

Avalanche: Prepare, Manage, and Understand Crisis Situations Using Social Media Analytics

Sven Schaust

AGT Group (R&D) GmbH
sschaust@agtinternational.com

Maximilian Walther

AGT Group (R&D) GmbH
mwalth@agtinternational.com

Michael Kaiser

AGT Group (R&D) GmbH
mkaiser@agtinternational.com

ABSTRACT

The recent rise of Social Media services has created huge streams of information which can be very valuable in a variety of scenarios. One specific scenario that has received interest is how Social Media analytics can be beneficial in crisis situations. In this paper, we describe our vision for a Social Media-ready command and control center. As motivation for our work, we present a short analysis of tweets issued in NYC during Hurricane Sandy in late October 2012 and we give an overview of the architecture of our event detection sub-system.

Keywords

Social Media Analytics, Crisis Management, Event Detection, Twitter, Natural hazard

INTRODUCTION

At the time of writing, on Twitter alone, more than 500 million posts are issued every day. A large part of these originate from private users who describe how they currently feel, what they are doing, or what is happening around them. We are only starting to understand how to leverage the potential of these real-time information streams. One scenario that has received growing interest in recent years is how Social Media analytics can be beneficial in crisis situations: How can Social Media improve situational awareness, and can it help to manage disaster?

In this paper we report on a semi-manual analysis concerning tweets issued in the New York Metropolitan region during the time when hurricane Sandy hit. We are looking to find evidence that supports the hypothesis that it is valuable for crisis management agencies to monitor Social Media streams during such a crisis. We then present a high-level overview of the “Avalanche” system, which we are currently building. Finally, we present technical details on our approach to detecting breaking events in Social Media streams, the first phase in our three phase approach.

RELATED WORK

Several studies have been conducted to evaluate the capabilities of micro-blogging services during natural disasters. In (Mendoza, Poblete and Castillo, 2010), Twitter usage is analyzed in the context of an earthquake which occurred off coast of Chile. The authors observed after the event that people tweeted about tsunami alerts, missing or deceased persons, etc. They also surveyed the trustworthiness of tweets and verified that false rumors were much more often questioned than confirmed truths. In (Vieweg, Hughes, Starbird and Alen, 2010), the authors evaluate the significance of Twitter as a contribution to the situational awareness picture of two natural disaster events (Oklahoma Grassfires and Red River flood in 2009). They checked for the occurrence of geographical information (geo-tags, street names, etc.) and discovered that about 82% of the users posted at least one tweet containing geo-location information. The authors of (Sakaki, Okazaki and Matsuo, 2010) propose a micro-blogging-based earthquake reporting system. Tweet analysis was performed to identify messages on Japanese earthquakes in real-time. The system proved to be reliable for earthquakes with a seismic intensity above 3 since 96% of these earthquakes were detected 80% of them while they happened. In (Heverin and Zach, 2010), the impact on Twitter during and after a deadly shooting killing four police officers in Seattle-

Tacoma, Washington, has been evaluated. The authors could show that Twitter was used by citizens and news media organizations to share event-related information during and after the event. Another important characteristic is the near-real-time nature of information provided by users since the situational awareness can only be improved if the information is provided temporally close to the event. Twitcident (Abel, Hauff, Houben, Stronkman and Tao, 2012) for example enables searching, and analyzing Twitter information streams during incidents. It listens to an emergency broadcast service which provides information about local incidents. Whenever a message comes in, it searches for related tweets which are semantically extended to allow for effective filtering during a user search. A case study was performed with Twitcident on approx. 97,000 tweets that got published around a storm event at Belgium festival (Terpstra et al., 2012). The results indicated an exponential rise of tweets during the storm event and how the provided information could be helpful for crisis management. Regarding command and control (C2) for crisis management, one example is the advanced C2 system for command post applications and first responders (Bakopoulos et al., 2011), that specifically leverages data fusion of available at-the-scene sensors, e.g. annotated video material or gas and fire sensors. While the approach tries to retrieve a meaningful situational awareness picture for resource planning and handling, it neglects inputs from other sources such as affected persons via social media.

