

Automatic Sub-Event Detection in Emergency Management Using Social Media *

Daniela Pohl
Institute of Information
Technology
Klagenfurt University, Austria
daniela@itec.uni-
klu.ac.at

Abdelhamid Bouchachia
Institute of Informatics
Systems
Klagenfurt University, Austria
hamid@isys.uni-klu.ac.at

Hermann Hellwagner
Institute of Information
Technology
Klagenfurt University, Austria
hellwagn@itec.uni-
klu.ac.at

ABSTRACT

Emergency management is about assessing critical situations, followed by decision making as a key step. Clearly, information is crucial in this two-step process. The technology of social (multi)media turns out to be an interesting source for collecting information about an emergency situation. In particular, situational information can be captured in form of pictures, videos, or text messages. The present paper investigates the application of multimedia metadata to identify the set of sub-events related to an emergency situation. The used metadata is compiled from Flickr and YouTube during an emergency situation, where the identification of the events relies on clustering. Initial results presented in this paper show how social media data can be used to detect different sub-events in a critical situation.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: Information Search and Retrieval

General Terms

Algorithms, Experimentation

Keywords

Emergency Management, Social Media, Clustering

1. INTRODUCTION

Disaster management consists of three phases: preparedness, response and recovery [12]. Preparedness focuses on monitoring and assessing the situation, whereas response is the relief reaction on a specific event during a disaster to stabilize the situation. Recovery focuses on actions to be taken

*The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n°261817, the BRIDGE project (www.bridgeproject.eu) under the Security Programme SEC-2010.4.2-1: Interoperability of data, systems, tools and equipment.

We also want to thank our colleagues Manfred del Fabro and Prof. Laszlo Böszörményi for their valuable discussion and support.

Copyright is held by the International World Wide Web Conference Committee (IW3C2). Distribution of these papers is limited to classroom use, and personal use by others.

WWW 2012 Companion, April 16–20, 2012, Lyon, France.
ACM 978-1-4503-1230-1/12/04.

after the emergency to get over it. During these phases, it is important to gain as much information as possible about the ongoing situation. This helps manage the aid teams and other resources, assess the situation, and make decisions. Since it is not possible (for the command center and relief teams) to have the 'eyes on the scene' from the very beginning, different information sources shall be used to acquire information about the incident.

Social media platforms (e.g., Flickr, YouTube, Twitter, Facebook) turn out to be a valuable technology for collecting data (e.g., continuous status update, context information) of different types (e.g., pictures, videos, text messages), making such technology very useful for emergency management. In particular, it allows discovering important sub-events, an important issue in emergency management. Studies show the importance of social media in disasters. For instance, Palen [10] states that the consideration of social media was very helpful in the Virginia Tech shooting and the Southern California wildfires. Furthermore, Yates and Paquette [14] describe how the social media is used as knowledge management tool enabling the aid organizations to share information, communicate, and collaborate. In other studies [2], multimedia, especially videos, have been shown to be a valuable source of information during an emergency. Thus, Flickr and YouTube can be important sources to be included in the analysis of such situations.

In this paper, we use Flickr and YouTube data to detect sub-events during a disaster. In particular, we investigate the application of textual metadata collected from Flickr and YouTube during an emergency situation to identify the sub-events. We apply a self-organizing map (SOM) based clustering approach to discover the sub-events.

This paper is structured as follows. Section 2 gives a short overview of related work. Section 3 presents the general framework for our present work. Section 4 outlines our approach for sub-event detection. Section 5 discusses the experiments and findings. The last section summarizes the work and indicates future investigations.

2. RELATED WORK

Social media in emergency cases is getting increasingly important, comparable to its intense utilization in private and commercial areas to communicate different situational, news, and contextual information.

