermath of mass convergence or emergency events is even bigger, showing bursts of messages describing the unfolding scenario, often complemented with images or videos [4]. It is therefore possible to exploit social sensors either to augment conventional emergency detection and monitoring systems that

rely on pre-deployed ad-hoc sensing equipment, or to substitute them in the lack of such equipment [5]. However, the sheer social media activity in the aftermath of emergencies renders manual analyses of data infeasible. Also, social media data is often noisy and bursty, with text often fragmented and unstructured [6]. Thus, automated analysis of such content asks for preprocessing steps before data can be effectively used. A data-filtering step is usually advisable and machine-learning algorithms can benefit from such operation allowing to achieve overall better performances.

ICTs now enable a new class of decision support systems and tools that aim at improving the capabilities of specialists in detecting and preparing a prompt and effective response to crises. Effectiveness of intervention is closely related to the knowledge of the event's intensity (i.e., its effects on people and infrastructures), whose evaluation can highly benefit from machine-enabled mining of the large volumes of data coming from social sensors. To this purpose, we extended our previous work on earthquake detection tools (EARS [7]) including a new level of filtering and a monitoring component and discussing the effectiveness of text analysis, regressive models and machine-learning classifiers in monitoring unfolding crises with the goal of rapidly detecting involved areas and estimate the extent of damage.

Starting from EARS we developed a high-level, domain-independent software architecture for the online detection and monitoring of emergency events. Based on the Human as a Sensor paradigm and using machine-learning techniques, we built a tool able to gather data from Twitter to detect earthquakes and to support early damage assessment. The modular design we pursued allows users to easily adapt and redeploy the system to various emergencies (e.g., earthquakes, flash floods, wildfires, riots, etc.). We propose a sharp methodology to perform a rough damage assessment after crisis based on machine-learning algorithms applied to User Generated Content (UGC). We show how it is possible to infer damage related information in the aftermath of the crisis by means of regressive analysis on textual content. Finally, in order to assess the solutions to the issues described so far, we implemented a prototypical version of all the previously introduced architectural components and we tested them on real-world data related to recent earthquakes and floods.

The remainder of the paper is organized as follows: Section II provides the architectural overview of the system. Sections III and IV introduce the message filtering and the earthquake detection tasks. Section V details the emergency monitoring

features of our proposed system and Section VI surveys related works. Lastly, Section VII draws conclusions and describes future work directions.

II. ARCHITECTURAL OVERVIEW

In this section, we give a high-level description of a real-time Web system for the detection and monitoring of emergency situations. As shown in Figure 1, the architecture is organized as a pipeline of components, each one performing a crucial step towards the acquisition and the analysis of UGC from social media. The first component of the system is the *Data Capturing* component which is responsible for the acquisition of emergency-related messages. Such operation is performed by querying a specific social media (i.e. Twitter) with keywords related to the kind of emergency to investigate. Analysis of collected data begins with the selection of relevant messages. This step is carried out by the *Message Filtering* component. A filtering step is necessary due to the noisy nature of social media reports. Indeed, not all the messages gathered by the *Data Capturing* component are actual reports of an outbreaking emergency [8]. Subsequently, selected relevant messages are aggregated and analyzed by the *Emergency Detection* component for the detection of an emergency situation. The detection is triggered via statistical analyses when the system records an exceptional growth in the number of (relevant) messages talking about an emergency. After the detection of an emergency, further analyses are carried out by the *Emergency Monitoring* component. The aim of these analyses is to exploit the content of the messages to increase the overall situational awareness and perform automatic preliminary damage assessments. Lastly, the *Alert Dissemination* component alerts decision makers and emergency responders about the detected emergency and provides additional information extracted from emergency-related messages. As shown in Figure 1, the *Message Filtering*, *Emergency Detection* and *Emergency Monitoring* components are colored in dark gray since they represent the core elements of the system. The analyses carried out by these components encompass some of the major engineering research challenges in this field. Such challenges are described in a separate section of this paper, together with proposed solutions and results. Instead, the *Data Capturing* component mainly presents implementation issues, which we briefly discuss in the remainder of this section together with the *Alert Dissemination* component.

