Context Analysis of Volunteered Geographic Information from Social Media Networks to Support Disaster Management: A Case Study on Forest Fires

Frank Ostermann, Institute for the Environment and Sustainability, European Commission Joint Research Center, Ispra, Italy

Laura Spinsanti, Institute for the Environment and Sustainability, European Commission Joint Research Center, Ispra, Italy

ABSTRACT

The increasing diffusion of integrated mobile devices connected with social networks has changed the way geographic information is collected, distributed and used. Several studies have already investigated the impact of social media during crisis events. Although networks of volunteers have demonstrated the ability to curate a large amount of information reliably, this approach faces issues of sustainability and scalability. Therefore, the authors propose a semi-automatic approach to extract volunteered geographic information from social media networks, to evaluate its quality, and thereby to render it useful during any crisis event. The system presented is novel in its approach in that it focuses less on individual pieces of information, and instead uses (geographic) context to determine quality and utility. This paper presents a successful case study on forest fires, but the system architecture is adaptable to different types of crisis events.

Keywords: *Disaster Management, Forest Fires, Geographic Context, Social Networks, Spatio-Temporal Analysis, Volunteer Geographic Informtion (VGI)*

INTRODUCTION

In this paper, we present the evaluation of an automated system to extract and process volunteered geographic information (VGI) from social media networks during crisis events. To this end, we conducted two case studies on forest fires in France, both with encouraging results.

The rationale behind our approach is as follows: The development and adoption of integrated mobile devices like smartphones and their connection with social networks has changed the way information is collected, distributed and used. In the past, information flowed only in one direction, from the "top" of administrations and government agencies to the "bottom" of the broad public. The main communication

DOI: 10.4018/jiscrm.2012100102

channels were traditional broadcasting media like newspapers, radio or television. Horizontal peer-to-peer communication was very limited due to the small reach of the available media (word-of-mouth and letters, later slightly improved through telephones and e-mail). This has changed with the emergence of new information and communication technologies, commonly referred to as the Web 2.0 (O'Reilly, 2005). The Web 2.0 services and platforms have given the public the opportunity to share freely various media through general social networks (Facebook, Google+) or more focused platforms, including text (Blogspot, Wordpress, Twitter), images (Flickr, Picasa, Panoramio), videos (YouTube, Vimeo) and maps (Google-Maps, GeoCommons, MapBox), thus enabling every user to seek and provide information and experiences. This development has also had an impact on the vertical flow of information, with public administrations and agencies adopting the new communication channels to disseminate information to and retrieve information from the public (De Longueville, Annoni, Schade, Ostlaender, & Whitmore, 2010; Palen & Liu, 2007; Puras & Iglesias, 2009; Roche, Propeck-Zimmermann, & Mericskay, 2011).

We can expect that these trends continue in the near future. With increasing wireless Internet access and availability of satellite-based navigation systems such as the Global Positioning System (GPS) in consumer communication devices, a related trend is highly likely: That the amount of public VGI will increase manifold during the coming years.

Several studies have already investigated the impact of (geo)social media during crisis events and show the value for relief workers or coordinators, and the affected population. Examples include wild fires in the United States and France, hurricanes in the United States, the 2010 earthquake in Haiti, and floods in the United Kingdom (Al-Khudhairy, 2010; Bressler, Jennex, & Frost, 2012; De Longueville, Smith, & Luraschi, 2009; Hughes & Palen, 2009; Liu & Palen, 2010; Schade et al., 2011; Sutton, 2010).

However, these studies have also shown that until now these newly created information back channels do not yet integrate well with traditional established emergency response protocols. The two main challenges for integration of VGI into disaster response are the large volume of information, and its unknown and potentially low quality. Consequently, in a crisis management context where inaccurate or incomplete information can have dire consequences including the loss of human life, skepticism by practitioners is understandable. Although networks of volunteers have demonstrated repeatedly the ability to curate a large amount of information reliably, this approach itself faces two main challenges: Sustainability and scalability. There is no guarantee that for any given crisis event, there will be a sufficiently numerous volunteer force. Further, it remains doubtful whether the volunteer approach scales well to the expected increase of information volume.

For these reasons, we propose a semi-automatic approach to filter VGI and evaluate its quality, and thereby render it useful during any crisis event. Human supervision of such a system will remain crucial, because VGI is intrinsically heterogeneous and unstructured: It originates from different persons, using different media such as photographs, text, or video, and authors often overcome device and software limitations in imaginative and unpredictable ways.

The system presented in this paper is novel in its approach in that it focuses less on individual pieces of information, and instead relies largely on (geographic) context to determine the quality and utility of volunteered information. The CONAVI system (CONtextual Analysis of Volunteered Information) is a fully operational proof-of-concept system based on prior research (Ostermann & Spinsanti, 2011). This paper investigates the following two research questions related to the evaluation of CONAVI:

1. Is the reality of forest fire events sufficiently represented in social media to be of any potential use at all, e.g. can we detect real world forest fires through CONAVI?

2. To what extent does the CONAVI system increase the quality of the information that is remaining at the end of the processing pipeline?

This paper structure is as follows: In the next section, we briefly provide the theoretical and conceptual framework for CONAVI, before we sketch the system itself in the section afterwards. The focus of the paper is in the section after, which describes the case studies and the results. Followed by a section which discusses the evaluation results and tries to answers the two research questions, before the paper concludes with an outlook on further work.

BACKGROUND - CREDIBLE AND RELEVANT VGI FOR DISASTER MANAGEMENT

In the previous section, we have introduced the basic premise of our research. In this section, we present the theoretical and conceptual background that led to the design of the CONAVI system, which we describe in the next section.

In order to address our two research questions on overall VGI utility and individual VGI quality, it is important to determine how and when VGI can be of potential use during the management of the disaster. The overall macro-behavior and management of disasters follow four phases: Mitigation, preparedness, response and recovery (see Figure 1).

The mitigation phase focuses on prevention, and includes the assessment of risk and implementation of long-term measures to reduce it. While social media platforms might improve communications generally, we do not see any specific utility at this phase. During the preparedness phase, the actors (agencies, organizations and the public) take concrete measures such as training of emergency crews and the public, stockpiling of relevant supplies, and more. The communication of actual risk to the public and to decision makers is important to help prevent some disasters such as forest fires, and it seems likely that observations from the public can improve the assessment of local risk. In case of an actual disaster, the response phase includes the execution of all necessary and planned measures. Affected citizens can assume important roles as providers and consumers of information: Because members of the public are usually respond first, VGI can be very helpful for decision-makers during the early assessment of the situation. Further, citizens can have

Figure 1. VGI during the disaster management cycles

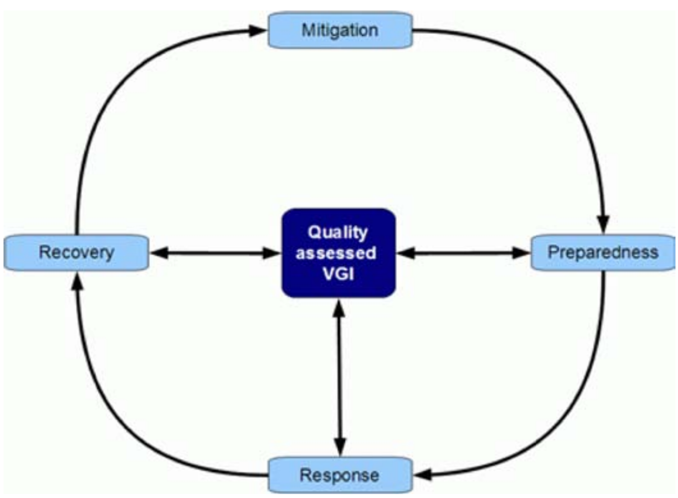

valuable local geographic knowledge and provide almost real-time localized updates, further increasing situational awareness. However, citizens also seek information from various sources, and rely heavily on their own social networks to validate and interpret official information. The idea of citizens as sensors (Goodchild, 2007) would entail that citizens become active participants for the collection of field information, instead of remaining passive victims. This assessment is supported by empirical evidence from research cited in the previous section and in Palen and Liu (2007). People use social media during crisis events for three main reasons: First, to get accurate, local and up-to-date information. Second, to cope with psychological stress resulting from the disaster. Finally, to offer their help to local people through organizing, curating and publishing information, for example on web portals. During the last phase of the disaster management cycle, the recovery phase sees the assessment and repairs of damage, and the creation and distribution of newly gained knowledge for the mitigation phase. Although much of this work relies on remotely sensed images and on the ground surveyors from public agencies and private insurances, reports and images taken by citizens might offer additional information.