MOTIVATION: WHAT DO PEOPLE TWEET WHEN DISASTER STRIKES? THE CASE OF “SANDY”

In order to better understand how people use Twitter in crisis situations, we collected and analyzed all tweets that could be linked to the New York metropolitan area via their location information during a seven day timeframe in October 2012, around the time when hurricane “Sandy” hit the region. Table 1 shows the total number of tweets issued in this region per day as well as the number of candidate tweets related to the hurricane. Candidate tweets are tweets that include at least one of the following keywords: “sandy”, “hurricane”, “storm”, “evacuation”, “flood”, “building”, “collapsed”, “power”, “outage” and “fire”. It can be observed that once the hurricane hit the NYC region on October 29th, the number of candidate tweets rose significantly from 42,000 messages on October 27th to over 200,000 messages on October 29th. Over 194,000 messages related to the hurricane were posted on October 30th. Tweets about floods, collapsed buildings, or power outages also rose significantly. It is also interesting that the overall number of tweets declined. The most likely explanation for this observation is that power outages that hit NY during this time made tweeting from some locations impossible. We additionally plotted the geo-located tweets about “flood” for the period 26th to 31st, showing how tweets significantly rose for certain areas once the storm hit the area on October 29th and 30th, see Figure 1.

Date	Total no. of Tweets	Candidate Tweets	“Flood”	“Fire”	“Building”	“Power”
2012/10/25	2,903,664	7,236	448	47	11	131
2012/10/26	3,051,310	24,620	956	62	36	461
2012/10/27	3,142,228	42,306	1,541	68	48	1,025
2012/10/28	3,169,843	108,408	5,194	1,055	195	2,102
2012/10/29	2,896,889	200,340	12,018	823	1,552	5,646
2012/10/30	2,635,138	194,247	34,332	3,773	2,270	14,542
2012/10/31	3,029,506	83,410	9,239	613	790	6,117

Table 1: Tweets in New York region with respect to total number of tweets, candidate tweets as well as tweets mentioning explicitly “flood”, “fire”, “building collapsed”, or “power outage”.

Tweets collected during the hurricane provide information, often including pictures, of events such as:

- A crane collapsing on a construction site near 57th street
- A part of an apartment house collapsing in Borough Park, Brooklyn
- A fire in Breezy Point, Queens
- Flooded tunnels, streets, and apartments in various areas
- Power outages in various areas

For these events we were able to identify tweets in the collected data that allow monitoring the event during the hurricane as well as provide post-incident report coverage showing the destruction caused by the events. Below, we list a few tweets for the crane event and the fire in Breezy Point which are suitable to illustrate how direct

observations contained in tweets can help to improve the situational awareness of operators:

Crane Tweets (overall 950 tweets were found for October 29th):

- 29.10.2012 18:41:56; Wow. Right down the street from me. #Sandy-damaged crane on new 57th St. hi-rise dangling in wind.
- 29.10.2012 18:46:20; Be careful on West 57th St as there is a crane dangling from the rooftop! #HurricaneSandy #Sandy #NYC
- 29.10.2012 18:50:31; From my window I can see the top of a crane hanging off, 60 stories up...not good news if that comes off #Sandy
- 29.10.2012 18:57:17; Curious to see what happens with the dangling crane on 57th between 6th and 7th Staying clear of that area for a while #HurricaneSandy

Breezy Point Fire Tweets (overall 1406 tweets were found for October 30th):

- 30 Oct 2012 01:51:11; A TV news crew covering the storm is trapped by rising water and nearby fire @ 147 Oceanside in Breezy Point - pls RT #sandy #fdny #nypd
- 30 Oct 2012 03:19:35; There are several fires burning in Breezy Point and Broad Channel, but the FDNY cannot reach them because of the flooding. #sandy
- 30 Oct 2012 06:00:58 ; Fire moving 130st street north and west toward Cronstant Ave in Rockaway. Fire at 209 street in Breezy. FDNY cannot get to Breezy. #sandy
- 30 Oct 2012 22:16:16 ; Never seen anything like this in my life. #sandy @ Breezy Point, NY <http://t.co/>

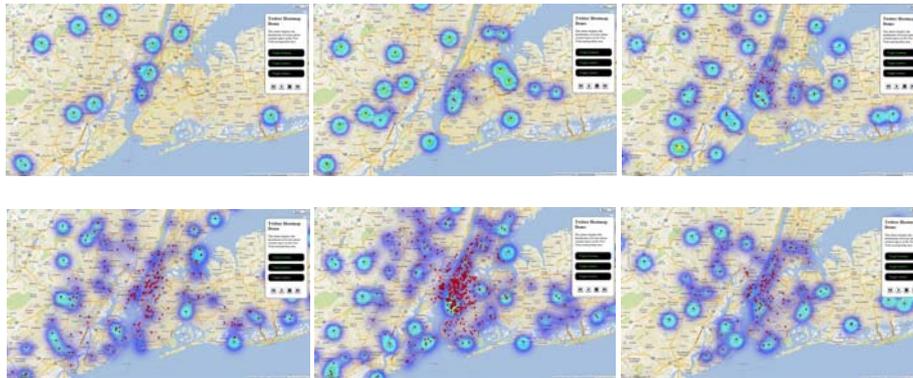


Figure 1: Geo annotated heat maps about “flood” plotted on Google maps for the NY Metropolitan area from upper left to lower right for 26th to 31st October.