Data Capturing. The online system operates on data collected in this phase. Errors at this stage, especially regarding the loss of data, have to be minimized since they will propagate throughout the system thus impairing its ability in detecting and monitoring emergencies. Our implementation of the *Data Capturing* component is based on the Twitter platform. The data acquisition capabilities of this component could be further expanded by implementing additional interfaces to other social media platforms such as Instagram, Weibo, etc. Among the methods offered by Twitter for information extraction, our implementation exploits the Streaming API to open a persistent connection with a stream of 140 characters messages, called tweets. By using this connection, new tweets matching the search criteria can be collected as soon as they are produced. In contrast with the Search API used in the systems described in [9] and [10], which gives access only to a subset of all the tweets produced, the Streaming API allows to capture all the

tweets matching the search criteria. The *Data Capturing* component must meet both *data completeness* and *data specificity* requirements. This represents a trade-off between the number of messages captured and their relevance to the emergency detection and monitoring tasks. Therefore the search criteria used to collect tweets must be carefully selected. In the remainder of this paper we present several results, mainly related to earthquake detection and monitoring. Such results have been obtained through analyses of tweets collected by our *Data Capturing* component with earthquake-related search keywords. We based the keyword selection process on previous studies in this field [7] as well as on other related works [9]. Furthermore, to guarantee the robustness and the reliability of this component we also implemented additional mechanisms to manage rate-limits and generic connection problems in the use of the APIs.

Alert Dissemination. Once an emergency has been detected and critical information has been extracted from the analyzed messages, such information must be promptly delivered to the various stakeholders. These might include decision makers, emergency responders and the involved population. Many studies, such as [11], already focused on the alert dissemination task. An extensive analysis of alert dissemination challenges is out of the scope of this work and we therefore provide only a basic implementation of this component, which is able to deliver automatic email messages upon the detection of an emergency. Our prototypical *Alert Dissemination* component also exploits a dedicated Twitter account to send direct messages to a limited set of users for testing purposes.

III. MESSAGE FILTERING

Using search keywords to query social media platforms allows the acquisition of messages potentially related to an emergency. However, not all the messages captured are actually related to an outbreaking emergency. A message selection procedure only based on the presence of certain keywords in the text is insufficient to ensure that the message is relevant [7]. As previously introduced, some messages can be misleading for the emergency detection and monitoring tasks and must be filtered out as noise. By noise we refer to the messages containing the search keywords but which are not related to the type of emergency under investigation. In [8] we identified two different sources of noise: (i) messages in which the keyword is used with a different meaning than the one related to the searched emergency event and (ii) messages in which the keyword refers to a past emergency. We collected the tweets shown in Figure 2 while looking for earthquake-related messages, and we report them as examples of these two kinds of noisy messages. Excessive levels of noise in collected messages lead to false detections by the system. However, filtering too much may result in the loss of useful messages and thus in the impossibility to detect important events. Therefore this task presents another crucial trade-off related to the accuracy of the filtering process. To overcome this trade-off we propose a solution employing data mining techniques to train a machine-learning classifier. The classifier exploits the characteristics of tweets to discriminate between relevant messages and noisy messages. Specifically, during the offline training phase the classifier is trained using two distinct sets of messages: those relevant and those not relevant (i.e. noisy messages) for an outbreaking emergency event. Messages of the training set

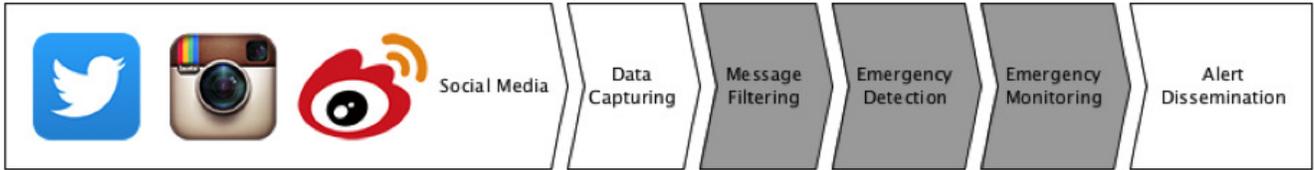


Fig. 1. System architecture



Fig. 2. Twitter examples of noise in collected messages

must be manually annotated to ensure the correctness of the training examples. In the online mode of operation, the trained classifier is able to predict the class (relevant or noisy) of any new message thus implementing the filtering functionality. The classifier bases its decision for the class to assign to a message on a series of features. Our analysis on the messages reporting actual emergencies has highlighted a few interesting characteristics that help distinguish between relevant messages and noisy messages. Relevant messages sent by eyewitnesses are generally very short, they present few punctuation and often contain slang or offensive words. This is due to the fact that social sensors reporting an emergency are usually scared and the contents of their messages tend to represent this emotional state [7].