Extensive research has been done on social media in disasters, but mostly focusing on microblogs, like Twitter [13], [8], [9]. For instance, Vieweg et al. [13] give hints for future

extraction methods concerning Twitter applications. The study identifies special types (warnings, damage reports, weather etc.) of messages related to emergencies (Red River floods 2009 and Oklahoma grassfires 2009). The work of Marcus et al. [8] introduces a system called TwitInfo for summarizing and visualizing events based on Twitter messages. The algorithm was evaluated based on soccer events and earthquakes. Mathioudakis and Koudas [9] describe the system TwitterMonitor for real-time trend detection in Twitter streams. Trends are labeled according to the frequency of words in Twitter messages.

Besides textual messages, like microblogs, also visual information is of great importance, especially in emergencies. Flickr, for example, is a valuable source of information to detect events, as the work of Rattenbury et al. shows [11]. Also, Palen shows the usability of Flickr and other photo sharing platforms as important sources [10]. The work of Liu et al. [7] gives a detailed overview of online photo sharing platforms during emergency cases. They analyze disasters and the corresponding activities on the Flickr platform.

Our work demonstrates that the emergency area can profit from information contained in Flickr and YouTube, and points out the usefulness of this information. The results indicate that it is possible to detect sub-events from this pool of information with simple clustering algorithms.

A similar approach for event detection can be found in the work of Becker et al. [1], but we do not focus on geo-referenced information or date/time constructs (like 'picture taken at') within the clustering approach, since in Flickr and YouTube this type of information is rarely available.

Identified sub-events must be summarized such that an overview of the situation can be generated. Examples for summarizing social events (e.g., famous weddings, soccer games) can be found in the work of del Fabro et al. [3], which is primarily the basis of our work.

3. FRAMEWORK

In an emergency situation, information from different perspectives is needed to enable the incident commanders to get the big picture of the situation. Identifying hot spots (sub-events, smaller crises) in disasters helps in assessing the situation and in decision making. Thus, images, videos, voice messages, sensor data, etc. directly from the response teams or from sensors in the field are important sources of information [5]. In addition, information assets published and shared by (casual) bystanders via social media can be valuable contributions to the operational picture.

Manual browsing through this pool of information and identifying sub-events is a cumbersome or often impossible task for people under time pressure. Therefore, we introduce a framework which allows to automatically analyze multimedia data from different social media sources in case of large-scale emergencies (Figure 1). The framework can be applied during a disaster to gain as much information as possible (this can be important for first responders before arriving personally at the scene) or after a disaster to analyze and evaluate the situation. The analysis is conducted using metadata (e.g., tags and title) associated with content found on social media platforms. Currently, the framework facilitates interfaces to YouTube and Flickr. In the future, the framework will be extended to other sources (e.g., Twitter) and additional repositories (marked light gray in Figure 1). The additional repositories include data directly collected by

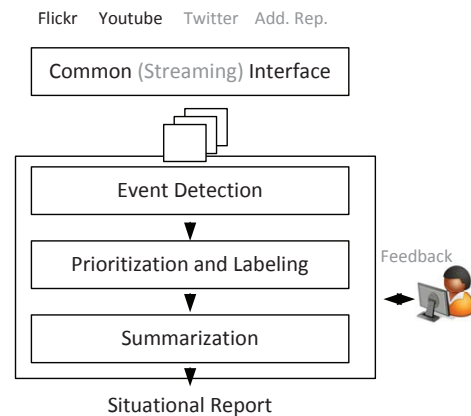


Figure 1: Multimedia (Metadata) Exploration Framework

(professional) first responders on the disaster scene or from news media (e.g. TV, press).

First, the framework performs a *pre-selection* of the data from different repositories using user-supplied keywords, like “UK riots 2011”. In the future, the pre-selection mechanism will be extended through a streaming interface where continuously new pictures, videos, etc. based on the inserted keywords will be processed. This will provide the opportunity to follow the situation in real-time, such that all interested parties continuously can gain new information of the scene. Such a dynamic approach is important for the command center to stay up-to-date. The current work shows static analysis as a proof-of-concept to identify sub-events through social media platforms, especially making use of Flickr and YouTube content.

