USE OF SOCIAL MEDIA AND OPEN SOURCE DATA TO ENHANCE SITUATIONAL AWARENESS IN THE AUSTRIAN CRISIS AND DISASTER MANAGEMENT

Dražen Ignjatović, Rainer Simon, Georg Neubauer
AIT - Austrian Institute of Technology GmbH
drazen.ignjatovic@ait.ac.at, rainer.simon@ait.ac.at, georg.neubauer@ait.ac.at

Johannes Pan, Clemens Gutschi
TU Graz - Graz University of Technology
johannes.pan@tugraz.at, clemens.gutschi@tugraz.at

Keywords

Social Media, Situational Awareness, Event Detection, Sentiment Analysis

Abstract

Advanced situational awareness is a central element of today's crisis and disaster management and the basis of improved decision-making. A prerequisite for optimised decision-making is to quickly obtain valid information about the specific crisis in the affected area. At the same time, a high demand for information arises from the media and the population. It is therefore necessary to collect, process and distribute valid information in the shortest possible time. In this context, the KIRAS project iLike focuses on the development of a demonstrator of an intelligent situation information portal that uses open-data information sources and channels for decision support. The demonstrator has two key central functionalities, namely event detection and sentiment analysis. This paper describes the architecture and key functionalities of the iLike demonstrator.

1. Introduction

In state crisis and disaster management (SKKM), the management organisations (e.g. Ministry of the Interior, police, civil protection and fire brigade) have the task of preventing both imminent and actual disasters, managing them as well as mitigating their effects. In this context, it is important to obtain valid information about the situation in the affected areas as quickly as possible in order to be able to make optimal decisions and initiate measures. At the same time, information is generated quickly and on a large scale in the media and among the population, which can be more or less helpful in managing the situation. In disaster situations, it must always be possible to assess the current situation under enormous time pressure and to include all available information in decision-making processes. Social media, instant messaging and other online communication channels that are widely used today have changed the way we share information. Today's online media happens in real time and across a variety of devices. User-generated content blurs the distinction between publisher and consumer and makes the quality of information and its sources difficult to assess. So far, there are hardly any solutions or concepts to efficiently and effectively integrate social media and other open data sources and channels into existing situational information systems.

This paper is based on an ongoing project called iLiKe. The aim of the project is to design an intelligent situation information portal that harnesses open-data information sources and channels for disaster management. The paper is structured as follows: Chapter 2 describes related work on which this paper builds. Chapter 3 presents the system architecture with some selected modules designed and implemented in the project. This is followed by Chapter 4, which shows results from two application examples. The paper concludes with a conclusion on the use cases and an outlook on further work.

2. Related work

Social media includes a broad range of social software platforms where people create, share, and exchange user-generated content. Social software are computer systems and applications that are facilitators or focus for social relationships (Shuler, 1994). User-generated content (UGC) is content that is made available online in a publicly accessible manner or to a group of people. (Vickery and Wunsch-Vincent, 2007).

Similarly, as mobile phone calls increase significantly during a disaster, Internet and social media usage also explodes. For example, Internet usage on the East Coast of the U.S. increased by over 114% in anticipation of Hurricane Sandy in 2012 (Whittaker, 2012). During the 2011 Japan earthquake, Twitter use in reference to the earthquake reached 1200 tweets per second (TPS) (Crowe, 2012); in 2012 during Hurricane Sandy, Twitter traffic in relation to the storm reached about 250 TPS; in 2013 after the Boston Marathon bombings, Twitter traffic in relation to the attack reached 700 TPS (Rovell, 2013) and in 2016 during the Munich attack twitter usage reached a max of 1400 tweets per minute related to the happening (Dyckmans, 2016). These examples show that social media can be an invaluable source of time-sensitive information during such a crisis. However, while disaster response, security, or humanitarian aid organizations have been seeking to leverage this information for more than a decade (cf. KIRAS QuOIMA), the increasing flood of social media platforms and messages remain a significant challenge, both for technical infrastructures, as well as our limited human capacity to process masses of information in a timely manner (Castillo, 2016), also known as the Data Rich, Information Poor (DRIP) syndrome (Goodwin 1996).