Experiments. We set up two different experiments on English and Italian tweets in order to assess the performances of the proposed filtering approach. The training set for the English language is composed of more than 5000 manually annotated tweets containing at least one occurrence of the “quake” or “shaking” (sub)strings. Instead, for the Italian language we collected more than 1400 tweets matching the “terremoto” (earthquake) or “scossa” (tremor) (sub)strings. We developed an ad-hoc Web annotation tool specifically designed for the annotation of tweet-based training sets. For the manual annotation we employed 3 human operators and each of them annotated both datasets. We included in the final training sets only those tweets which received the same annotation by all 3 annotators. We also designed 24 distinct features based on the results of our previous analyses about tweets’ characteristics. Such features take into account many structural characteristics of tweets like words count, the presence of mentions, ‘RT’ string in case of retweets, urls, punctuation, uppercase letters and slang/offensive words. We ran feature selection and ranking algorithms to only include in our classifiers the most influential features. Algorithms employed for the feature selection

TABLE I. CONFUSION MATRIX FOR ENGLISH TWEETS

Actual Class	Predicted Class	
	Relevant	Non Relevant
Relevant	1197	501
Non Relevant	400	3371

TABLE II. CONFUSION MATRIX FOR ITALIAN TWEETS

Actual Class	Predicted Class	
	Relevant	Non Relevant
Relevant	654	52
Non Relevant	88	618

are *Information Gain* and *Pearson’s Correlation Coefficient*. As a result 9 features were selected for the English language and 7 features were selected for the Italian language. Both the English and Italian classifiers were generated with the Weka framework [12] using the decision tree J48, corresponding to the Java implementation of the C4.5 algorithm [13] with a 10-fold cross validation. We measured classifiers’ performances by means of standard evaluation metrics such as True Positives (TP) count, True Negatives (TN) count, False Positives (FP) count and False Negatives (FN) count. Classification results are reported in Tables I and II, presenting the so-called confusion matrices. Greyed-out cells in the confusion matrices highlight the numbers of correct classifications (TP and TN), while the other cells represent the numbers of incorrect classifications (FP and FN). Overall filtering results show an Accuracy of 83.5% (4568 correctly classified tweets out of 5469 total tweets) for English tweets and an Accuracy of 90.1% (1272 correctly classified tweets out of 1412 total tweets) for Italian tweets.

IV. EMERGENCY DETECTION

The detection of an emergency is triggered by an exceptional growth in the number of relevant messages captured by the system. The better the filtering phase, the easier is the task of emergency detection. Event detection in social media is a topic that has been widely studied for a broad variety of purposes. Among commonly adopted event detection techniques are Bayesian statistics [9] and peak detection algorithms [14]. In our system we adopt an approach based on a burst detection algorithm. A burst is defined as a large number of occurrences of a phenomenon within a short time window [15]. Figure 3 displays a rug plot of the arrival times of earthquake-related relevant tweets in Italian language, as well as a histogram plot showing their frequency per minute, during a 3.4 magnitude earthquake occurred at 15:47:49, August 9 2014, in Tuscany regional district (Italy). After the occurrence time of the earthquake, indicated by the red vertical dashed line, a big burst of tweets was recorded. These bursts are

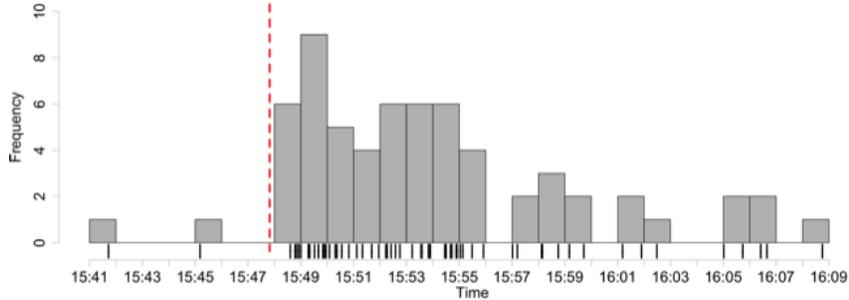


Fig. 3. A burst of tweets reporting a moderate earthquake in Italy

caused by the large number of messages shared on social media by the people who actually felt the shaking. This “bursty” behavior is not constrained only to earthquakes, but the same applies to many other kinds of emergency situations, possibly serving as a red flag for the occurrence of an emergency.

