The JORD System - Linking sky and social multimedia data to Natural and Technological Disasters

Kashif Ahmad
University of Trento
Trento, Italy
kashif.ahmad@unitn.it

Micheal Riegler
Simula Research Laboratory
Oslo, Norway
michael@simula.no

Ans Riaz
University of Trento
Trento, Italy
ans.riaz@unitn.it

Nicola Conci
University of Trento
Trento, Italy
nicola.conci@unitn.it

Duc-Tien Dang-Nguyen
Dublin City University
Dublin, Ireland
duc-tien.dang-nguyen@dcu.ie

Pål Halvorsen
Simula Research Laboratory
Oslo, Norway
paalh@simula.no

ABSTRACT

Being able to automatically link social media information and data to remote-sensed data holds large possibilities for society and research. In this paper, we present a system called JORD that is able to autonomously collect social media data about technological and environmental disasters, and link it automatically to remote-sensed data. In addition, we demonstrate that queries in local languages that are relevant to the exact position of natural disasters retrieve more accurate information about a disaster event. To show the capabilities of the system, we present some examples of disaster events detected by the system. To evaluate the quality of the provided information and usefulness of JORD from the potential users point of view we include a crowdsourced user study.

1. INTRODUCTION

During the last few decades, satellite data has been widely used to explore the surface of the Earth and its environment. Being able to cover a large spatio-temporal area, remote sensed data has been proved very effective and inherently useful in different applications like disaster management [9]. Recently, NASA released a dataset from the longest running satellite program called Landsat[1]. Landsat has been heavily used by researchers in the last decades where useful techniques have been developed, e.g., automatic detection of forest fires, sandstorms or floods [6, 8]. The latest Landsat dataset contains images shot by the Landsat 8 satellite, and it is available via Amazon Webservices[2]. Such a dataset provides many opportunities for society, and researchers can benefit from remote sensed data to develop new applications integrating satellite images in detection, retrieval and monitoring services.

Nevertheless, remote-sensed data also brings some challenges. For example, satellite images have a low temporal frequency, and more importantly, they only give a birds'-eye view of an actual event. For instance, nowadays, floods can be detected in satellite images (St-George seen from NASA MODIS satellite[3]) as shown in Figure 1, but it becomes difficult to understand the impact of the event on the people living in the area, both from practical and societal view point.

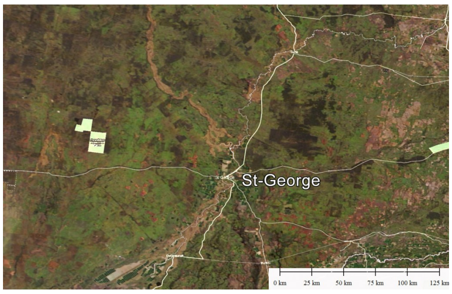

Figure 1: Based on a satellite image (flood in St George), it is almost impossible to give a clear statement about the impact on the environment and society (image: NASA MODIS Satellite)

On the other hand, over the last few years social media has emerged as an important source of information and rapid communication in emergency situations [14]. Particularly, Twitter has proved to be very effective in dissemination of news about natural disasters [1]. Moreover, as demonstrated in [12], there are many situations in which news agencies could not provide information at all or in time, about emergency situations, simply due to the unavailability of reporters in the area. In such cases, social media plays an important role [3].

A rather recent trend is to combine content from social media with remote sensed data, e.g., in the MediaEval benchmark initiative with the "placing" task[4] where researchers try to predict geo-coordinates for images and videos from Flickr. Additionally, a task to build a system that links social multimedia to events, which can be detected in satellite images, has been introduced as a challenge at ACM MM 2016[5]. This clearly shows that the multimedia research community is interested in this new direction.