This is the reason why tools and methods are needed to support disaster response teams in taking advantage of the new opportunities offered by social media. Although researchers have explored various methods to summarize, visualize or use data for analysis, first responder organizations have not yet been able to take full advantage of research advances, largely due to the gap between academic research and deployed, working systems. The process from collecting social media content, to filtering relevant information, to presenting actionable results to the end user is the goal of iLike. The project aims to provide an indication of a possible solution to bridge this gap by introducing a concept of such a system, but also by building on prior work done in the field. The most relevant research work focuses on semantic integration of disparate information sources into a common operational picture (COP) (Ulicny et al. 2013), identifying of disruptive events from social media (Alsaedi et al. 2015), situational awareness enhancement through social media analytics (Snyder et al. 2019) followed by a report on using of social media for enhanced situational awareness and decision support (Homeland Security, 2014).

3. Implementation of the iLike PoC

3.1. System architecture

This section describes our overall vision of intended workflows for the proposed solution, from raw social media data collection, to methods for content filtering and analysis, to the interfaces, presentation modes and alerting mechanisms provided to the end user. It emphasizes the fact that the maturity of the proof of concept is currently reaching a Technology Readiness Level (TRL) of 4. Planned future developments after the end of the project iLike envisage a higher maturity of the system and are discussed in chapter 5.

Figure 1. represents the high-level architecture of the system, which is divided into three main parts. The first part is the harvesting of social media data, which in this case is done by using software from PublicSonar[1]. A set of customisable filters is applied to the entire search space (a defined set of social

---

[1] https://publicsonar.com/about-publicsonar/

media and open news sources on the Internet) to pre-filter the data and prepare it for analysis. This is also the first step in the planned workflow – *data harvesting*. The second part of the architecture is the analysis of the harvested data by implementing advanced algorithms to extract information from a broader dataset on one hand and to apply sophisticated algorithms to extract the most relevant information from it on the other. A large amount of data is incomprehensible to a human user, so it must be made compact and easy to digest. This represents the second step in the planned workflow - *filtering the relevant data*. The final part of the architecture is the presentation of the analysis results to the end user, which is done using Hexagon[2] software. When designing such a system, it is very important to keep the user in mind and provide a robust and user-friendly interface. This is also the last step in the planned workflow - *providing analytics of the data*.

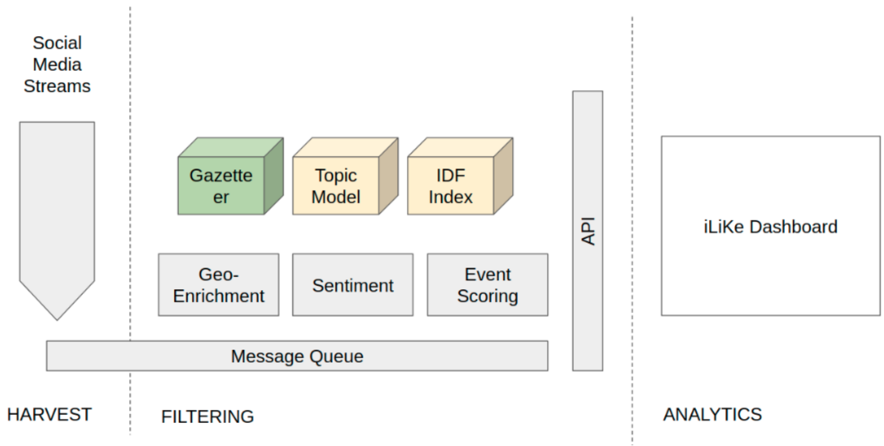

Fig.1. iLiKe high-level architecture.

3.2. Development approach

The iLiKe project follows an agile development approach characterised by incremental co-design and involves stakeholders such as crisis managers, authorities or representatives from LEMAs in several phases. In the initial phase, sample data was collected from the Ski World Cup slalom in Schladming 2020. The collected data was used to develop and train models for data analysis. In two parallel streams of activity, initial requirements for the iLike PoC system were gathered in stakeholder interviews, while at the same time sample data was profiled in an exploratory analysis and a number of selected technical approaches were tested in response to the information gathered from stakeholders. In addition, both developers and stakeholders participated in a physical event in Kulm in Styria in Austria during the 2020 Ski Flying World Cup to gain a better understanding of network traffic and messaging dynamics during a large-scale event. The experience and information gathered during this event was used to define the technical requirements for the PoC, which were validated by the stakeholders in order to develop software that fits to the end user needs in real-life situations. The identified requirements were prioritised from both stakeholder and developer perspectives and implemented and reviewed in monthly sprints. The agile approach allowed for simultaneous development of software demonstrators to test specific filtering and analysis methods and validate them against the collected data. To further elaborate the proposed solution, the project also developed a conceptual user interface design, as a basis for discussion for how such a potential