Experiments. Building on the “bursty” characteristics of emergency reports, we adapted the burst detection algorithm originally proposed in [15] and we applied it to the detection of earthquakes in Italy. The detection is performed solely from Twitter data. The detection of a burst is based on the current frequency of relevant messages recorded during a short-term sliding time window. A burst is detected, and consequently the detection of an emergency is triggered, when such frequency exceeds a given threshold. The threshold to trigger a burst depends on a reference frequency calculated over a long-term sliding time window. In our experiments we tried different combinations of settings and we achieved the best detection results with the following settings:

- short-term sliding time window: 1 minute
- long-term sliding time window: 1 week

The threshold for the current frequency of relevant messages is set as ten times the *reference frequency*. This technique is exploited to measure how much the current instantaneous message arrival rate (computed every minute) is large with respect to the average arrival rate (computed over a one week window). We tested our earthquake detection procedure with a dataset of Italian tweets collected over 70 days from 2013-07-19 to 2013-09-23. Over this testing period we were able to detect 47 earthquakes with this procedure. To validate our detections we exploited authoritative earthquake reports released by the Italian National Institute of Geophysics and Volcanology (INGV), which is the government agency responsible for monitoring seismic events in Italy. We therefore cross-checked all our social media-based detections against those obtained by INGV with their seismic network. We classified our earthquake detection results as in the following:

- True Positives (TP): earthquakes detected by our procedure and confirmed by INGV
- False Positives (FP): earthquakes detected by our procedure, but not confirmed by INGV
- False Negatives (FN): earthquakes detected by INGV, but not by our procedure

True Negatives (TN) are meaningless, as it would mean counting the number of earthquakes that did not happen and that our procedure did not detect. In addition we also computed the following evaluation metrics:

- Precision: ratio of correctly detected earthquakes among the total number of detected earthquakes
- Recall: ratio of correctly detected earthquakes among the total number of occurred earthquakes
- F-Measure: harmonic mean of Precision and Recall

Table III shows the final results of social media-based earthquake detections. Results show that the detection of earthquakes with a magnitude < 3 is very difficult (F-Measure $< 50\%$). This is because the majority of these earthquakes are only detected by equipment and not by people. For events with a magnitude equal to or greater than 3.5, results show an overall good performance of the system (F-Measure $> 75\%$) which seems to suggest the effectiveness of the proposed solution. This is significant given that seismic events of a magnitude around 3 are considered “light” earthquakes and are generally perceived only by a very small number of social sensors. These results are promising, especially considering that the proposed technique is adaptable to other emergency scenarios (flash floods, wildfires, riots, etc.) where automatic detection equipment, playing the role of seismographs for seismic events, might not be available. Furthermore, the textual content of tweets often conveys many other kinds of information, such as the presence/lack of damage in a specific location [16], [6]. Mining such content can indeed provide a deep insight into the evolving scenario.

V. EMERGENCY MONITORING

The characteristics and the contents of emergency reports can be exploited not only for automatic emergency detection, but also for the acquisition of valuable insights on the unfolding situation. Emergency monitoring, situational awareness and automatic damage assessment are vast research fields currently attracting much interest from Academia and practitioners alike. To date, several solutions have been proposed to address issues in these fields, however a number of crucial issues still remain unsolved. Yet, the development and adoption of such techniques may provide a great benefit to decision makers and emergency responders and have a huge impact on emergency management procedures.