To this aim, we propose a system called *JORD*, which is to the

[1] http://landsat.usgs.gov/
[2] https://aws.amazon.com/public-data-sets/landsat/

ICMR '17, June 06-09 2017, Bucharest, Romania
ACM xxx-x-xxxx-xxxx-x/xx/xx.
http://dx.doi.org/xxxxxxxxxxxxxxxxx

[3] http://www.aerometrex.com.au/blog/?p=289
[4] http://multimediaeval.org/mediaeval2016/placing/
[5] http://www.acmmm.org/2016/wp-content/uploads/2016/03/ACMMM16_GC_Sky_and_the_Social_Eye_latest.pdf

best of our knowledge the first one, that is able to combine satellite images with information from social media (e.g., Twitter, Flickr and Youtube) and news items collected in real-time (as fast as possible after the event occurs). It also provides query refinement by automatically generating queries in all local languages that are relevant to the position of a disaster. With a system combining multimedia mining [10], retrieval [2] and linking [5] methods, we are able to tell a much clearer and more useful story to the users. JORD is able to link social multimedia data to events that are detected in satellite images. For instance, JORD is able to link an automatically detected sandstorm to social media reports found in Flickr, Youtube, Twitter and web pages. The linked information then can be used by the users of the system to get a clearer picture of the event and its possible impacts on the environment and the society. Moreover, the proposed system can also be helpful to follow up and monitor the process of rebuilding of infrastructure and other rehabilitation activities.

The main advantages of JORD are that it (i) collects data about events autonomously, continuously, and automatically in real-time from a disaster database (ii) automatically filters irrelevant information (iii) is able to link this information automatically with satellite and social media data using information retrieval and (iv) allows query refinement for the retrieval of information using automatic language detection and translation.

For the demonstration purposes, we will give a live demo of the system where it will be demonstrated how our system is able to collect and explore different multimedia data along with satellite images for natural and technological disasters. We will also show how JORD can facilitate users to find information about disasters from different sources in one go.

The rest of the paper is organized as follow: Firstly, we provide a detailed description of the methodology adopted to develop the proposed system. After that, we present the system architecture with more focus on the technologies used in implementation. This is followed by a detailed evaluation of the system using crowdsourcing to get real user feedback. At the end, we give a conclusion and present some ideas about future work.

2. METHODOLOGY

Figure 2 provides a block diagram of JORD. It is composed of three main parts (highlighted with different colors in Figure 2): (i) multimedia data retrieval from social media, (ii) filtering and (iii) linking social media data with remote-sensed data and events. We provide a detailed description of all parts of the system which includes the natural disasters we target, query refinement through translation, multimedia data retrieval from social media, filtering and linking multimedia data from social media with remote sensed data.

2.1 Target Events

As one can see in Figure 2, JORD first extracts a list of natural and technological disasters from the EM-DAT database [7] in real-time. This means that as soon as a new event occurs in the database, JORD starts collecting and linking information about it. EM-DAT is an international disaster database (supported by the World Health Organisation WHO) that provides information of natural and technological disasters that have occurred all over the world as fast as possible. We test our system with 82 recent natural and technological disasters. However, it is important to point out that JORD is able to collect, retrieve and link an unlimited number of events, and can operate in both a live mode as an quasi autonomous system and an on demand mode controlled by a user. Moreover, the events are not limited to a certain continent or region.

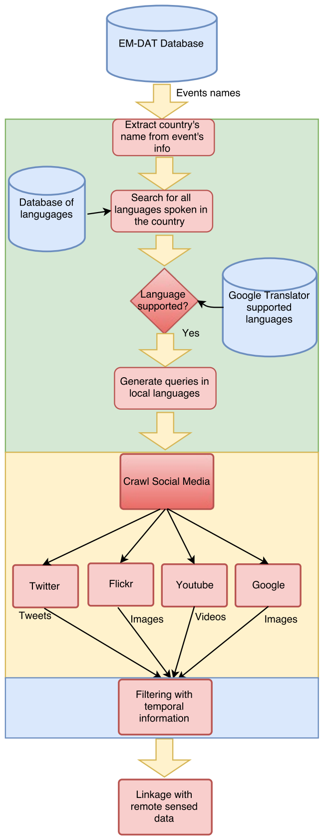

Figure 2: Block diagram of the proposed JORD system.

2.2 Query Refinement and Translation

The EM-DAT database provides the time period and location information for each disaster. Our system utilizes this information for the retrieval of related social multimedia data, such as videos, images and text. Furthermore, JORD supports query refinement by collecting additional tags for the underlying events by determining and translating them to the local languages of the area where the disaster occurred. By looking at news or social media posts, one can easily observe that during natural and technological disasters the local community usually initiates the process of spreading news about an event. This normally happens by commenting, posting and sharing information using different channels and media types right after an event occurs. Furthermore, people living in the proximity of the event usually tend to post and share information in their local languages. Based on these observations, our system automatically generates queries in all local languages that are relevant to the position of a disaster.