Based on this assessment, we investigate the use of VGI during the response phase first. As mentioned in the introduction, public authorities traditionally generate, manage, update and disseminate information in accordance with established protocols and within closed information systems. Any information used or disseminated must be reliable and quickly verifiable, and this validation process has to distinguish between consolidated facts and unverified opinions. This is a complex task, and a proper understanding of the semantics of the volunteered information is essential. Due to this, emergency response agencies have been reluctant to use VGI for critical decision-making. From their perspective, the main weaknesses are a lack of managerial control over the dissemination process, and unknown reliability and trustworthiness of the information (Jennex, 2010). In contrast, mainly NGOs and volunteer groups are pioneering the effort to utilize VGI.

For understanding the semantics of VGI, it is worth looking at the basic heuristics that humans use to deal with new information (Metzger et al. 2010): What do others know and say (social confirmation), what do we know about the situation the information is referring to (expectancies), and how does the information relate to my current needs? The first two heuristics try to determine overall credibility of the new information, while the last one relates to its relevance.

Relevance is a well-established quality metric in the information (retrieval) sciences and denotes how well some information meets the needs of a particular user or use case, e.g. how relevant is the message for determining the extent of a forest fire. The CONAVI system adapts the relevance heuristic in two steps: First, by filtering incoming VGI based on their probability of being on-topic. Second, by assessing the overall information content of a VGI item, in others words its potential to improve situational awareness of a user.

A large body of research addresses the issue of credibility by assessing the information's source characteristics of trustworthiness and expertise. Our approach instead uses a complementary approach and focuses on using the spatio-temporal location and geographic context (semantics) of a VGI item. The CONAVI system enriches VGI with information about its geographic context, such as population density, vegetation, geomorphology, and known disaster events, and compares these "expectancies" ("What do I already know?") with the actual information contained in the VGI. The social confirmation heuristic is adapted by clustering all potentially relevant VGI spatio-temporally ("What do others say?").

The CONAVI system addresses all criteria of successful crisis emergency response systems, i.e. to facilitate clear communications, to improve information and knowledge transfer and use, to improve efficiency and effectiveness of decision-making, and to prevent or at least mitigate information overload through information management (Jennex, 2010).

The next section provides an overview of the current CONAVI system implementation, according to its four main tasks: First, the retrieval and storage of VGI from various sources. Second, the enrichment of the VGI with further information about content, location, and geographic context. Third, the spatio-temporal clustering of the high-quality VGI. Fourth, the representation of the results.

IMPLEMENTATION - CONAVI SYSTEM ARCHITECTURE AND COMPONENTS

This section can only give brief overview, since the paper's scope is limited and focused on the evaluation through the case studies. For more information see (Ostermann & Spinsanti, 2011; Spinsanti & Ostermann, 2011). Additionally, the overall research objective is not to implement a robust, operational system, but instead a proof-of-concept prototype that allows us to evaluate the approach. For this reason, we limited the geographic scope to four European countries (Italy, France, Spain, Portugal), in order to take full advantage of the in-house knowledge available through the European Forest Fire Information Service (EFFIS). The following Figure 2 shows the principal processing steps for using VGI (left column), the state of the VGI (middle column) and the respective CONAVI modules (right column). The last step has not yet been implemented in CONAVI and relies on external modules such as map mash-ups.

The following Figure 3 shows a more detailed version of the implemented system, with the components aligning Figure 3's first three phases, going from left to right instead of top to bottom for reasons of layout. The last phase of Dissemination and Alerting is currently realized through an external web page reading and displaying data from the CONAVI storage. Thus, the top-most layer represents VGI data sources (left), external analysis algorithms (middle), and output (right). The middle layer is the implemented CONAVI modules, with Sensor (left), Analyzer (middle), and Clusterer (right). All data is stored in an Oracle DBMS (lower layer):

CONAVI Sensor - VGI Sources and Retrieval

The CONAVI Sensor acts as passive sensor that reads observations posted or broadcast by citizens and stores them. This is clearly distinct from portals where anyone can actively provide information, and in the terminology of Harvey (2012), our input data is not *volunteered* but merely *contributed*. The distinction is not

Figure 2. Overview of processing steps, information, and CONAVI modules

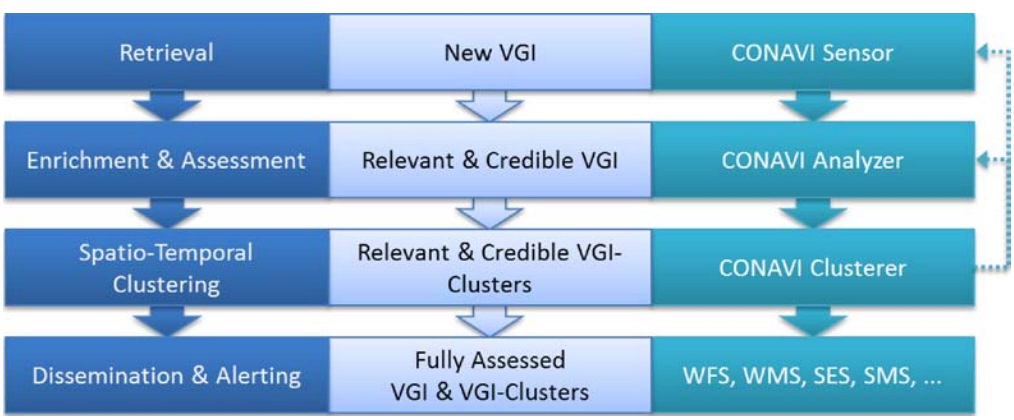

Figure 3. Overall CONAVI system design

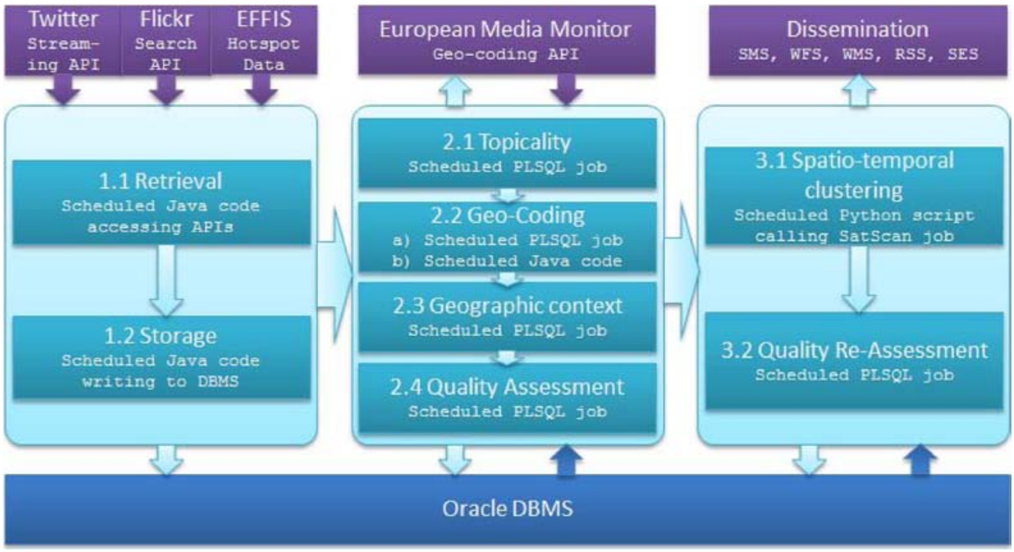

crucial to the CONAVI system, however, and therefore we continue to use the more established term VGI.

We decided on Twitter and Flickr as sources for the following reasons: They have a well-documented API that allows detailed queries without any rate-limits, they have a large potential content base related to forest fires, and they represent textual and visual VGI. Other sites did not have a sufficiently developed or maintained API (e.g. ImageShack), would provide a redundant type of VGI (images), have a comparatively small content base (e.g. YouTube), or have access restrictions on shared content (e.g. Facebook).