TABLE III. SOCIAL MEDIA-BASED EARTHQUAKE DETECTION VALIDATION

Magnitude	Earthquakes	TP	FP	FN	Precision	Recall	F-Measure
> 2.0	404	17	30	387	36.17%	4.21%	7.54%
> 2.5	102	16	30	86	34.78%	15.79%	21.62%
> 3.0	26	13	17	13	44.33%	50.00%	46.43%
> 3.5	11	9	3	2	75.00%	81.82%	78.26%
> 4.0	7	5	0	2	100.00%	71.43%	83.33%
> 4.5	2	2	0	0	100.00%	100.00%	100.00%

A. Acquiring situational awareness from textual content

The textual content of emergency reports represents a huge mine of up to date information often coming from eyewitnesses. It has been shown that the contents of emergency reports evolves in time as the events unfold [17]. It is therefore possible to monitor the textual content of the messages shared in social media on the Web to extract knowledge about the unfolding emergencies. Natural Language Processing (NLP) techniques can be employed for a better understanding of textual contents and for the identification and extraction of those textual bits that convey critical and actionable information [16]. Term-clouds are one of the commonly adopted visualization techniques to highlight recurring terms in a corpus. The combination of NLP techniques and term-clouds can yield important results towards the acquisition of a greater situational awareness during crises. Specifically, a corpus of emergency related social media messages can be collected in a real-time fashion via the techniques previously discussed. NLP techniques can be employed to extract terms recurring in the corpus of messages. Such terms can be further contrasted with a statistical baseline created from common messages not related to an emergency. This operation helps eliminate the so-called stop-words (such as articles, prepositions, etc.) as well as other generic recurring terms which are not distinctive of the unfolding emergency. The remaining most frequent terms can be visualized via term-clouds to give an idea of the unfolding scenario. These steps might be repeated in time in order to keep the term-clouds up to date and to follow the evolving situation.

Experiments. We applied this technique to Twitter data related to an strong earthquake and a flooding occurred in Italy. The first case study was the 5.9 magnitude earthquake struck in the Emilia-Romagna region on 20 May 2012 at 04:03 local time, causing 26 deaths and widespread damage¹. We generated the term-clouds at different times analysing the historical corpus of tweets related to the earthquake. The first messages collected talk about an earthquake shake (“scossa di terremoto”) and strong earthquake (“Terremoto forte”), as well as about the towns struck by the event (“Terremoto a Milano”, “Terremoto a Bologna”). The term-cloud generated 20 minutes after the first shake, already brings useful information about the situation: people report being afraid (“paura”), roaming in the streets (“strada”, “gente in strada”) and witnessing damage (“danni”, “botta”) to buildings. We got much information about the strong aftershock (“assestamento”) which occurred at 05:02, and from 06:00 it is possible to observe tweets reporting about casualties (“vittima”, “morto”, “morto a Bondeno”) and collapsed buildings (“crollo”). What can be observed in the tweet corpus is the frequent presence of place names that can help the reader to understand the area that was mainly hit by the seismic shake [6]. Such characteristic improves the value of

the tweet content that, through term-cloud, can be highlighted and endorsed.

In a second, online experiment we tested the technique during a recent severe flash flood which struck the city of Genoa on 9 October 2014 around 23:00 local time². Also in this case we observed a scenario evolution, from the worrying situation caused by the river level due to rains, to flooded areas extensively damaged, with results comparable to those got with the earthquake experiment. During the last part of the crisis, messages were carrying more significant damage related information, talking about cars being swept by water and people moving to higher floors of their houses. In addition to the online extraction of term-clouds for increasing situational awareness, other possible exploitations of NLP techniques on the textual contents of social media emergency reports are the development of machine-learning classifiers trained to automatically detect messages talking about damage [16].

B. Quantifying the impact of disasters

Qualitative approaches described in the previous part of this section give significant contributions to the understanding of unfolding emergencies. However, quantitative approaches could provide even more valuable insights. Being able to exploit the characteristics and the content of emergency reports to promptly produce estimations of a disaster intensity, i.e., its effects on people and things, can prove valuable to those trying to mitigate its consequences. In order to find relations between disaster impact and social media messages, a numeric quantification of disaster intensity is required. Typically, such quantification is evaluated in the days after the disaster and involves specialized personnel who assess its consequences by visiting the locations struck by the event. This information is often available only when the topical part of an emergency is already over. In addition to such numeric quantification, techniques capable of mapping characteristics and content of social media reports into an estimation of disaster intensity, are also required. Statistical regressive and predictive models can go a long way towards disaster impact estimation. These are statistical learning techniques based on offline analyses of historical data. The result of the offline learning process is a model which can be employed in an online mode of operation on unseen data.