[2] https://www.hexagonmi.com/

future integrated system could function from the end-user's perspective. To ensure that the user interface met the needs of the end users, the best practices for user interface design were followed, mainly considering the ten usability heuristics (Nielsen, 2005.). In addition, all design ideas were validated by obtaining feedback from the stakeholders involved, using their experience in the field.

4. iLike Functionalities

Two key functionalities were identified in conversations with stakeholders as the most useful to trial:

- **Event detection**, a functionality that, once activated, would perform continuous background monitoring of media traffic, and create alerts when a new topic trend emerges suddenly. Envisioned use cases for such a system are general incident alerting, or targeted monitoring of events that involve large crowds, and where social media use is prevalent, e.g. sports or music events.
- **Sentiment mapping,** a set of spatio-temporal visualizations (e.g. to complement an existing operations dashboard, in the form of printed reports) offering a picture of how the overall "sentiment" of social media discourse evolves over a geographical area, over time. Sentiment analysis has been used previously in the context of disaster management and has also become available as a product in the Open Source Intelligence (OSINT) market.[3] For a comprehensive overview of possible technical approaches, see e.g. Kaur et al. (2017). Such visualizations can be useful in different contexts, for example to assess the impact and public sentiment towards policy decisions or disaster response measures.

4.1. Event Detection

To implement event detection, our demonstrator combines two established natural language processing techniques: topic models and the *term frequency-inverse document frequency* algorithm (TF-IDF). Both techniques are *data-driven* rather than *rule-based.* This means that operators do not need to define specific alert conditions or query terms beforehand. Instead, the techniques rely on a *reference dataset* of social media messages that is harvested a priori, over an incident-free time. The reference dataset then serves as an example to the system for what counts as "normal" message traffic. In our demonstrator, an "event" is characterized by a burst of messages that deviate from the known reference message corpus by some defined criteria, such as vocabulary used, traffic amount, etc.

Topic models are probabilistic models used to uncover underlying semantic structure in a document collection. The idea behind topic models is that each document in a collection can be represented as a vector of different *topics*, where each topic is defined as a distribution over a fixed vocabulary of terms (Blei and Lafferty, 2009). Our system generates a topic model from the reference dataset (using an open-source implementation of the *Latent Dirichlet Allocation* algorithm, cf. Blei et al., 2003) and compares incoming messages against the model. Do new messages match one or more of the topics that were common during the reference period? Or is there no good match with any of them? A sudden spike of "off-topic" messages may indicate a sudden change in theme in the social media discourse and will trigger a notification.

TF-IDF is a well-known algorithm for measuring the "importance" of a word to a document in a document collection (Manning et al., 2008). The basic premise is that the importance is proportional to the relative frequency of the word in the document (term frequency), and inversely proportional to the frequency of the word in the whole collection (inverse document frequency). Thus, words that

---

[3] See e.g. https://www.hensoldt-analytics.com/

are frequent in a document, but rare in the collection overall, receive high TF-IDF scores. Our system uses the reference dataset for comparison, and computes TF-IDF for words found in new messages. As a result, the algorithm automatically foregrounds words that did not occur (or occurred rarely) during the reference period, especially if they become more frequent during the observation period.