The retrieval module accesses Twitter and Flickr APIs using a list of selected keywords. We developed several concepts related to forest fire with input from domain experts, e.g. fire, area and vegetation affected, actors and actions. The set of query keywords was extensive on purpose, as is possible to observe in Table 1. The keywords were selected from regional South-West Europe languages (French, Italian, Spanish (Catalan and Castilian) and Portuguese) plus German and English to include tourists.

The number of received VGI items was still very high, and stretched the limits of our computational infrastructure. An analysis of the 2010 and 2011 data allowed us to narrow down the number of search keywords and reduced noise for 2012. The number of VGI items retrieved per month for 2010 and 2011 is shown in Table 2.

As can be seen, the number of Tweets has increased more than two-fold within one year. The number of Flickr images is not directly comparable between 2010 and 2011 because of refined search terms (see the upcoming sections for details).

CONAVI Analyzer - VGI Enrichment and Assessment

As mentioned in the previous sections, the CONAVI Analyzer enriches the VGI with further information, including topicality based on its content and geographic context based on location. Although both Twitter and Flickr offer to use geographic search parameters, an analysis of our data set shows that only a small fraction of the content (1% for Tweets and

Table 1. Concepts and initial (2010) search keywords

Concepts	Keywords
Fire	Evacuação, evacuació, evacuación, evacuation, évacuation, evacuazione, Evakuation, feu, Feuer, Feuerwehrmann, feux, fire, foc, fogo, fuego, fuoco, incendi, incendie, incendio, incêndio, Waldbrand, wildfire
Area	Ettari, hectárea, hectáreas, hectares, Hektar
Vegetation	Boschivo, broussailles, brush, brushwood, forêt, florestais, forest, forestal, forestale, forêt, Gestrüpp, macchia, maleza, maquia, maquis, matagal, mato, shrub, sterpaglie, Unterholz, vegetação
Actors	Bombeiro, bomber, bombero, canadair, Canadair, elicotteri, elicottero, firefighter, helicopter, hélicop-tér, helicoptero, helicóptero, Helikopter, Hubschrauber, Löschflugzeug, pompier, pompiere, pompieri, vigile, vigili
Actions	Alarm, alarm, alarme, alert, alerta, alerte, allarme, allerta, Warnung

Table 2. Retrieved VGI items per month

Year	2010		2011	
Month	Number of Tweets (in Millions)	Number of Flickr Images	Number of Tweets (in Millions)	Number of Flickr Images
June	---	---	4.9	14,300
July	2.6	179,000	6.7	22,000
August	2.9	164,000	8	21,000
September	3.7	152,000	7.1	14,000
October	---	---	6.2	11,100

20% for Flickr images) is actually geo-coded. Further, the geographic location provided is unreliable, since it depends on the hardware specifications of the device, on the software configuration, the user settings, and the geo-coding done by the social media platform. Using a user's profile location as a substitute is also problematic (Hecht, Hong, Suh, & Chi, 2011). Further, the VGI might be about a place that is different from its origin. As consequence, the CONAVI Analyzer needed a geo-coding module. A geo-coder looks for toponyms (place names) in the text body or metadata of the VGI, and uses a geographic gazetteer to resolve that toponym into coordinates. CONAVI uses the administrative level of communes from the Geographical Information System at the Commission (GISCO)1 database, containing more than 57,000 toponyms in 292 provinces for the four considered countries (Italy, France, Spain and Portugal). Originally, the geocoding step preceded the topicality assessment. However, in our case the computational cost to geocode every VGI item proved too high with more than one million matching attempts per Tweet. While there are computation solutions to this problem, these were out of scope for our proof-of-concept implementation. However, a quick manual assessment of a sample showed that the current CONAVI Sensor calibration led to a high proportion of noise in the form of off-topic VGI. This is not surprising, since keywords can have several meanings (homonymy) depending on the context, or they can be used figuratively. For example, *fire* is used colloquially instead of *wildfire*, but has a range of other meanings. Therefore, the CONAVI Analyzer assesses the topicality first to filter out noise, before searching for toponyms. In the case of crisis events, social media content that does not contain any information on location has no immediate value, so only VGI that was on-topic and geo-coded was passed on to the next module, which enriches the VGI with additional relevant information such as vegetation cover or population density. In the following paragraphs, we provide some more details.

Topicality

The short and unstructured nature of Tweets and Flickr image annotations lead to particular difficulties for established natural language processing tools (Cheng, Caverlee, & Lee, 2010; Finin et al., 2010). Considering that the CONAVI system has to classify incoming VGI as either on-topic or off-topic, it seemed sufficient to look for particular keyword occurrences. In order to establish rules, we manually annotated a set of around 6,000 Tweets and classified them as being about forest fires, or not (including those about other types of fires). Then we programmatically extracted and counted all keyword occurrences. From these, we derived simple rules to determine topicality (see the upcoming sections for details). Later tests with several machine-learning algorithms (J48, Naive Bayes) using the Weka software suite supported our rule set. Only those VGI likely to be about forest fires is sent to the geo-coder (2-3% of the daily retrieved data).

Toponym and Location

The CONAVI Analyzer attempts to geocode the VGI from toponyms in the text body (Tweets), or the title and tags (Flickr images). Despite the simple premise "First find toponyms, then look them up in a gazetteer", geo-coding is not a trivial task, but greatly complicated by variations of the same toponym in different languages, variations in spelling due to special characters, and the same toponym often belonging to many distinct places (disambiguation). We experimented with several third-party web-based geo-coding services, but the results were not satisfactory. One particular challenge was the multilingualism of our data. In the end, we implemented and evaluated a simple search for toponyms: It is based on string matching unigrams of the VGI content with the toponyms contained in the GISCO database of place names for our area of interest (Spain, Portugal, France, Italy). It uses the most detailed common level, the communes, as well as Province names (if no commune was found), and a regular expression to avoid errors such as partial matches, case sensitivity, special characters and the informal writing style. To test this geo-coder, we manually classified the same set of 6,000 Tweets used for the topicality as containing a valid toponym within our geographic scope, containing a valid toponym out of our geographic scope (e.g. New York), or containing no valid toponym. The results were good enough and are reported in the following section. As a further measure to increase geo-coding reliability, the VGI items are sent to an in-house geo-coder from the European Media Monitor (EMM) that has a smaller gazetteer but more sophisticated geo-coding algorithms, and the results compared.

Context

Finally, and central to the CONAVI approach, is enriching the VGI with geographic context by looking up characteristics of the locations identified. In principle, these could be any characteristics found in Spatial Data Infrastructures (SDI) or other databases with a geographic component. In the case of forest fires, examples of relevant context information are distances to known hot spots or forest fires, the population density and predominant vegetation type. Geographic and temporal proximity to a known hotspot or fire increases the credibility considerably, while the absence of any combustible vegetation would decrease it. The population density influences mainly relevance, but also credibility: Forest fires are most dangerous in sufficiently forested terrain with a moderate population density, e.g. urban-rural border. These were the most obvious choices and of course, others could be useful like smartphone penetration, socio-demographic parameters, history of forest fire events, infrastructure conditions, and others. Since our geo-coding is based on the communes in the GISCO dataset, we aggregated raster data sets on population density (DGUR-Degree of Urbanization) and land cover (CORINE data 2006) through zonally spatial analysis. The distance to hot spots used the latest MODIS2 data from the EFFIS3, downloaded at regular intervals and used in a spatial query.

Integrated Context Score

Finally, the CONAVI Analyzer integrates all three scores into a single Integrated Context Score (ICS):

$$ICS = (FOREST\_SCORE + POPULATION\_SCORE + HOTSPOT\_SCORE)/3 \tag{1}$$

$$\mathrm{FOREST\_SCORE} = \begin{cases} 1; & forest\ cover \geq 0.5 \\ forest\ cover{*}2; & forest\ cover < 0.5 \end{cases} \tag{2}$$

$$\mathrm{HOTSPOT\_SCORE} = \begin{cases} 0; & hotspot\ distance > 50\, or\ null \\ 1 - \dfrac{hotspot\, distance}{50}; & hotspot distance \leq 50 \end{cases} \tag{3}$$

$$\mathrm{POPULATION\_SCORE} = \begin{cases} \dfrac{population\ density}{200}; & 0 < population\ density \leq 200 \\ \left(2 - \dfrac{population\ density}{200}\right); & 200 < population density \leq 400 \\ 0; & population\ density > 400 \end{cases} \tag{4}$$

The scoring functions are the result of discussions with domain experts and literature studies. However, we are aware that they are (yet) arbitrary and very likely have a lot of room for improvement. This type of sensor and system calibration needs grounding in and evaluation with real world data, which is the motivation for the case studies presented in the sections below.