Experiments. We performed experiments for the estimation of earthquake intensity. Earthquake intensity indicates the local effects and potential for damage produced by an earthquake. It is verified after the earthquake has occurred with direct surveys on the field. Earthquake intensity is also estimated by exploiting information collected from Web surveys delivered to samples of the involved population. In our experiments we

¹http://en.wikipedia.org/wiki/2012_Northern_Italy_earthquakes

²<http://www.washingtonpost.com/blogs/capital-weather-gang/wp/2014/10/10/torrential-rain-causes-deadly-flash-flooding-in-genoa-italy-video/>

TABLE IV. EARTHQUAKE INTENSITY ESTIMATION RESULTS

Region	R ²	MAE	RMSE
Central & North America	0.4820	0.66	0.82
Central & South America	0.7769	0.38	0.49
Asia & New Zealand	0.5917	0.47	0.72

employed a set of linear regressive models to map tweets' characteristics into the intensity estimations produced by US Geological Survey's (USGS) Web surveys. We collected 90 days of earthquake-related tweets in the English and Spanish language using emergency related keywords, and we extracted 45 distinct features from them. This set of features is organized into 4 main categories. The first category of features includes the number of tweets containing the #earthquake (or similar) hashtag and the average word and character count of tweets. The correlations between the spatial distribution of tweets around the epicenter and the earthquake intensity were assessed whenever geographic data was present. Features belonging to the second category includes the number of distinct accounts that tweeted about the earthquake, the average distance of home accounts from the epicenter, the number of different regions from where the earthquake was perceived. Among features of the third category we computed the rise in the number of tweets shared with regards to a statistical baseline, the average number of tweets per minute in the aftermath of the earthquake, the longest streak of tweets having a maximum delay of 5 seconds between one another. Finally, for features of the fourth category we measured the occurrence of certain prototypical words used to describe earthquake consequences. Such features are used to map the intensity of the 7283 earthquake reported by USGS during the 90 days time window on a 10 levels scale reporting the intensity of each event. Intensity measures the effects of the earthquake on the objects (buildings) located in the affected areas, and may vary from zone to zone according to the ground conditions and conformation. We performed different experiments on 3 geographical regions: Central & North America, Central & South America and Asia & New Zealand. We used multiple linear regression and we ended up in a linear fitting of earthquakes on a 10 degrees linear scale that represents the earthquake intensity estimated using tweets gathered. Results are evaluated with the coefficient of determination metric (R^2), mean absolute error (MAE) and root mean squared error (RMSE). The coefficient of determination is the most used evaluation metric in regression analysis and indicates the "goodness of fit" between the various regressors (i.e. our features) and the estimated quantity (i.e. earthquake intensity). It ranges from 0 to 1, where a value of 1 represents a perfect fit. Table IV shows strong correlations between the characteristics of the messages shared in social media and the intensity of worldwide earthquakes. We achieved a R^2 value of 0.7769 with $RMSE = 0.49$ and $MAE = 0.38$ for the best performing region. This is a good result considering that Web survey's intensity estimations are expressed in a 1 to 10 range of values. A mean absolute error of 0.38 over such range corresponds to a 3.8% error. Despite significant reductions in R^2 values, results from both other regions still present low errors, the worst being equal to 6.6%. Such errors still reflect accurate estimations.

VI. RELATED WORK

In this section we survey related works in the field of social media emergency management. In recent years, several initiatives have been developed both in scientific and in application environments with the aim of exploiting information available on social media. Among these, many works have focused on the analysis of such data for the detection of emergencies. For instance, researchers in [9] had the goal of creating an early warning system for the real-time detection of earthquakes and tornadoes in Japan. The proposed system was able to timely detect 67.9% (53 of 78) of the earthquakes which occurred over two months. U.S. Geological Survey (USGS) efforts towards the development of an earthquake detection system based solely on Twitter data are described in [14]. The proposed solution is evaluated with different settings according to the sensitivity of the event detection module. However, even with its best configuration the system could only detect 48 globally distributed earthquakes out of the 5,175 earthquakes reported during the observation time window. USGS recently announced the official employment of a twitter detection system named TED (Tweet Earthquake Dispatch). As explained by USGS, such detection system proved to be more responsive than those based on seismographs in regions where the number of seismographic stations is low. In [7] we presented the design of a system for the detection and the assessment of the consequences of earthquakes. The proposed solution employs data mining and natural language processing techniques on social media data to enhance situational awareness after seismic events. We also applied text analysis on collected tweet using NLP techniques for filtering operations and for early damage assessment [16], [6]. In [18] authors introduce a tool for decision makers based on spatial data while in [19] authors showed how SOA can be used to integrate different services to support disaster management. Other works related to the emergency management have studied communication patterns and information diffusion in social media in the aftermath of disasters. The study described in [20] shows how social media can be used as a reliable source of spatio-temporal information. Researchers investigate Twitter activity during a major forest fire in the south of France in July 2009. Other similar studies have been carried out in [21], [22] and [23] showing the importance of social media in the communications after a disaster. These studies encourage the exploitation of this information and motivate studies such as the one that we are proposing. In [24] authors propose a complete analysis of emergency management approaches and how to mine and aggregate social media data for further processing.

