Event Representation and Visualization from Social Media

Xueliang Liu[1,⋆] and Benoit Huet[2]

[1] Hefei University of Technology, Hefei, China
[2] EURECOM, Sophia Antipolis, France
liuxueliang@hfut.edu.cn, huet@eurecom.fr

Abstract. The user generated content, available in massive amounts in social media, is receiving increased attention due to its many potential applications. One of such applications is the representation of events with multimedia data. This paper addresses the problem of retrieving and summarizing events on a given topic, and propose a novel and original framework for leveraging social media data to extract and illustrate social events automatically on any given query. The problem is tackled in three steps. First, the input query is parsed semantically to identify the topic, location, and time information related to the event of interest. Then, we use the parsed information to mine the latest and hottest related events from social news web services. In the end, for each event, we retrieve both relevant tweets on Twitter and compelling images from Google image search. The resulting documents are shown within a vivid interface featuring both event description, tag cloud and photo collage.

1 Introduction

The amount of social media data posted daily by online users has grown rapidly in recent years and continues to do so. How to leverage the explosion of this vast amount of data to benefit web users at large is, however, still an open and challenging problem with numerous potential application.

An event is one of the most important cues for people to recall past moments of one's life. The reminder value of an event makes it extremely helpful in organizing data. Thus, events-based media analysis has recently drawn much attention within the multimedia research community.

In this paper, we address the problem of retrieving and summarizing events from social media data, and propose a novel and original framework to extract and illustrate social events automatically on a given query. To solve the problem, we use natural language processing algorithm to parse the input query semantically and extract the most popular events from social news web service. Then, the obtained events are depicted through a multi-modal faceted representation composed of a textual description, a tag cloud and a photo collage, providing

⋆ The work was performed while the first author was working towards his PhD at EURECOM.

B. Huet et al. (Eds.): PCM 2013, LNCS 8294, pp. 740–749, 2013.

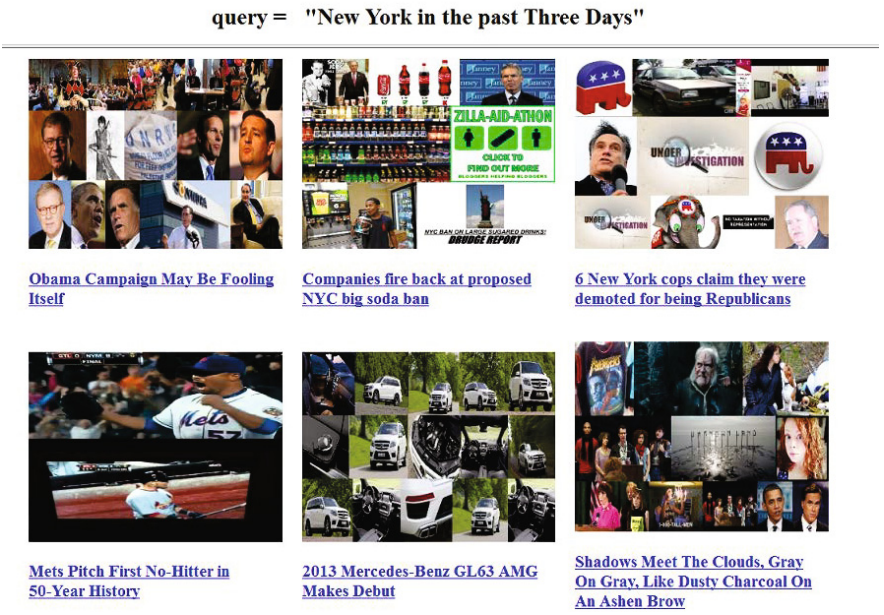

Fig. 1. The snapshot of a query example results

the viewer with a rich informative panel about events. The novelty of our proposed framework is twofold. Firstly, as opposed with other methods [15] that mine events from social media data directly, we extract events from social news website Digg.com. This has considerable advantages: no additional time cost, computation and storage are needed and the results can be achieved online. Furthermore, the built-in collective intelligence of social news web services can be used to rank the hottest events. Secondly, we collect relevant media documents from different media source and use tag cloud and photo collage to provide vivid interface illustrating as visually as possible the events. This enables content from different view to be shown at the same time. The result is generated automatically on any given query. A subset of the events for an example query can be found in Figure 1.