CONAVI Clusterer - Spatio-Temporal VGI Clustering and Event Detection

While the Analyzer addresses individual VGI, the Clusterer looks for patterns and (social) confirmation, and tries to identify forest fire events. With spatio-temporal clustering techniques, it is possible to determine which information relate to a potential event (Spinsanti & Ostermann, 2010). The parameters for this clustering need to be carefully chosen and tested in order to reduce false positives and more importantly avoid false negatives, with adjustments for the population density if it varies significantly over the study area. Information on detected patterns can improve the results of the Sensor module (by focusing on particular geographic areas for example) as well as the Analyzer module (by adjusting the scores).

We had to rely on external software for the Clusterer, since Oracle does not provide sufficient support for spatio-temporal clustering. After testing various software (CrimeStatIII, packages of R, ArcGIS, QGIS), we settled on using SatScan4. It has been published about and is in use, and offers the widest variety of possible scan methods, including Space-Time Scan Statistics, Bernoulli and Discrete Poisson Models. In regular intervals, data is exported from the Oracle database, fed into SatScan, processed, the outputs parsed and uploaded back into the database. The detection of events in a near-real time stream of information is a challenging task that needs further investigation. For the moment, each cluster is assigned a preliminary score based on the confidence reported by SatScan. It can be used to rank likely events, which can be investigated further by a human domain expert. However, as a brief investigation has shown, the number of false positive clusters can be high, and further ranking measures are needed. We experimented with several during the case studies, as the next section reports.

RESULTS - CASE STUDIES

This section is the paper's overall focus, presenting the two case studies that provided the data to evaluate the current CONAVI approach and calibrate the processing parameters. To recapitulate, the two key research questions are:

1. Is the reality of forest fire events sufficiently represented in social media to be of any potential use at all, e.g. can we detect real world forest fires through CONAVI?
2. To what extent does the CONAVI system increase the quality of the information that is remaining at the end of the processing pipeline?

The first section deals with a pilot case study conducted early during the project focusing on the first question. The second section shows in-depth the validation case study conducted with data from the second forest fire season in 2011 and addresses both research questions.

Pilot Case Study: France 2010

This case study refers to data harvested at the beginning of the forest fire season 2010 from 16th of July (the first big fire of the season) until 30th of September 2010. During this period, we collected around 8 million Tweets and metadata on almost 700,000 Flickr images using the keywords shown in Table 1.

Clearly, the data volume is too big for even a cursory manual evaluation with a completely random sample. Therefore, we purposely filtered the data geographically (see Table 3). First, we decided on a region to examine. There are several reasons for choosing France: One is to take advantage of the seasonality of forest fires and reduce the amount of potential noise from the rest of the year. Another is the need to have a suitable number of fires recorded by EFFIS. Further, it allows reducing linguistic ambiguities: In all our monitored countries (Italy, France, Spain and Portugal), people produced the majority of VGI in the country's native language. Among Italian, Spanish and Portuguese sets of keywords used to retrieve VGI, there are several common terms. The most important is *incendio*, which means *(forest) fire* in all the three languages. Moreover, Spanish or Portuguese are the official languages in Latin America, introducing VGI outside our European focus. While French is also spoken in Quebec and many African countries, the absolute amount of non-European French VGI is low, and French seems to be 'good' VGI producers5.

The first step was to select only VGI with a French keyword (*feu, incendie, foret, hectares, alarm, alerte, evacuation, canadair, pompier*). The primary objective of this proof-of-concept case study was to find out whether the approach is sensible at all, and therefore it is necessary to know whether the reality of forest fires is represented sufficiently in social media content.

Table 3. Processing steps and data volume for 2010 pilot study

Processing Steps Applied	Data Volume
(0) Keyword filtered retrieval from API	8 million Tweets 700 thousand Flickr images
(1) Filtering for French keywords	611,274 Tweets 61,697 Flickr images
(2) Filtering for *incendie* keyword	6,754 Tweets 458 Flickr images
(3) Successfully geocoded VGI	1,123 Tweets 293 Flickr images
(4) Filtering for location in France	437 Tweets 243 Flickr images

The analysis of manually annotated Tweets (see previous section) showed that almost all French Tweets about forest fires contained the keyword *incendie(s)*. Consequently, we reduced the subset for the French keywords by considering only those with *incendie* in their text, description, or tags, and those with explicit geocoding in the form of coordinates or those with implicit geocoding in the form of toponyms. Because at the time of the case study our own geo-coding was still in development and the sample was sufficiently small, we used Yahoo!Placemaker service with French language settings.

The remaining 680 VGI items were the input for the spatio-temporal clustering. We used SatScan's space-time permutation model, which only needs a case file, i.e. the VGI data itself. There are several options for running the space-time permutation model (Kulldorff, Heffernan, Hartman, Assunção, & Mostashari, 2005), first and foremost whether to base the spatial location of the scanning windows on the VGI cases or on other locations, such as known locations of forest fires. From a conceptual viewpoint, this resembles the choice to look whether we can detect clusters in the data without prior knowledge about possible events, or whether the known events are represented in the data. We opted for both methods to compare them. Regarding other options, we chose the default parameter values (unrestricted cluster size but no spatial cluster overlap), and a set of modified parameters (maximum cluster radius 50 kilometers and spatial overlap possible). Thus, in total there are four sets of results. Each result set consists of a number of likely clusters ($p<=0.0001$ estimated from 9999 Monte Carlo simulation runs).

The clusters analysis reveals two interesting aspects. First, three clusters are identical in all four cases. Second, some clusters contain both types of media (Tweets and photos), but many clusters contain almost exclusively one of the two. Further, we compared the clusters detected with the official fires registered in the EFFIS system. There were 6 major fires in the considered time window in France. Table 4 shows the accuracy of the clustering for the 4 different cases. As True Positives we count the number of clusters that match a known fire, while False Negatives are clusters that match no known fire. Undetected fires are real-world EFFIS-reported fires that were not detected at all.

These results show that without prior knowledge, the system detects 50% of the known fires when restricting spatial overlap of clusters (Case 1), and all of them (and a large number of potentially False positives) when allowing for spatial overlap. An analysis of the False Positive clusters in Case 3 leads to the hypothesis that there may be fires not reported by EFFIS, offering a potential improvement for alert systems based on remote sensing. With prior knowledge, the system is able to detect all known forest fires when we allow for spatial overlap of clusters thus confirming the social network activity around disaster reports.

These were encouraging results and motivated further improvements on the system, examined in the following case study of 2011 data.

Table 4. Accuracy of clustering methods; # one fire had two clusters associated

Case	Total Number of Clusters	True Positives	False Positives	Undetected Fires
1 (default, all locations)	6	3	3	3
2 (default, fire locations)	4	3	1	3
3 (modified, all locations)	74	7#	67	0
4 (modified, fire locations)	8	6	2	0

Validation Case Study: France 2011

The validation case study examines the first question analogously to the pilot case study, i.e. how well the retrieved and processed VGI represents the real world forest fires. In addition, a sampling and manual annotation at several moments of the workflow investigates the second question, i.e. how much CONAVI improves the overall information quality and thereby the utility of VGI for use in a crisis management context.

At the time of this second case study, the CONAVI system was fully implemented, including all sub-modules of the Analyzer. The following Table 5 gives an overview of the processing and corresponding data volumes. All processing steps and results are explained in detail in the following sections.

Retrieval and Resulting Dataset

The case study is confined geographically to France for the same reasons explained in 2010 case study and temporally restricted to the main forest fire season ranging from July to September 2011. All the French keywords were used in this case study. The Sensor calibration for Flickr images was refined based on the results from the manual evaluation of VGI, and changed Flickr API parameters (the original large number of search terms was not supported anymore).

The 3 months dataset contains 21.9 million Tweets and 54,000 Flickr images. Similar to 2011, we filtered this dataset by selecting only those containing at least one of the French keywords. This Processing Step 1 reduced the VGI items to 659,676 Tweets and 39,016 Flickr images.