The translated queries are used in the next part of JORD, which is responsible for collecting multimedia data. During data analysis, we observed that the combination of translated and the original queries results in a larger amount of retrieved data per query (for some events, a search based on only English queries results in very little or no result). Moreover, user feedback revealed that the information retrieved with translated queries are more relevant to underlying events.

2.3 Social Media Platforms

In recent years, social networks, such as Twitter and Flickr, have emerged as important sources of information that report events in real-time, and provide a much broader story [13, 4]. With respect to Twitter, for example, the majority of the users use Twitter for breaking news, and thereby become a part of the process by commenting, posting and sharing information at the moment when an event occurs [1]. Similarly, Flickr and YouTube facilitate users to share audio-visual contents about an event whenever it happens. Therefore, it is important to include social media information in our system, and to get an overview of social media content for an event, we crawl four different social media platforms, i.e., Twitter, YouTube, Flickr, and also Google image search. It is important to mention that queries in local languages, generated by our system during query refinement phase, are supported by three platforms including Twitter, YouTube and Google. Flickr supports queries in English only.

2.4 Filtering and Linkage with satellite Data

Another challenge that we tackle is how to ensure that the retrieved information is relevant to the event. To solve this problem, we filtered the downloaded data using time information provided by EM-DAT database for each event. The final step of our system is the use of geo-location information to link the retrieved social media data and the remote-sensed images (before and after a disaster). To this aim, we rely on Google Earth and the geo-location information.

3. SYSTEM ARCHITECTURE

Figure 3 shows the block diagram of the architecture of our proposed system, which is composed of three main components: (i) Front View with a user friendly interface allowing users of the proposed system to interact with the tool according to their need (2) Controller, which is basically composed of the crawler with query refinement capabilities; (3) Database Model, which represents the structure of the database to store retrieved data. A reliable database, and a user friendly interface is always very crucial in such applications. To this aim, in order to host huge amounts of data, we carefully design the structure of the database for the proposed tool. Moreover, we tried our best to keep the database as simple as possible by avoiding the complexity of foreign keys and associations. Every kind of information required in the platform is stored in a proper way by allocating a separate table for it. All in all, we have 4 tables in our database, one containing information of events for which our crawler use to search data in social media platforms. Similarly, for each type of data, retrieved from social media platforms, we have a separate table, which facilitates information retrieval and report generation. Our database also keeps meta data along with the images, videos and tweets downloaded from social media. Similarly, the front view is designed in a way to fully facilitate users with less efforts on the user's side. A number of functionalities have been provided, to facilitate users, in the tool. For instance, users can search for a particular disaster as well as for a particular type of data. As far as the controller unit is concerned, it is composed of the crawler, which searches for images, videos and tweets in social media, with query refinement capabilities, and is implemented in PHP. To this aim, respective APIs have been integrated in the crawler to download multimedia data along with meta data, and store it in the database. The crawler can be configured in two different ways (automatically in a certain interval or manually) regarding how often the social media should be crawled). For query refinement and translation, the Google Translator Api[6] is used in combination with a database of spoken local languages per country in our pipeline. The list of local languages along with

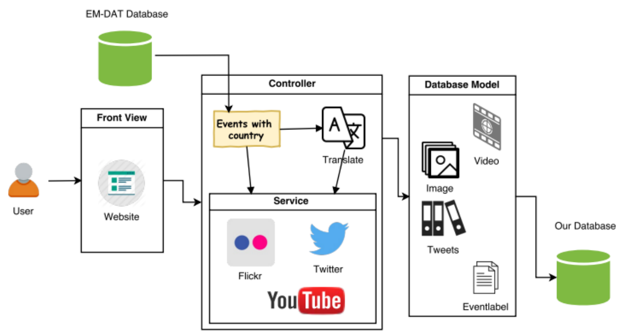

Figure 3: System architecture of JORD.

the country names are retrieved from InfoPlease[7]. Infoplease is a free encyclopedia almanac, atlas, dictionary, and thesaurus. Thus, JORD automatically extracts country information, where the disaster has been occurred, from the EM-DAT database and checks for the local languages spoken in that country, and subsequently generates new queries in local languages via Google Translator.