The results returned via the pre-selection mechanism are the basis for the next step, namely *sub-event detection*. Sub-events are events during a disaster which are separated from other events w.r.t. time or location. For example, during an earthquake, in one place a bridge might collapse, while at the same time in another location also some buildings might be critically damaged. This separates an event into ‘smaller sub-events’ in terms of location.

After the identification of the sub-events, it is necessary to analyze them. This is performed via *labeling* and the assessment of the resulting sub-events. Through a *prioritization* mechanism, different representatives (multimedia documents) of the identified sub-events are selected (*summarization*). The selected documents are included into a situational report which is presented to the user to give an overview of what is going on.

4. SUB-EVENT DETECTION: A CLUSTERING APPROACH

We define a sub-event as an event that occurs at a specific time or location during an emergency case which needs (immediate) emergency response.

The sub-event detection in our framework is performed using a clustering approach based on a *Self Organizing Map* (SOM). SOM is a special case of a neural network without any hidden layer (see Figure 2) [6]. It maps input vectors into a lower-dimensional map. Inputs, which are closely re-




Figure 2: Self Organizing Map (SOM), with a 3x3 map resulting in 9 clusters (adopted from [6])

lated in the original data set, are also closely related in the lower-dimensional map and hence mapped to the same map-unit [6]. The map-unit is described by a weight vector, also called codebook vector.

In our case, the input vectors contain the so-called term frequency-inverse document frequency (tf-idf) of relevant words based on all including documents (as in Becker et al. [1]). Irrelevant words are removed by a stopwords list. The resulting input vectors are normalized and used for the SOM training.

The resulting map-units contain a codebook of input vectors that are located in the nearest neighborhood of the current map-unit, describing it. In the best case, each unit represents a specific sub-event. To identify relevant units, we use the number of hits per unit and we label the resulting SOM units based on the codebook vectors (representing the most frequent words).

The results based on the implemented prototype (realized with WEKA [4] and MATLAB) for sub-event detection are described in the following section.

5. EXPERIMENTS

Our experiments with social media show a high potential for the usage of the data in large-scale emergency management, especially for sub-event detection.

The algorithms operating on social media are also applicable to data collected from (professional) first responders or news media, since the used metadata from the social media is similar to such other data sets.

5.1 Data Sets

The data sets are constructed by means of the pre-selection mechanism described previously. YouTube delivers the first 1,000 relevant videos and Flickr returns the first 4,000 relevant pictures. For our experiments, we chose four emergency situations of the last year. This results in four metadata datasets (see Table 1).

Table 1: Data sets used for sub-event detection

Name/Abbr.	Period	Pics/Videos
Mississippi Flood/MF	04-19 May 2011	2039/442
Oslo Bombing/OB	22 Jul. 2011	31/222
UK Riots/UK	06-10 Aug. 2011	178/274
Hurricane Irene/HI	23-29 Aug. 2011	455/700

The metadata which is considered for each multimedia document are *title*, *description*, and *the corresponding tags*.

We did not use any location and time metadata especially since most of the pictures and videos do not supply such information. The resulting information is filtered based on a date period inserted by the user. Afterwards, the tf-idf mechanism is applied, resulting in input-vectors that are used to train the SOM.

5.2 Findings

By determining clusters from the datasets, it is possible to identify sub-events using simple machine learning mechanisms. Table 2 and Table 3 illustrate the results obtained. Due to limited space, only the results of UK and OB are described in detail. Some of the words (marked with a *) are represented in their principal part (e.g., Polit*: politic, political; Pol*: police; loot*: loot, looting).