Filtering and Enrichment

Topicality

Processing Step 1 still leaves a large number of VGI, and we suspect a high level of noise (off-topic) and irrelevant (no toponyms or outside study area) VGI. In Processing Step 2, we assign a topicality score to each VGI item. The score is categorical and contains the classes {A,B,C,D} in decreasing topicality. Category A is very likely about forest fires and contains at least two keywords, one from the set {*foret(s), hectares*} and the other from the set {*feu(x), incendie(s)*}. Category B is probably about forest fires and contains {*incendie(s)*}, while categories C and D are probably not about forest fires, with C containing only one word from the set {*feu(x), hectares, foret(s)*}, and D being all the remaining cases.

The topicality scoring resulted in 2,197 VGI items being Category A, and 23,487 VGI items being Category B. Because the remaining categories C or D were equally large, we

Table 5. Processing steps and data volume for 2011 case study

Processing Steps Applied	Data Volume
(0) Keyword filtered retrieval from API	21.9 million Tweets 54,000 Flickr images
(1) Filtering for French keywords	659,676 Tweets 39,016 Flickr images
(2) Calculating topicality for each VGI and filtering for high scores	25,684 VGI items
(3) Successfully enriched VGI	5,770 VGI items
(4) Spatio-Temporal Clustering	129 clusters containing 2,682 VGI
(5) Excluding smaller clusters (<6 items)	75 clusters containing 2,565 VGI
(6) Filtering for keywords in clusters	11 clusters containing 469 VGI

decided to pass on only Category A and B VGI items to the next processing step.

To find out more about the overall information quality after Processing Step 2, we took a random sample of 2% VGI items using the Oracle DBMS sample function (resulting in 533 VGI items). We then annotated them according to the following criteria: First, whether they were on topic at all. Second, whether they contained any information that might contribute to situational awareness, i.e. what (information about the forest fire), where (detailed information on the location), and who (any actors or persons affected by or fighting the fire). Each situational awareness component present in a VGI item raised that item's situational awareness score by 1 (i.e. each VGI item has a situational awareness score ranging from 0 to 3). Of the 533 VGI items, only 43 (8%) turned out to be about forest fires. By summing up all situational awareness points, we hope to compare information between samples. For this sample, the total score is 66, indicating that most VGI items have little information content.

Toponyms

The 25,684 VGI items with topicality category A or B were sent to the geo-coding module that operates as described in the previous section. It tries to match the VGI content with 32,435 unique commune and 96 province names in France. It found 5,770 VGI items containing toponyms, of them 655 distinct commune names and 39 distinct province names. The external EMM geo-coder reported 3,192 results for French toponyms. However, the geographic scope can range from national level ("France"), to commune names. To be comparable with the GISCO, we limit it to province or commune level, resulting in 1,562 VGI items with toponyms. Table 6 shows the details:

From a comparison of the coverage of the two sets, we found that only in 123 cases (2%), EMM retrieves results, while the GISCO approach retrieve nothing. At this stage of the implementation, we decided to ignore the possibility to include the EMM search in the system for such a small set.

Taking another random sample from this set (1%), we annotated 509 VGI items for the accuracy of the geo-coding and their information content. Out of the 509, 377 (74%) were correctly geo-coded. However, still only 62 (12%) were on-topic, i.e. about forest fires. This is an increase of 50% over the previous sample, with the total situational awareness score almost doubling to 121.

Geographic Context

All of the content that CONAVI was able to geocode was enriched with geographic context as described in the previous sections, i.e. population density, distance to known hotspots and forest cover of the commune. Out of the 5,570 geo-coded VGI, the Analyzer enriched 97% VGI with population density and forest cover, but only 62% had a recorded hotspot within a 72h temporal window. This is probably due to incomplete hotspot data (possibly caused by cloud cover), because we have confirmed fires for days without any hotspots. Moreover, the hotspot distance has a median of 368 km, so it

Table 6. Details of geo-coding approaches

Topicality	VGI Type	CONAVI	EMM
A	Tweets	453	78
	Flickr	18	---
B	Tweets	5,151	1,469
	Flickr	148	15

is not useful in most cases. We consider it to be important for the context analysis only when is less than 50km (as defined in equation 3), that is only 4% of the geo-coded Tweets.

At a first simple analysis, we were not able to observe a statistical correlation of population density and forest cover with the fire clusters; we believe that a finest spatial granularity could be needed. In any case, this was not a specific goal for this study, so further statistical analysis could find more interesting results.

Clustering

Similar to the 2010 study, we used SatScan for the spatio-temporal clustering. Based on the 2010 experiences and some additional testing, the parameters were as follows: Maximum temporal extent of a cluster 10% of study period (i.e. 9 days), no geographic overlap, and no restriction on maximum cluster size. However, the amount of input data was much higher, almost tenfold. This results in 129 clusters containing 2,682 VGI items. While this is another significant reduction, there are clearly more clusters than fire events, indicating a high remaining level of noise. Therefore, we took yet another random sample out of this set, 495 VGI items. Of these, 373 (75%) are correctly geo-coded, but only 67 (14%) are on-topic, and the information content score of 120 is also similar to the before the clustering. It seems that the clustering itself has not had the desired effect.

Analysis and Interpretation of Results

In an effort to improve the results, we excluded small clusters with five or fewer VGI items. Still, 75 clusters remained, with the biggest cluster consisting of 638 VGI items. However, this cluster was not about a forest fire but a fire incident that received a lot of media attention because six persons died in the fire. Three more clusters have between 100 and 200 VGI items, with the remaining cluster sizes ranging from six to 100. There are two choices for further filtering of clusters: By applying a context score threshold, or by re-applying a topicality score threshold.

Context Score Filter

One option for reducing the large number of clusters is to look at their geographic context. Fourteen clusters are located in communes that have no forest cover at all, and 37 clusters (49%) are in areas with less than 20% of forest cover. Further, 65% of the clusters have a radius of 0, meaning that all VGI belonging to the clusters is geocoded to the same commune. Sixteen clusters have not a single hotspot associated with them, while 24 clusters have all of the VGI items associated to a hotspot.

Therefore, promising filters are threshold values for the distances to known hotspots, or the percentage of forest cover in the communes. However, as stated before, the distance to hotspots is difficult to interpret because of some obvious omissions and incompleteness of the hotspot data set. On some days with known forest fires, no hotspots were recorded. Thus, for some VGI there was no hotspot on the same day, while the distances are very large for others. The forest cover is only an approximation and aggregation and not sufficient alone for discriminating clusters. In the case of VGI items geocoded on the level of a province name, the numbers are more uncertain. The large variance and standard deviations for forest cover and hotspot distance on the level of clusters let us decide to rely mostly on the median for calculating a cluster context score as follows:

$$\text{cluster\_context\_score} = (\text{median}(\text{forest\_score}) + \text{median}(\text{population\_score}) + \max(\text{hotspot\_score}))/3 \quad (5)$$

The results for cluster context scores > 0.33 are shown in Figure 4, with the clusters actually related to a forest fire colored in red.

Figure 4. Histogram of cluster scores larger than 0.33; true positives in red, false positives in black

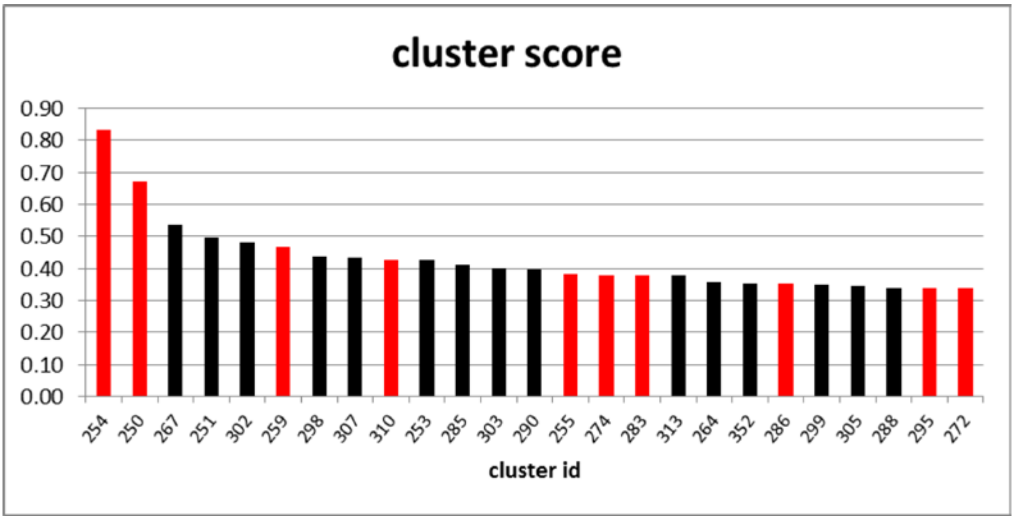

This set contains all the clusters about forest fires except one, but also a number of clusters that are not about forest fires (false positives).