THE THREE PHASES OF C2 RELATED CRISIS HANDLING

Avalanche assumes a three phases approach to Crisis Management, which will be detailed in the following.

Phase 1: Natural disaster that strike urban areas usually cause many smaller emergencies e.g. fires, collapses of buildings or power cuts (see also the analysis presented in the last section). If the information contained in Social Media is to be useful for emergency services, such individual events should be automatically determined and classified according to their type (fire, power cut etc.) by the underlying Social Media Analysis platform and presented to the operator in a concise manner. The operator should have the possibility to select whether he wants to dismiss an event or monitor it further. In our approach the automatic detection of breaking events constitutes the first step in our three-staged effort to support situational awareness, see Figure 2.

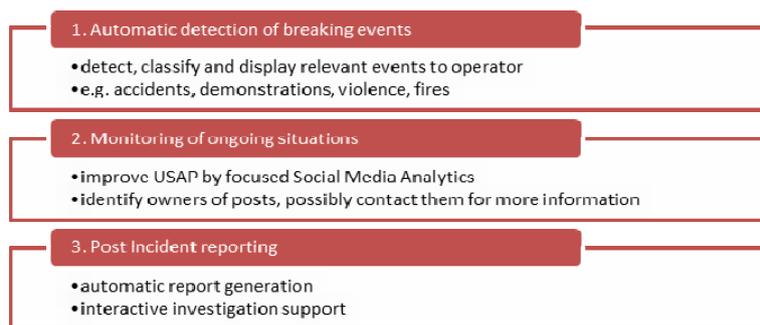


Figure 2: The three Key Phases for C2 – Detect Breaking Events, Monitor Events, and Post-Incident Reporting.

Phase 2: If an operator decides to monitor an event the second stage is reached. Here the system automatically determines whether previously unknown, relevant information is shared about the already detected events on Social Media sites. If this is the case the operator is made aware of it, in a dedicated window. In some cases, even an already detected event can raise an alarm that the operator has to react to, for example in cases where casualties are reported. Note that social media is typically only one of many data sources at the C2, hence a data fusion process is used to present the operator a combined situational awareness for each monitored event.

Phase 3: The third phase in our approach starts not until the worst of the crisis is over. It is concerned with the generation of automatic or interactive reports of all relevant information found on Social Media sites. For some agencies this is required by law. Besides that, this step also allows identification of witnesses of the events that can be contacted via their Social Media accounts.

EVENT DETECTION ARCHITECTURE

As detailed in the last section, breaking event detection is the first of three phases in our setup. Its purpose is to alert the C2 operator about incidents of which he is unaware at this point in time. Our system monitors all posts on Twitter issued within a given geographic region and identifies places that show a high amount of activity. In a second processing step, we analyze the resulting spatio-temporal clusters of posts with a Machine Learning component in order to detect whether they constitute real-world events or not. The detected events are then displayed to the operator on a map, at the location where they happen and while they happen. In this paper, we will focus on architectural considerations of our system. For an in-depth description of our algorithm and an evaluation of its classification performance, please refer to (Walther and Kaiser, 2013).

Requirements for social media analytics architectures generally include support for large incoming data streams (“Big Data”) which have to be processed in real-time. In our case, because we focus on a specific geographic region, this challenge is somewhat reduced. Nevertheless, we deal with a significant number of tweets that have to be stored and processed. Additionally, we were from the outset considering to scale the scenario to multiple and/or larger regions. Hence, our architecture supports real-time processing of tweets and centers around a MongoDB database in which we keep tweets for 24 hours, intermediate results or tweet clusters (event candidates) for 48 hours, and final results (events) for 7 days.

Our architecture consists of a set of components in order to fetch and process social media posts. A high-level overview is given in Figure 3. The components outside of the actual event detection platform are provided by social media sites. In general, such networks offer APIs in order to access their contents. Examples are Facebook, Twitter, and Foursquare. Besides adding search terms and many other options, many APIs provide the possibility to filter posts based on geo-locations. As mentioned earlier, this allows for focusing on certain geographical areas and hence limits the amount of data to be processed.

The architecture of the event detection platform itself consists of several independent modules to query and update the employed NoSQL database. Each module is designed in a way so that multiple instances can be run in parallel, should it be necessary. The modules themselves can be divided into four groups, namely, fetching and preprocessing, clustering, event detection, and visualization.