Table 2: UK Riots 2011: Clustering results

Cluster (#hits)	Top 4 Words
Cluster 1 (151)	Polit*, Anarch*, <i>Salford</i> , <i>Manchester</i>
Cluster 2 (118)	<i>Birmingham</i> , UK, peopl*, burn*
Cluster 3 (104)	<i>London</i> , loot*, riot*, pol*
Cluster 4 (60)	<i>London</i> , <i>Birmingham</i> , loot*, riot*
Cluster 5 (10)	Polit*, <i>Manchester</i> , str*, <i>Salford</i>
Cluster 6 (9)	Polit*, Anarch*, <i>Salford</i> , <i>Manchester</i>

The results of the UK Riots in 2011 show mainly sub-events *related to* (marked as italic-bold) locations, like Manchester (Salford), Birmingham, and London. Cluster 4 shows also some common aspects of London and Birmingham concerning looting and riots (more keywords identifying the cluster are: fire, UK, etc.). The Clusters 1 and 6 are similar concerning the first four keywords, but differ in weights for other keywords (like violence, UK, etc. which are not shown here).

The same phenomena can be noticed in other datasets concerning MF and HI, whereas these datasets require more clusters to identify sub-events. In HI, clusters related to New York or North Carolina were identified, whereas in MF, clusters for Berwick, Vicksburg, or Memphis were found. This lies in the large scale of such emergency cases, where large areas of the US (some federal states) were affected.

During the Oslo bombing, everything went very fast and in a very short time interval. Most of the data collected is related to the bombing itself and not to the shooting on the island. Cluster 1 in Table 3 refers to the shooting; of course not all data is of the shooting itself, as this was an isolated sub-event and therefore less data was collected. The other clusters, instead, give some information on the attack itself, e.g., that it has something to do with government, injury, car, and explosion.

Table 3: Oslo Bombing 2011: Clustering results

Cluster (#hits)	Top 4 Words
Cluster 1 (131)	terror, attack, <i>shoot*</i> , kil*
Cluster 2 (59)	governm*, <i>Oslo</i> , expl*, bomb*
Cluster 3 (47)	injur*, <i>car</i> , peopl*, kil*
Cluster 4 (16)	expl*, <i>Oslo</i> , governm*, bomb*

The experiments show that large-scale events (like riots or hurricanes) are easier to handle in comparison with smaller ones (like Oslo bombing within half a day). One reason is

that in such a case few data is available due to the fast occurrences of (encapsulated) happenings.

The experiments show promising results in using YouTube and Flickr data as source for sub-event detection in emergency cases. To gain more specific information from the data sets some improvements and refinements in processing the metadata are needed.

5.3 Discussion

The experiments show promising results in sub-event detection by only using textual metadata information. The results show hot spots of activities that are located in a disaster area. This gives an immediate insight of what is going on and, especially, where. It also offers the possibility to perform after-the-fact analysis for understanding a disaster and illustrating hot spots.

In our work we focus on real large-scale events (e.g., floods), as these events offer a wide spectrum of detailed information. But we want to keep an eye on how these methods can be applied to smaller events (e.g., Oslo bombing) too. The used data gives hints that there is a high potential to uncover more specific sub-events from the metadata, in addition to the current results.

We also did experiments w.r.t. changing specific parameters for influencing the results. Therefore, we increased the amount of words used for clustering (UK: from 15 to 40). This shows the need of considering nominal phrases or semantic equal words. For example, a cluster belonging to UK results in: *salford, polit*, dem*, demonstr**. The words *dem** and *demonstr** have the same semantics - belonging to the same concepts describing demonstration, demo, and demonstrators. Our prototype currently realizes a rudimentary process of natural language processing. Hence, some extensions are needed to handle additional similar phrases.

6. CONCLUSIONS

In this paper, we show how multimedia metadata of social platforms can be used to identify sub-events in large-scale emergency cases. Research showed that social media is becoming increasingly important in emergency situations for providing a general overview. Manually browsing through this information by persons working under stress in an emergency situation is an impossible way of handling such information. Hence, we suggest a framework for automatically identifying important sub-events which need (immediate) response by relief units. This framework can be used in the command center to identify and plan important activities. We perform the sub-event detection via a SOM-based clustering approach. Our experiments using Flickr and YouTube data show interesting results in using information from social media. In future work, we are planning to extend our framework, focusing on improvements on natural language processing, adding additional data sources (e.g., Twitter), and enriching it with streaming facilities.