Considering that the spatial context statistics have an undetermined degree of inaccuracy (because of the geocoding, the aggregated values, and possible omissions in the hotspot data), it seems advisable to use the geographic context as ancillary information only and examine more closely a filtering/ranking of the clusters by keyword occurrences.

Topicality Filter

After analyzing the keyword frequencies in the clusters, it was apparent that *incendie* was the most frequently used (accounting for 82% of all keyword occurrences), and it was present in all clusters except one. If we exclude the keywords meaning *fire* (*incendie(s), feu(x)*), the most widely used keyword was *hectares (4%)*, followed by *forêt (3%)*.

From the manual annotation of Tweets, we know that the combination of a fire-related and an area-related keyword is a very good indicator of topicality. Therefore, we filter out the clusters that do not contain *hectares* at least once in one VGI of the cluster. Eleven clusters remain and can be compared manually to the real world fires reported by EFFIS. Another filtering option is to exclude clusters that have only VGI with topicality equal to B (that means only the word *incendie* is present in that VGI, but no any other word related to forest). In this case, 15 clusters remain - the same set filtered before plus four additional clusters.

Similar to the pilot case study and addressing the first research question, we compared the clusters detected by CONAVI with the fires registered in the EFFIS system. There were eight major fires reported by EFFIS in the considered time window, with burned areas varying between 40 and 240 hectares. First comparing cluster locations and time extent and then content, we could establish that CONAVI detected six out of the eight fires. One of the two fires CONAVI did not detect was a relatively large fire with 137 hectares of burned area. A manual web search for any kind of news about a fire in that location in the same time extent did not reveal any information. Geographically, the fire occurred in a big military area for heavy weapon drill, the *Camp Militaire du Poteau*, with public access prohibited. It is possible that there was a wild fire, or weapon training or testing reported as burned area. Colleagues from EFFIS reported a hotspot and a burned area existing from the

Table 7. Details of the filtered clusters

Cluster ID	No of VGI	Province Toponym (No of VGI Containing it)	Commune Toponym (No of VGI Containing it)	Cluster Radius (in Km)	Duration of the Fire (by EFFIS)	Number of Fire Events in the Same Area (by EFFIS)
250	178		Lacanau (178)	0	5 days	1
254	127	Gironde (24)	Bordeaux (45), Cestas (47), Gradignan (4), Luxey (6)	40.42	5 days	3
274	28	Pyrénées-Orientales (5)	Collioure (7), Perpignan (11), Vingrau (5)	46.34	2 days	2
310	9	Hérault (1)	Gabian (8)	30.97	1 day	1
259	29	Lozère		0		
306	7		Latour-de-carol (3), Porta (4)	4.59	4 days	1
286	14	Aveyron (3)	Alban (11)	47.81	2 days	2
272	13		Marie (13)	0	6 days	1
283	11		Luri (11)	0	3 days	1
255	41	Haute-Corse (41)		0	2 days	?
295	12		Solaro (1), Ghisonaccia (2), Ghisoni (3), Palneca (6)	25.59	2 days	3

satellite images. The second fire that was not detected by CONAVI shows up with a manual case insensitive search using the province toponym on Flickr. The commune toponym was present in the Flickr 'description' field that was not used in the study, because very often large chunks of copied text are present, introducing too much noise. There are also a few Tweets on the fire appearing in a cluster discarded by SatScan, because it overlaps geographically with another cluster occurred one month before in the same area. Allowing geographic overlap, however, would introduce much more noise in the form of smaller clusters, as the pilot case study has shown. Regarding the five clusters that cannot be associated with known fires from EFFIS, a manual check revealed that in fact they refer to forest fires, with four of them reporting burned areas greater than 40 hectares. Table 7 lists all clusters.

In two cases the retrieved toponyms were wrong: 'Alban' and 'Marie instead of 'Saint-Alban', 'Sainte-Marie'. These cases result in a wrong geographic positioning of the fire. Six of the clusters have a cluster radius larger than zero, and five larger than 25 km.

In summary, the topicality filter performs well detects almost all major fires, indicating that forest fires well represented in the social media content and detectable by automated methods. Concerning the information content, the results are a significant improvement: Of the 469 VGI items, 431 (92%) are on-topic, with a total situational awareness score of 859.

DISCUSSION AND CONCLUSION

In this section, we will discuss the utility of social media VGI for disaster management on two different levels: First, the utility of the CONAVI system itself, the lessons learned and the potential for improvement. Based on this knowledge and these experiences, we then attempt to discuss the overall utility of social media VGI for disaster management, its opportunities and challenges.

Discussion on the Utility of the CONAVI System

It is safe to say that the CONAVI system performed well in the case studies, detecting all EFFIS-reported fires plus additional ones. We can answer the first research question positively: CONAVI is able to detect the spatio-temporal "reality" of forest fires in social media content. The following two maps show the reduction in data volume from all geo-coded VGI with high topicality score (processing step 3) to clusters of VGI on specific forest fires events (processing step 6):

The map in Figure 5 shows (a) MODIS hotspots, (b) geocoded VGI, (c) vulnerability. Concerning (a), it shows only the MODIS hotspots registered from July-September and likely to be caused by vegetation burning. Concerning (b), the VGI is already pre-filtered according to keyword occurrences, and shows only those likely to be about forest fires, with the dots sized being proportional to the number of VGI from that commune. Finally, (c) is a simple measure derived from population density and predominant vegetation, e.g. low population density and sparse vegetation means a low vulnerability, while high population density and dense vegetation results in high vulnerability

Figure 5. Map showing raw data for France

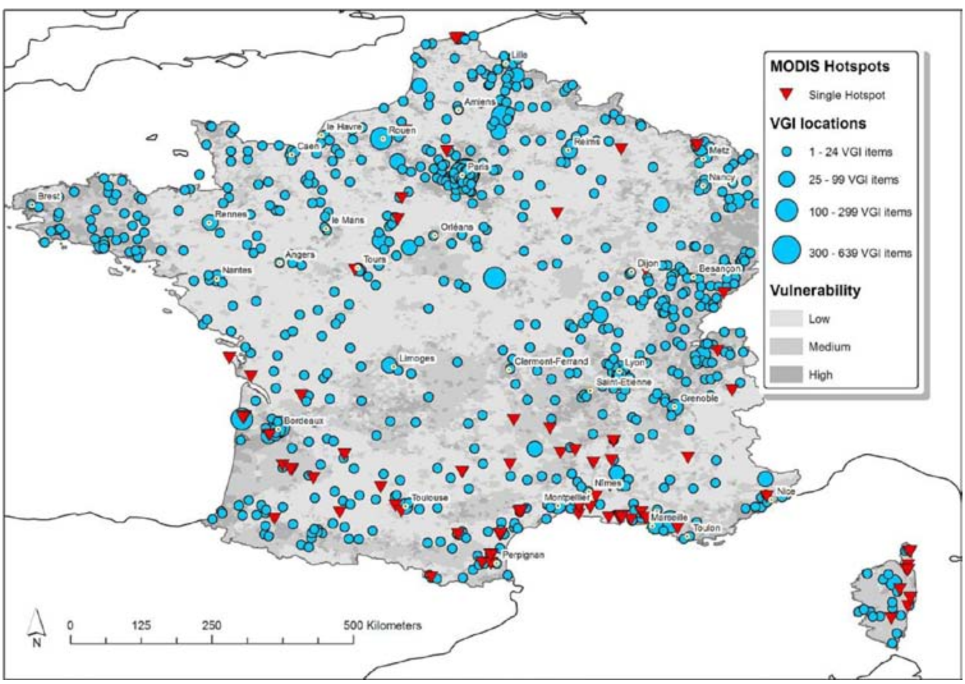

The Map in Figure 6 shows (a) forest fires, (b) processed and filtered VGI clusters, and (c) again vulnerability. The forest fires (a) are divided into two classes: Those reported only by CONAVI, and those reported by both. The dots are sized proportional to the number of VGI in that cluster originating from that commune, with locations belonging to the same cluster sharing the same color.