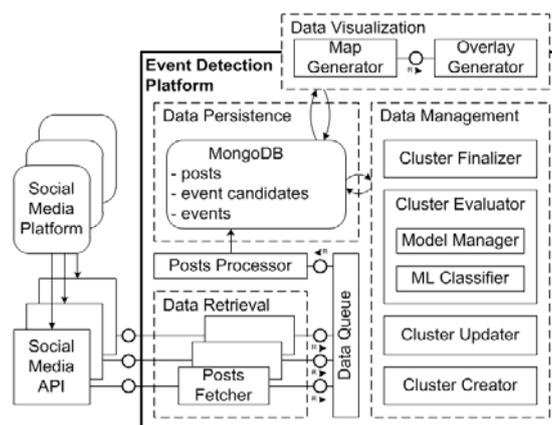


Figure 3: Architecture Overview of the System.

Fetching social media posts is an ongoing process which is executed in the data retrieval component. Each incoming post is put into a data queue. Hereafter, a posts processor takes the posts from the queue, transforms them into a unified format, and pushes the results to the central NoSQL database (MongoDB).

In a next step, a cluster creator merges collected close-by posts to create event candidates. It uses a function which determines the proximity of posts based on their temporal and spatial distance. E.g., if more than three tweets have been posted within a radius of 200 meters in the last 30 minutes; those tweets form a new event candidate. A cluster updater inserts incoming posts into existing clusters or merges similar clusters.

With the event candidates at hand, a cluster evaluator uses machine learning algorithms to decide about each candidate as to be or not to be a real-world event. The evaluator's classifier has been trained with a set of classified event candidates. The cluster finalizer enhances detected events with additional information, such as referenced photos, videos, etc. For more details, please refer to (Walther and Kaisser, 2013).

In a final step, the detected events are displayed in a command and control center. This allows for the creation of an improved unified situational awareness picture leveraging all available sources including social media.

CONCLUSION AND FUTURE WORK

We have presented our vision for a social media-ready command and control center during the following three phases: before the event, during the event, and after the event. We briefly presented a first semi-manual analysis of the tweets during hurricane Sandy, showing that the provided information can be used as a valuable input for a C2 operator. We also provided a brief description of the architecture of our system as far as the first phase, event detection, is concerned. The second and third phases of our system are still work in progress and thus not presented here. For post-incident handling, we currently develop a system that enables automatic generation of summaries for all relevant events in a given time frame.

REFERENCES

1. Abel, F., Hauff, C., Houben, G.-J., Stronkman, R. and Tao, K. (2012) Semantics + Filtering + Search = Twitcident. Exploring Information in Social Web Streams, *Proceedings of the 23rd Conference on Hypertext and Social Media*, pp. 285 – 294, ACM.
2. Bakopoulos, M. and Tsekeridou, S. and Giannaka, E. and Tan, Z.-H. and Prasad, R. (2011) Command & Control: Information Merging, Selective Visualization and Decision Support for Emergency Handling, *Proceedings of the 8th International ISCRAM Conference*, Lisbon, Portugal.
3. Heverin, T. and Zach, L. (2010) Micro-blogging for Crisis Communication: Examination of Twitter Use in Response to a 2009 Violent Crisis in the Seattle-Tacoma, Washington, Area, *Proceedings of the 7th International ISCRAM Conference*, Seattle.
4. Walther, M. and Kaisser, M. (2013) – Geo-spatial Event Detection in the Twitter Stream, *Proceedings of the 35th European Conference on Information Retrieval*, Lecture Notes on Computer Science, Springer.
5. Mendoza, M., Poblete, B. and Castillo, C. (2010) Twitter Under Crisis: Can we trust what we RT?,

Proceedings of the First Workshop on Social Media Analytics, pp. 71 – 79, ACM.

6. Sakaki, T., Okazaki, M. and Matsuo, Y. (2010) Earthquake Shakes Twitter Users: Real-Time Event Detection by Social Sensors, *Proceedings of the 19th International Conference on World Wide Web*, pp. 851- 860, ACM.
7. Terpstra, T. and Stronkman, R. and de Vries, A. and Paradies, GL. (2012) Towards a realtime Twitter analysis during crises for operational crisis management, *Proceedings of the 9th International ISCRAM Conference*, Vancouver, Canada.
8. Vieweg, S., Hughes, A.L., Starbird, K. and Palen, L. (2010) Micro-blogging during two Natural Hazards Events: What Twitter may Contribute to Situational Awareness, *Proceedings of the 28th International Conference on Human Factors in Computing Systems*, pp. 1079 – 1088, ACM.