VII. CONCLUSIONS AND FUTURE WORK

In this paper we discussed techniques for the detection and monitoring of emergencies by exploiting social media data and we proposed some possible solutions to overcome the most relevant challenges in the field. We showed how Web systems can be used to gather data from pervasive social sensors and we provided extensive experimental results derived from the employment of the proposed techniques, mostly in the field of earthquake emergency management. Overall results are promising and seem to encourage the adoption of such techniques. Among the discussed challenges, online emergency monitoring is the one currently deserving more investigation. Techniques for the extraction of knowledge from

the textual and multimedia content of messages, as well as disaster intensity quantification, might give a strong contribution to emergency management procedures. In particular we are going to investigate the potential offered by crisis mapping by exploiting the presence of named entities of locations in tweet corpus using NLP tools with Named Entity Recognition procedures, to link tweet to map regions. This approach may provide better results than term-clouds, or at least complementary results, if we consider the possibility to link tweets and location at a fine grain.

All the techniques introduced to date rely on text analysis and on metadata that complements tweets. We believe that analysis of such data is critical in order to get insights into unfolding emergencies. However, extending such analyses to multimedia content (such as photos and videos) shared in the aftermath of disasters may further improve current social media-based emergency management systems. Indeed, the importance of images towards the assessment of the consequences of disasters has long been asserted, as they can improve situational awareness especially when such imagery data can be coupled with geographic and temporal information [4]. Commonly adopted procedures rely on very high resolution (VHR) satellite images. However, the multimedia content of social media data could be exploited as a complementary source of images in the aftermath of crisis and disasters. In fact, we observe a growing number of disaster-related messages carrying multimedia content, such as pictures of damaged buildings, wildfires, flooded regions, etc. Image classification and clustering techniques should be used for detecting and grouping images carrying sensitive information. Image classification techniques can help selecting only the most informative images, thus reducing the amount of data to be analyzed. Furthermore, being able to automatically group similar images, such as the ones showing the same damaged building, can greatly contribute to the understanding of the unfolding scenario.