4.2. Sentiment Map

We implemented a sentiment mapping demonstrator, comprising two core components: sentiment analysis and automatic geo-coding. **Sentiment analysis** was realized using two of the most popular open-source software packages for this purpose: VADER (Hutto and Gilbert, 2014), and TextBlob.4 Both assign a score to each message that quantifies whether message stance is generally positive (score between 0 and 1), or negative (score between –1 and 0). To determine the score, both VADER and TextBlob rely on term lexica and static rules. TextBlob additionally includes an alternative machine learning implementation, which uses a naïve Bayes classifier trained on movie reviews. Due to the nature of the implementation, different analyzers are required for different languages. In our case, we used TextBlob to process German-language messages, whereas VADER was only used in the experimenting stage, with a small sample of English messages.
To locate messages on a map, we obtained geographic coordinates from two sources. On the one hand, we relied on coordinates that some message sources (e.g. Twitter) embed into message metadata. On the other hand, we implemented a component for **automatic geo-coding** of message texts. The component applies *Named Entity Recognition* (NER), using the spaCy open-source natural language processing framework2 to automatically extract place names mentioned in message texts. Identified place names are then matched against a local index of the GeoNames dataset, a corpus of places, their names, geo-coordinates and other properties. It must be noted that granularity and reliability of both methods has limitations: coordinates embedded in messages often represent a pre-set location chosen by the user; while NER often works well on the city and regional level, but typically fails to yield exact results beyond that (unless, e.g., street names are mentioned in the message, and the system was set up with a street geocoding index).

4.3. Demonstrators and Validation

As mentioned above, the iLike PoC remains at a lower TRL level, which means that generally the individual system components were not integrated into a unified system. In order to validate the proposed approach, small-scale test setups were implemented and tested with the sample data collected earlier. The data collected at the ski race in Schladming provided a particularly good test opportunity for event detection: during the race, a prankster ran across the finish area, triggering timekeeping just before one of the racers crossed the finish line. Naturally, the event triggered corresponding social media reactions, not usually seen during normal skiing races.

We synthesized a reference dataset from messages recorded before the prankster incident and augmented it with messages from some time after the incident, removing all messages that referred to the prankster, in order to build our reference models. The technical test setup for the trial consisted of an ElasticSearch[5] index to store the messages and simulate the message traffic over time; a simple backend application (developed in the Python programming language) to match "incoming" messages against the reference models; and a graphical frontend showing key metrics from the event detector: the rate of off-topic messages vs. total messages; terms most common in the total messages

---

[4] https://textblob.readthedocs.io/en/dev/

[5] https://www.elastic.co/

overall; terms highest scored by the TF-IDF algorithm. In addition, the application also showed the images included in the off-topic messages, if any. Our initial tests revealed that the setup was useful in, first, identifying the fact that an unusual event had occurred, by measuring the ratio of off-topic vs. total messages; second, by exposing the most descriptive terms for the event and, thus, providing immediate insight into the nature of the incident. This was especially useful in the period immediately succeeding the event, where social media traffic was still characterized by a mix of "normal" reporting on the ski event and posts about the incident. Only in a later phase, the incident had received sufficient attention, so that it dominated the discussion as a whole; and additional filtering would no longer yield significant benefit over, for example, simply plotting the most frequent terms in all messages.) Sentiment analysis was evaluated with a second test setup. A dataset was collected separately, over a period of time, by harvesting messages containing terms from a list of keywords related to extreme weather conditions, such as storms, heavy rain, flooding etc. In our demonstrator, sentiment analysis and NER were connected through a message queue which picks up batches of incoming messages for processing; and a web application which visualizes the result as an interactive map and as a timeline in a browser-based user interface. As a showcase, raw result data was also exposed through a RESTful API, and imported into the Hexagon system, where it is made available to the end user for additional capabilities such as documentation, ranking of the analysed data, and so on.

5. Conclusion and Outlook

The work presented in this paper reflects the development of the demonstrator at TRL level 4 achieved in the project iLike. Further development, evaluation and refinement beyond the project is needed to be able to finally assess the extent to which the approach of using and combining social media and open source data can improve situational awareness in crisis and disaster management. There are many solutions to extract data from sources like Twitter (PublicSonar, SYNYO[6]), but only a very limited number of them extracts information that is suitable for decision support. Therefore, we plan to explore more use cases with different stakeholders in different domains such as pandemics and terrorist situations. In the next step, we want to expand the functionalities and maturity level of the demonstrator on the one hand and use it for international use cases on the other hand. For this purpose, we want to work closely with national stakeholders to identify suitable international funding opportunities, such as the EU's Horizon Europe funding programme. Of course, each iLike functionality and further developments have to comply with the legal standards of the General Data Protection Regulation and – with regard to the processing of personal data for the purposes of the prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties – with the national provisions in the Austrian Data Protection Act (DSG), based on the Directive EU 2016/680. As far as personal data is affected, the procession has to be lawful, for example according to reasons of Art 6 General Data Protection Regulation.