4. SYSTEM EVALUATION

Figure 4 shows a sample output for a flood query in Kenya. As can be seen in Figure 4, JORD combines the information from different sources, and presents a more detailed story about the disaster. To evaluate the system, in terms of if the retrieved multimedia contents about the disaster events are correct and useful for the users, we conducted a crowdsourcing study on Microworkers. We asked workers to give their opinion about the retrieved multimedia contents related to underlying events. We paid each worker 1.50 USD and tried to be as fair as possible regarding the discarding of workers. As shown in [11], controlling and discarding workers too much can lead to an undesired outcome of the study, which we tried to avoid by accepting almost every worker if they did the task in a reasonable way (which means did not try to cheat).

For the evaluation, we asked the crowd-workers five different questions: The first question, (i) "Do you think the system provided information was useful?", aims to get feedback from workers about the usefulness of the collected multimedia data using a scale from one to five from not useful to very useful. (ii) "From these three possible events, which one do you think has been the one presented to you?" (a list of three possible event names, which also contains similar types of events but in different location and time period, is presented to the user, two are wrong, one is the event retrieved by the system). This question was used to evaluate if JORD can help the user to understand the retrieved event and if the presented information is useful. The third question (iii) asked the workers to state how useful they find the three different type of multimedia content (images, tweets and videos) from a scale from one to five from not useful to very useful. The fourth question (iv) was a simple yes or no question where we asked the workers if they would use such a system or not. The final question (v) was an open question where the crowdworkers had to reason their yes or no from the previous question. This question we used to filter out workers that did the task not correctly. We checked each answer manually for each worker. If the answer made sense and showed that the worker answered it carefully we accepted it. If not we did not include it in the final evaluation. In total, we had 385 workers working on the task. We discarded 36, because they did not answer the open question in a way that it appeared that they conduced the task properly.

[6] https://cloud.google.com/translate/docs/

[7] http://www.infoplease.com/ipa/A0855611.html

This resulted in answers from at least five distinct users per event[8]

The first question (i) where the workers had to state if they find the system useful or not, had an average of 4.47 for all workers. This is a clear indication that workers find the system and the provided information useful.

For the second question (ii) where the workers had to chose the correct event out of three possible ones, only 19 workers out of 349 failed to correctly recognize the event presented to them. A closer investigation showed that all of them had just the country wrong but gave a correct answer about the disaster. This shows two important things. Firstly, that the retrieved information of JORD is accurate and can help users to get more information about events and secondly, that connecting it to satellite images is important to improve understanding of the event in terms of location.

For the third question (iii) where the workers had to report how useful they find different types of multimedia content we got an average of 4.23 for images, 4.08 for tweets and 4.44 for videos. Based on this, we can see a tendency that users find videos most useful and tweets least. We think that this might be due to the fact that a video usually contains more information than a text or image and that it helps people more to understand and experience the current situation.

The last two questions (iv) and (v) are evaluated together. For the first question (iv), 336 from 349 workers (around 96%) stated that they find the system useful. This is a promising indicator that such a system would be useful and used by users.

5. CONCLUSION

In this demo paper, we presented a demo of JORD, which is able to autonomously and automatically retrieve social media data from various social platforms and links it to remote-sensed data. We demonstrated that queries in local languages that are relevant to the position of natural disasters retrieve more accurate information about a disaster event. We also showed that JORD automatically finds interesting events and used these events to query social media. The evaluation of the JORD system was carried out through crowdsourcing study where workers were asked to evaluate the usefulness of the system and to identify an event presented to them with collected images, tweets and videos. The evaluation indicates that JORD works very accurate without human input, and it can be used to collect and merge a large number of event based data for technological and environmental disasters from different sources. In future, we plan to include more different social media platforms and also to improve the representation of the retrieved and linked information to the users.

6. REFERENCES

[1] F. Atefeh and W. Khreich. A survey of techniques for event detection in twitter. *Computational Intelligence*, 31(1):132–164, 2015.

[2] S.-F. Chang. New frontiers of large scale multimedia information retrieval. In *Proceedings of the ACM International Conference on Multimedia Retrieval (ICMR)*, pages 5–5, 2016.