The above maps also answer the second research question, i.e. to what extent the CONAVI system improves the overall quality of the social media VGI. As we have shown, the last processing step of assessing the topicality of the detected clusters improved the information quality most significantly. The spatio-temporal clustering finds many clusters of content on fire events (not only forest fires), which can have an enormous echo in social media VGI if they occur in centers of major cities, or cause several fatalities. The implications are that although geographic context can improve situational awareness and assessment of remaining clusters, it is topicality scoring and filtering that contributes the most to overall information quality, either before and/or after the spatio-temporal clustering phase.

Crucial for a successful processing with CONAVI is of course the choice of parameters. In the case studies, relatively simple parameters and approaches worked already well. We can expect an even better performance with improved geo-coding methods that are more error-tolerant and disambiguate better. For example, the geocoding mechanism currently uses the centroid of the commune or province detected. As a province can be a large area, the accuracy is reduced. Further improvements include the incorporation of machine learning techniques for VGI classification to support the simple rule-based current method, with the addition of data on the sources, and with

Figure 6. Map showing processed data for France

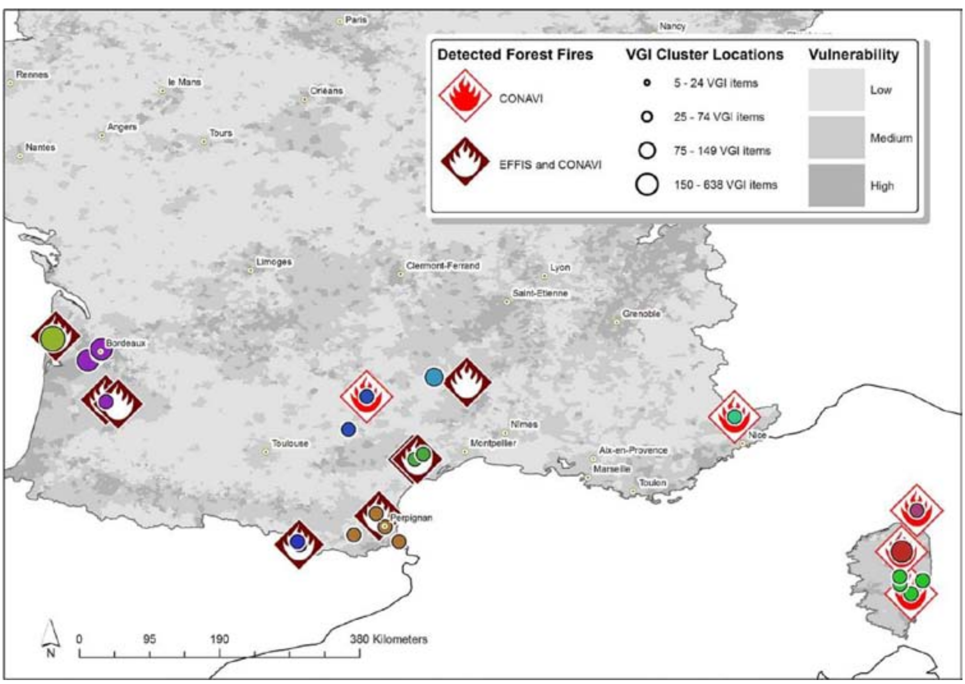

more geographic context data (e.g. forest fire risk indices).

An important part of any evaluation is the assessment of how well the case study performance represents an overall performance of the system in real-time. CONAVI has shown that it is possible to analyze a multi-lingual single-topic crisis event type with standard off-the-shelf hardware and medium-sized enterprise DBMS: Looking at Figure 3, every module that is not a PLSQL job runs on a single Intel Xeon X5550 machine with 12 GB RAM without significant CPU usage. A computational bottleneck is the search for toponyms within the Oracle DBMS, but only because CONAVI implements the simplest approach possible, i.e. brute-force string matching. We expect a fully optimized and indexed DBMS to perform much better, without even including the implementation of advanced techniques such as map-reduce. Thus, from a resource point of view, the CONAVI system provides valuable information with few investments necessary compared to traditional remote sensing and emergency response infrastructure.

The CONAVI system has been designed with the specific use case of forest fires in mind, but the system architecture shows that it is comparatively easy to adapt it to different types of crisis events. The input of domain experts is indispensable to provide information on useful keywords, valuable context information and parameters for the spatio-temporal clustering. Since the full original data is retained, users of the system can easily adjust the displayed information by narrowing it down through keyword and geographic queries, or weighting and filtering of the topicality and geographic context scores. Currently, this is only possible through directly manipulating the DBMS and not through a graphic or web user interface, but this limitation is easy to overcome. There are number of portals providing inspiration, such as Twitcident (http://twitcident.com/).

Conclusions on the Utility of Social Media VGI for Disaster Management

The results from our case studies show that social media content encloses potentially useful information, and can act as additional communication channel for the affected population. The latter suffers from the sheer amount of content and its unstructured nature. Very often, neither the available hardware nor software allow individual members of the public to search social media content efficiently, and make sure all important information is received and read. Therefore, the integration and dissemination of social media content is an important and valuable contribution to the overall disaster management effort. For the dissemination of results from systems such as CONAVI, various avenues are open, including broadcasting via dedicated social media channels, SMS, and web maps. CONAVI has already implemented the latter approach experimentally: It submits the highly likely candidates and clusters to a web application developed by the EFFIS. However, the evaluation is still under way, and this application has not yet been released outside the intranet. Another option would be to send alerts during the response phase through the Sensor Event Service (SES) (Bröring et al., 2011), which can be used to push sensor data (including detected events) to subscribed clients. Its functionality includes the definition of user defined filter criteria, as well as the support of multiple channels for sending the actual notification messages. For decision makers, the use of VGI can have also important legal consequences that should be considered, and administrations might be reluctant to agree on a tight integration of VGI and official data because of such liability issues. Therefore, a loose integration like the Web 2.0 Broker (Núñez-Redó, Díaz, Gil, González, & Huerta, 2011) seems a more promising approach. It keeps the social media VGI content and official authoritative data separate but allows the combined use and display of VGI and SDI. This would also the CONAVI system to become more flexible in respect to the fast-changing nature of social media. With new platforms emerging and existing platforms changing their APIs and term of usage, the Sensor module needs human supervision for adjustment and calibration. For example, the Flickr platform has seen a decrease in use during recent years, with other photo-sharing platforms and social networks like Facebook surpassing it (Douglas, 2011; Offer, 2011). Similar to the input, the output can also easily be adapted to suit the needs of different user groups such as decision makers or members of the public. However, while decision makers work collaboratively using large desktop screens or wall-projections to coordinate crisis response, citizens are likely to be mobile, employing devices with a small screen, and looking for concrete information on evacuation routes, shelters, or the whereabouts of friends and family members. Thus, for the citizens, textual communication via SMS or micro-blogs might be sufficient.

Finally, we argue here that it is not only possible and useful to analyze social media content during crisis events, but in any case necessary. While some argue that social media is a self-correcting medium and the "wisdom of the crowds" will eventually single out and delete false information, this may happen too late in a time-critical situation like a crisis event. Therefore, the authorities charged with managing the disaster can try to counter false information - if they know about it. Since the raw data is still available and can be retrieved in drill down searches, the sources of misinformation can be found. This, however, points to important ethical questions of privacy and consent to use. Unless the user-content has been volunteered for a specific purpose, either through a portal or the use of certain (hash-) tags, all social media content should be considered "contributed" and not "volunteered" (Harvey, 2012).

Summing up the results, for the moment people do not use extensively the automatic geographic references for their information, but this could be encouraged especially in emergencies. Anyway, when reporting an event, such as a forest fire, toponyms are mostly used in the text. The most common behaviour on Twitter is to spread around information as it arrives without any further editing, and although hashtags are used they are not yet a good practice. Finally, different social media have different rules to retrieve, extract, and assess information. A fine-tuning of the system has to consider these differences.

We conclude that the echo of forest fire events in the social web has enough volume to be listened to, and the system can be tuned to receive only the signals we are interested in.

ACKNOWLEDGMENTS

This work profited greatly from the expertise of a number of colleagues at the JRC, most notably EFFIS and EMM. We would like to thank Massimo Craglia from the Digital Earth and Reference Data Unit in particular for invaluable input and support.