REFERENCES

- [1] A. Zhou, W. Qian, and H. Ma, "Social media data analysis for revealing collective behaviors," in *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2012, pp. 1402–1402.
- [2] A. P. Sheth, "Citizen sensing, social signals, and enriching human experience," *IEEE Internet Computing*, vol. 13, no. 4, p. 87, 2009.
- [3] M. Srivastava, T. Abdelzaher, and B. Szymanski, "Human-centric sensing," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 176–197, 2012.
- [4] S. Dashti, L. Palen, M. P. Heris, K. M. Anderson, S. Anderson, and S. Anderson, "Supporting disaster reconnaissance with social media data: A design-oriented case study of the 2013 colorado floods," in *The 11th International Conference on Information Systems for Crisis Response and Management (ISCRAM)*, 2014.
- [5] B. Theilen-Willige, F. E. Behari, H. A. Malek, M. Chaibi, A. Charif, C. Nakhcha, M. A. Ougougdal, M. Ridaoui, and E. Boumaggard, "Remote sensing and gis contribution to the detection of areas susceptible to natural hazards in the safi area, w-morocco," in *Information and Communication Technologies for Disaster Management (ICT-DM), 2014 1st International Conference on*. IEEE, 2014, pp. 1–5.
- [6] S. Cresci, A. Cimino, F. Dell'Orletta, and M. Tesconi, "Crisis mapping during natural disasters via text analysis of social media messages," in *Web Information Systems Engineering-WISE 2015*. Springer, 2015.
- [7] M. Avvenuti, S. Cresci, A. Marchetti, C. Meletti, and M. Tesconi, "Ears (earthquake alert and report system): a real time decision support system for earthquake crisis management," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2014, pp. 1749–1758.
- [8] M. Avvenuti, S. Cresci, M. La Polla, A. Marchetti, and M. Tesconi, "Earthquake emergency management by social sensing," in *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2014 IEEE International Conference on*. IEEE, 2014.
- [9] T. Sakaki, M. Okazaki, and Y. Matsuo, "Earthquake shakes twitter users: real-time event detection by social sensors," in *Proceedings of the 19th international conference on World wide web*. ACM, 2010, pp. 851–860.
- [10] M. A. Cameron, R. Power, B. Robinson, and J. Yin, "Emergency situation awareness from twitter for crisis management," in *Proceedings of the 21st international conference companion on World Wide Web*. ACM, 2012, pp. 695–698.
- [11] S. Nilsson, J. Brynielsson, M. Granåsen, C. Hellgren, S. Lindquist, M. Lundin, M. Narganes Quijano, and J. Trnka, "Making use of new media for pan-european crisis communication," in *Proceedings of the Ninth International Conference on Information Systems for Crisis Response and Management (ISCRAM 2012)*. ISCRAM, 2012.
- [12] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, "The weka data mining software: an update," *ACM SIGKDD explorations newsletter*, vol. 11, no. 1, pp. 10–18, 2009.
- [13] J. R. Quinlan, *C4. 5: programs for machine learning*. Elsevier, 2014.
- [14] P. S. Earle, D. C. Bowden, and M. Guy, "Twitter earthquake detection: earthquake monitoring in a social world," *Annals of Geophysics*, vol. 54, no. 6, 2012.
- [15] R. Ebina, K. Nakamura, and S. Oyanagi, "A real-time burst detection method," in *Tools with Artificial Intelligence (ICTAI), 2011 23rd IEEE International Conference on*. IEEE, 2011, pp. 1040–1046.
- [16] S. Cresci, M. Tesconi, A. Cimino, and F. Dell'Orletta, "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.
- [17] S. Roy Chowdhury, M. Imran, M. R. Asghar, S. Amer-Yahia, and C. Castillo, "Tweet4act: Using incident-specific profiles for classifying crisis-related messages," in *10th International ISCRAM Conference*, 2013.
- [18] K. Derbal, F. Ibtissem, K. Boukhalfa, and Z. Alimazighi, "Spatial data warehouse and geospatial decision making tool for efficient road risk analysis," in *Information and Communication Technologies for Disaster Management (ICT-DM), 2014 1st International Conference on*. IEEE, 2014, pp. 1–7.
- [19] K. Ait Abdelouhab, D. Idoughi, P. Kolski *et al.*, "Agile & user centric soa based service design framework applied in disaster management," in *Information and Communication Technologies for Disaster Management (ICT-DM), 2014 1st International Conference on*. IEEE, 2014.
- [20] B. De Longueville, R. S. Smith, and G. Luraschi, "Omg, from here, i can see the flames!: a use case of mining location based social networks to acquire spatio-temporal data on forest fires," in *Proceedings of the 2009 international workshop on location based social networks*. ACM, 2009, pp. 73–80.
- [21] A. L. Hughes and L. Palen, "Twitter adoption and use in mass convergence and emergency events," *International Journal of Emergency Management*, vol. 6, no. 3, pp. 248–260, 2009.
- [22] A. Murakami and T. Nasukawa, "Tweeting about the tsunami?: mining twitter for information on the tohoku earthquake and tsunami," in *Proceedings of the 21st international conference companion on World Wide Web*. ACM, 2012, pp. 709–710.
- [23] P. Earle, M. Guy, R. Buckmaster, C. Ostrum, S. Horvath, and A. Vaughan, "Omg earthquake! can twitter improve earthquake response?" *Seismological Research Letters*, vol. 81, no. 2, pp. 246–251, 2010.
- [24] M. Imran, C. Castillo, F. Diaz, and S. Vieweg, "Processing social media messages in mass emergency: A survey," *ACM Computing Surveys (CSUR)*, vol. 47, no. 4, p. 67, 2015.