5.1. Acknowledgment

The research leading to these results have received funding from the FFG KIRAS program, under grant Agreement no. 873496.

[6] https://www.synyo.com/

6. References

Blei, David & Ng, Andrew & Jordan, Michael. (2003). Latent Dirichlet Allocation. Journal of Machine Learning Research. 3. 993-1022. 10.1162/jmlr.2003.3.4-5.993.

Blei, D.M. & Lafferty, J.D (2009) Topic models. In: Srivastava, A., Sahami, M. (eds.) Text Mining: Classification, Clustering, and Applications. pp. 71-93. Chapman & Hall/CRC, London.

Castillo, C. (2016) Big Crisis Data – Social Media in Disasters and Time-Critical Situations, Cambridge University Press, New York, USA, 2016, ISBN 978-1-107-13576-54.

Crowe, A. (2012) Disasters 2.0: The application of social media systems for modern emergency management. 1st edn. CRC Press.

Dyckmans, P. (2016) 4 Dinge, die wir nach der Amoknacht über Twitter gelernt haben, BR, 08. 2016, Der Münchner Amoklauf in Social Media: 4 Dinge, die wir nach der Amoknacht über Twitter gelernt haben | Netz | Themen | PULS (br.de).

Goodwin, S., 1996. Data rich, information poor (DRIP) syndrome: is there a treatment?. [online] PubMed. Available at: <https://pubmed.ncbi.nlm.nih.gov/10158370/> [Accessed 2 April 2021].

Homeland Security. (2014). Using Advanced Social Media for Advanced Situational Awareness and Decision Support [Ebook]. Retrieved from https://www.dhs.gov/publication/using-social-media-enhanced-situational-awareness-decision-support.

Hutto, C.J. & Gilbert, E.E. (2014). VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. Eighth International Conference on Weblogs and Social Media (ICWSM-14). Ann Arbor, MI, June 2014.

Kaur, H., Mangat, V. and Nidhi (2017). A survey of sentiment analysis techniques. 2017 International Conference on IoT in Social, Mobile, Analytics and Cloud (I-SMAC), 2017, pp. 921-925, doi: 10.1109/I-SMAC.2017.8058315.

L. S. Snyder, M. Karimzadeh, C. Stober and D. S. Ebert, "Situational Awareness Enhanced through Social Media Analytics: A Survey of First Responders," 2019 IEEE International Symposium on Technologies for Homeland Security (HST), Woburn, MA, USA, 2019, pp. 1-8, doi: 10.1109/HST47167.2019.9033003.

Manning, C.D., Raghavan, P. & Schutze, H. (2008). Scoring, term weighting, and the vector space model. In: Introduction to Information Retrieval, Cambridge University Press.

N. Alsaedi, P. Burnap and O. Rana, "Identifying disruptive events from social media to enhance situational awareness," 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM), Paris, France, 2015, pp. 934-941, doi: 10.1145/2808797.2808879.

Ngak Chenda „Social media a news source and tool during Superstorm Sandy.", CBS News, October 2012. http:/www.cbsnews.com/news/social-media-a-news-source-and-ool-during-superstorm-sandy.

Nielsen, J. (2005). Ten usability heuristics [pdf]. Retrieved from https://pdfs.semanticscholar.org/5f03/b251093aee730ab9772db2e1a8a7eb8522cb.pdf.

Rovell, D. (2013) Boston bombing led 2013 tweets, ESPN, Dec 12 2013, https://www.espn.com/nba/story/_/id/10124297/boston-marathon-bombings-most-tweeted-sports-event-2013.

Schuler, D. (1994) Social computing, Communications of ACM, 37(1), 28-29.

Ulicny, B., Moskal, J., & Kokar, M. (2013). Situational Awareness from Social Media. In Semantic Technologies for Intelligence, Defense, and Security. George Mason University.

Vickery, G., Wunsch-Vincent, S. (2007) Participative Web and user-created content: Web 2.0 wikis and social networking. Organization for Economic Cooperation and Development (OECD).

Whittaker, Z. (2012) Internet usage rocketed on the East Coast during Andy: Report. ZDNet Online, Oct 2012.