7. REFERENCES

- [1] H. Becker, M. Naaman, and L. Gravano. Learning Similarity Metrics for Event Identification in Social Media. In *Proc. of the 3rd ACM Inter. Conf. on Web Search and Data Mining*, WSDM '10, pages 291–300, New York, USA, 2010. ACM.
- [2] F. Bergstrand and J. Landgren. Information Sharing Using Live Video in Emergency Response Work. In *Proc. of the Information Systems for Crisis Response and Management Conf. (ISCRAM 2009)*, 2009.
- [3] M. del Fabro and L. Böszörmenyi. Summarization and Presentation of Real-Life Events Using Community-Contributed Content. In *18th Inter. Conf. on Advances in Multimedia Modeling*, pages 630–632. Springer Berlin Heidelberg, 2012.
- [4] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The WEKA Data Mining Software: An Update. *SIGKDD Explor. Newsl.*, 11:10–18, November 2009.
- [5] J. Lachner and H. Hellwagner. Information and Communication Systems for Mobile Emergency Response. In W. Aalst, J. Mylopoulos, N. M. Sadeh, M. J. Shaw, C. Szyperski, R. Kaschek, C. Kop, C. Steinberger, and G. Fliedl, editors, *Information Systems and e-Business Technologies*, volume 5 of *Lecture Notes in Business Information Processing*, pages 213–224. Springer Berlin Heidelberg, 2008.
- [6] D. T. Larose. *Discovering Knowledge in Data: An Introduction to Data Mining*. John Wiley & Sons, Inc., Hoboken, New Jersey, 2005.
- [7] S. Liu, L. Palen, J. Sutton, A. Hughes, and S. Vieweg. In Search of the Bigger Picture: The Emergent Role of On-Line Photo-Sharing in Times of Disaster. In *Proc. of the Information Systems for Crisis Response and Management Conf. (ISCRAM 2008)*, 2008.
- [8] A. Marcus, M. S. Bernstein, O. Badar, D. R. Karger, S. Madden, and R. C. Miller. Twitinfo: Aggregating and Visualizing Microblogs for Event Exploration. In *Proc. of the 2011 Annual Conf. on Human Factors in Computing Systems*, CHI '11, pages 227–236, New York, USA, 2011. ACM.
- [9] M. Mathioudakis and N. Koudas. TwitterMonitor: Trend Detection over the Twitter Stream. In *Proc. of the 2010 Inter. Conf. on Management of Data*, SIGMOD '10, pages 1155–1158, New York, USA, 2010. ACM.
- [10] L. Palen. Online Social Media in Crisis Events. *EDUCAUSE Quarterly (EQ)*, 31(3):76 – 78, 2008.
- [11] T. Rattenbury, N. Good, and M. Naaman. Towards Automatic Extraction of Event and Place Semantics from Flickr Tags. In *Proc. of the 30th Inter. ACM SIGIR Conf. on Research and Development in Information Retrieval*, SIGIR '07, pages 103–110, New York, USA, 2007. ACM.
- [12] A. Sagun. *Advanced ICTs for Disaster Management and Threat Detection: Collaborative and Distributed Frameworks*, chapter Efficient Deployment of ICT Tools in Disaster Management Process, pages 95–107. IGI Global, 2010.
- [13] S. Vieweg, A. L. Hughes, K. Starbird, and L. Palen. Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness. In *Proc. of the 28th Inter. Conf. on Human Factors in Computing Systems*, CHI '10, pages 1079–1088, New York, USA, 2010. ACM.
- [14] D. Yates and S. Paquette. Emergency Knowledge Management and Social Media Technologies: A Case Study of the 2010 Haitian Earthquake. *Inter. Journal of Information Management*, 31(1):6 – 13, 2011.