REFERENCES

Al-Khudhairy, D. H. A. (2010). Geo-spatial information and technologies in support of EU - crisis management. *International Journal of Digital Earth*, *3*(1), 16. doi:10.1080/17538940903506014.

Bressler, G. H., Jennex, M. E., & Frost, E. G. (2012). Exercise 24: Using social media for crisis response. *The World Financial Review*. Retrieved from http://www.worldfinancialreview.com/?p=1600

Bröring, A., Echterhoff, J., Jirka, S., Simonis, I., Everding, T., & Stasch, C. et al. (2011). New generation sensor web enablement. *Sensors (Basel, Switzerland)*, *11*(3), 2652–2699. doi:10.3390/s110302652 PMID:22163760.

Cheng, Z., Caverlee, J., & Lee, K. (2010). You are where you tweet: A content-based approach to geolocating twitter users. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management* (pp. 759–768). Toronto, ON: ACM.

de De Longueville, B., Smith, R. S., & Luraschi, G. (2009). "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*. doi:http://doi.acm.org/10.1145/1629890.1629907

De Longueville, B., Annoni, A., Schade, S., Ostlaender, N., & Whitmore, C. (2010). Digital Earth's nervous system for crisis events: Real-time sensor web enablement of volunteered geographic information. *International Journal of Digital Earth*, *3*(3), 242–259. doi:10.1080/17538947.2010.484869.

Douglas, N. (2011, August 16). Flickr Is dying. Here's the graph. *Slacktory*. Retrieved from http://slacktory.com/2011/08/flickr-is-dying-graph/

Finin, T., Murnane, W., Karandikar, A., Keller, N., Martineau, J., & Dredze, M. (2010). Annotating named entities in Twitter data with crowdsourcing. In *Proceedings of the 11th Annual Conference of the North American Chapter of the Association for Computational Linguistics*, Los Angeles, US.

Goodchild, M. F. (2007). Citizens as voluntary sensors: Spatial data infrastructure in the world of Web 2.0. *International Journal of Spatial Data Infrastructures Research*, *2*, 24–32.

Harvey, F. (2012). To volunteer or to contribute locational information? Towards truth in labeling for crowdsourced geographic information. In D. Sui, S. Elwood, & M. Goodchild (Eds.), *Crowdsourcing geographic knowledge: Volunteered Geographic Information (VGI) in theory and practice* (pp. 31–42). Berlin, Germany: Springer. doi:10.1007/978-94-007-4587-2_3.

Hecht, B., Hong, L., Suh, B., & Chi, E. H. (2011). Tweets from Justin Bieber's heart: The dynamics of the location field in user profiles. *Proceedings of the 2011 Annual Conference on Human Factors in Computing Systems* (pp. 237–246). Vancouver, Canada: ACM.

Hughes, A. L., & Palen, L. (2009). Twitter adoption and use in mass convergence and emergency events. In J. Landgren & S. Jul (Eds.), *Proceedings of the 6th International ISCRAM Conference*, Gothenburg, Sweden.

Jennex, M. E. (2010). Implementing social media in crisis response using knowledge management. *International Journal of Information Systems for Crisis Response and Management*, *2*(4), 20–32. doi:10.4018/jiscrm.2010100102.

Kulldorff, M., Heffernan, R., Hartman, J., Assunção, R., & Mostashari, F. (2005). A space–time permutation scan statistic for disease outbreak detection. *PLoS Medicine*, *2*(3), e59. doi:10.1371/journal.pmed.0020059 PMID:15719066.

Liu, S. B., & Palen, L. (2010). The new cartographers: Crisis map mashups and the emergence of neogeographic practice. *Cartography and Geographic Information Science*, *37*, 69–90. doi:10.1559/152304010790588098.

Núñez-Redó, M., Díaz, L., Gil, J., González, D., & Huerta, J. (2011). Discovery and integration of web 2.0 content into geospatial information infrastructures: A use case in wild fire monitoring. In *Proceedings of the IFIP WG 8.4/8.9 International Cross Domain Conference on Availability, Reliability and Security for Business, Enterprise and Health Information Systems (ARES'11)* (pp. 50–68). Berlin, Germany: Springer-Verlag. Retrieved from http://dl.acm.org/citation.cfm?id=2033973.2033980

O'Reilly, T. (2005, September 30). What Is web 2.0 - Design patterns and business models for the next generation of software. *O'Reilly.com*. Retrieved from http://oreilly.com/web2/archive/what-is-web-20.html

Offer, J. (2011, June 19). The slow decline of Flickr. *Codehesive*. Retrieved from http://www.codehesive.com/index.php/archive/the-slow-decline-of-flickr/

Ostermann, F. O., & Spinsanti, L. (2011, April 18-21). A conceptual workflow for automatically assessing the quality of volunteered geographic information for crisis management. In S. Geertman, W. Reinhardt, & F. Toppen (Eds.), *Proceedings of 14th AGILE International Conference on Geographic Information Science2011* (pp. 1–6), Utrecht, Netherlands: Association of Geographic Information Laboratories for Europe.

Palen, L., & Liu, S. B. (2007). Citizen communications in crisis: Anticipating a future of ICT-supported public participation. In *Proceedings of the Computer Human Interaction 2007*, San Jose, CA (pp. 727–726).

Puras, J. C., & Iglesias, C. A. (2009). Disasters 2.0. application of web2.0 technologies in emergency situations. In J. Landgren & S. Jul (Eds.), *Proceedings of the 6th International ISCRAM Conference*, Gothenburg, Sweden.

Roche, S., Propeck-Zimmermann, E., & Mericskay, B. (2011). GeoWeb and crisis management: Issues and perspectives of volunteered geographic information. *GeoJournal*, 1–20. doi: doi:10.1007/s10708-011-9423-9.

Schade, S., Díaz, L., Ostermann, F., Spinsanti, L., Luraschi, G., Cox, S., … de Longueville, B. (2011). Citizen-based sensing of crisis events: Sensor web enablement for volunteered geographic information. *Applied Geomatics*, 1–16.

Spinsanti, L., & Ostermann, F. O. (2010, October 11-13). Validation and relevance assessment of volunteered geographic information in the case of forest fires. In C. Corbane, D. Carrion, M. Broglia, & M. Pesaresi (Eds.), *Proceedings of the 2nd International Workshop on Validation of Geo-Information Products For Crisis Management*, Ispra, Italy (pp. 101–108). Luxembourg: Publications Office of the European Union.

Spinsanti, L., & Ostermann, F. O. (2011). Retrieval of volunteered geographic information for forest fires. In M. Melucci, S. Mizzaro, & G. Pasi (Eds.), *Proceedings of the 2nd Italian Information Retrieval (IIR) Workshop*, Milan, Italy (p. 4).

Sutton, J. (2010). Twittering Tennessee: Distributed networks and collaboration following a technological disaster. In S. French (Ed.), *Proceedings of the 7th International Conference on Information Systems for Crisis Response and Management*, Seattle, WA.

ENDNOTES

1 http://epp.eurostat.ec.europa.eu/portal/page/portal/gisco_Geographical_information_maps/introduction
2 Moderate Resolution Imaging Spectroradiometer (MODIS) sensor on NASA's Earth Observing System (EOS) Aqua and Terra satellites.
3 http://effis.jrc.ec.europa.eu/data
4 http://www.satscan.org/
5 http://wearesocial.net/blog/2011/01/twitter-usage-europe/

Frank Ostermann is a post-doctoral researcher at the European Commission Joint Research Centre in the Digital Earth and Reference Data Unit, currently researching the integration of user-generated geographic information and UAV data within existing spatial data infrastructures in the context of crisis management. Prior to joining the JRC, Frank has worked as post-doc at the University of Zurich on analysis and visualization of human spatio-temporal activity patterns, and the retrieval and evaluation of volunteered geographic information from social media platforms.

Laura Spinsanti is a post-doctoral researcher at the European Commission Joint Research Centre in the Digital Earth and Reference Data Unit, currently researching the integration of user-generated geographic information and UAV data within existing spatial data infrastructures in the context of crisis management. Prior to joining the JRC, Laura has worked as post-doc at the Database Laboratory at Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland. She worked on semantic trajectory representation, geographic semantic enrichment and ontology application for knowledge discovery.