[3] M. Cheong and V. Lee. Twittering for earth: A study on the impact of microblogging activism on earth hour 2009 in australia. In *Proceedings of the Asian Conference on Intelligent Information and Database Systems*, pages 114–123. Springer, 2010.

[4] X. Dong, D. Mavroeidis, F. Calabrese, and P. Frossard. Multiscale event detection in social media. *Data Mining and Knowledge Discovery*, 29(5):1374–1405, 2015.

[8]The created dataset including the crowdsourcing study, the events and the results from the crowdsourcing study can be requested by mailing one of the authors. We cannot make it public available due to copyright restrictions.

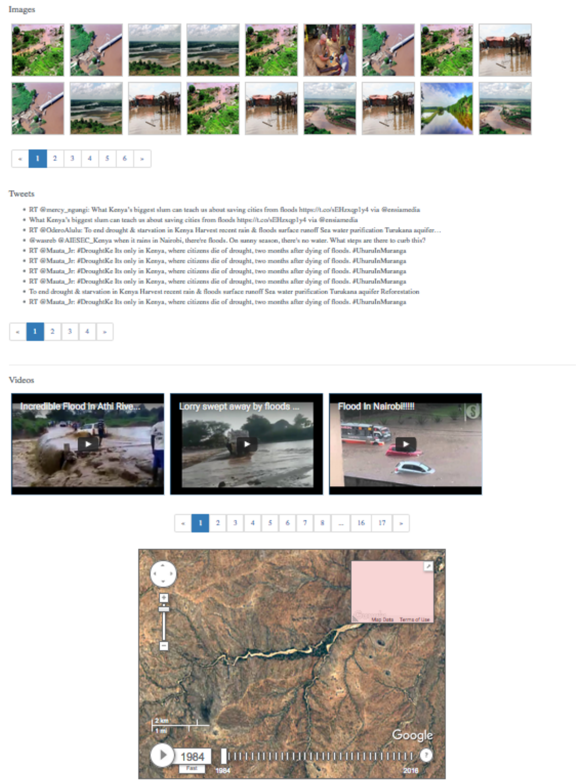

Figure 4: A sample output of JORD in terms of retrieved images (at the top), Tweets, Videos and the satellite data.

[5] B. Erol, J. J. Hull, and D.-S. Lee. Linking multimedia presentations with their symbolic source documents: algorithm and applications. In *Proceedings of the eleventh ACM International Conference on Multimedia (MM)*, pages 498–507. ACM, 2003.

[6] A. Fisher, N. Flood, and T. Danaher. Comparing landsat water index methods for automated water classification in eastern australia. *Remote Sensing of Environment*, 175:167–182, 2016.

[7] D. Guha-Sapir, R. Below, and P. Hoyois. Em-dat: International disaster database. *Catholic University of Louvain: Brussels, Belgium*, 2015.

[8] J. Kansas, J. Vargas, H. G. Skatter, et al. Using landsat imagery to backcast fire and post-fire residuals in the boreal shield of saskatchewan: implications for woodland caribou management. *International Journal of Wildland Fire*, 25(5):597–607, 2016.

[9] S. Lewis. Remote sensing for natural disasters: Facts and figures. *Science and Development Network*, 2009.

[10] T. Manjunath, R. S. Hegadi, and G. Ravikumar. A survey on multimedia data mining and its relevance today. *IJCSNS*, 10(11):165–170, 2010.

[11] M. Riegler, V. R. Gaddam, M. Larson, R. Eg, P. Halvorsen, and C. Griwodz. Crowdsourcing as self-fulfilling prophecy: Influence of discarding workers in subjective assessment tasks. In *14th Workshop CBMI*, pages 1–6, June 2016.

[12] B. Stelter and N. Cohen. Citizen journalists provided glimpses of mumbai attacks. *The New York Times*, 30, 2008.

[13] C. Tzelepis, Z. Ma, V. Mezaris, B. Ionescu, et al. Event-based media processing and analysis: A survey of the literature. *Image and Vision Computing*, 2016.

[14] J. Yin, A. Lampert, M. Cameron, B. Robinson, and R. Power. Using social media to enhance emergency situation awareness. *IEEE Intelligent Systems*, 27(6):52–59, 2012.