This paper is organized as follows: we review the related work in Section 2, and describe our proposed method in Section 3. Experimental results are presented in Section 4. Finally, we summarize the paper and propose future work in Section 5.

2 Related Work

Nowadays, Twitter has become one of the most important tools for people to share their interest, their personal issues, their views and their experiences as well as comment on other's. Some research has been done to find events directly from Twitter post [15,14]. In [15], the authors studied how to employ a wavelet-based techniques to detect events from Twitter stream. A similar method can

be found in [4] to detect events from Flickr time series data. In [14], the authors investigate how to filter the tweets to detect seismic activity as it happens. In [2], a system is proposed to detect emerging topics from social streams and illustrate the topics with the corresponding information in multiple modalities. Quack *et al.* [13] presented methods to mine events and object from community photo collections by clustering approaches. In [4], a wavelet based approach is proposed to detect events from social media data. Pan and Mitra [12] developed a system to combine the popular LDA model with temporal segmentation and spatial clustering for automatically identifying events from a large document collection. In [7], the authors focused on building a Naive Bayes event models which classify photos as either relevant or irrelevant to given events.

Illustrating events with media addresses the problem of how to leverage vivid multi-modal content to share experience. In [10], the authors studied users' uploading behaviors on Flickr and matched concert events with photos based on different modalities; such as text/tags, time, and geo-location. It results in an enriched photo set which better illustrates events. A similar framework involving more modalities is proposed in [3] for enriching event descriptions. In [6], to improve the users' attention when reading news articles, a system was proposed to help people reading news by illustrating the news story. In [8], an unsupervised approach was presented to describe stories with automatically collected pictures. In this framework, semantic keywords are extracted from the story, and used to search an annotated image database. Then a novel image ranking scheme automatically choose the most important images. In [17], a Text-to-Picture system was developed that synthesizes a picture from natural language text without limitation. The system firstly identified "picturable" text units by natural language processing, then searched for the most likely image parts conditioned on the text, and finally optimized picture layout conditioned on both the text and image parts.

A web service with similar illustration functionalities can be found in EventBurn[1]. It creates a summary of a given hot event from popular services like Twitter, Facebook, and Flickr, but fails to extract events automatically from social media streams.

3 Our Proposal

Our framework, shown in Figure 2, extracts and illustrates public events by leveraging the social media data directly. Since events can be defined as something happening at a given location and time, we start by parsing the query to identify its topic, location and time information with natural language processing algorithm. Rather than detecting events from Twitter data directly, events can be obtained by crawling and scraping social news web service. This saves time, computation and storage compared to alternative detection processes. To provide a vivid illustration for each event, we retrieve the relevant tweets from Twitter, and show the collected data with tag cloud. In addition, we also retrieve

[1] `http://www.eventburn.com/`

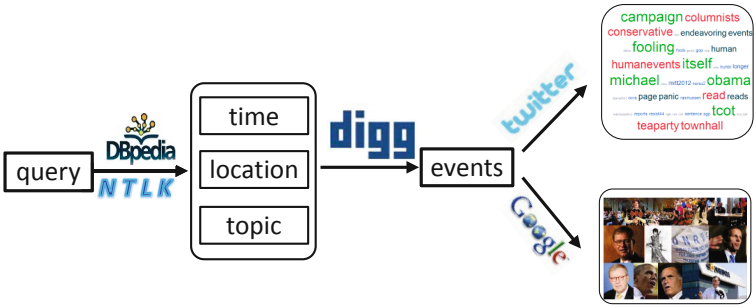

Fig. 2. Overview of the proposed framework

photos with Google image search engine, and summarize the results with photo collage/montage.

3.1 Semantic Query Parsing

The three basic properties of events are location, time, and topic, as stated in [10]. To identify the meaning of a given query input, we would like to extract the information in the three dimensions. Here, we assume that the query input is a noun phrase headed or tailed with complements, such as "New York in the past three days". We extract the structured data from this noun phrase, where there is a predictable organization of entities and relationships. The question of extracting structured data from text has been well studied in Natural Language Processing (*NLP*). This process, composed of 3 steps is performed using *NLTK*[2], a well-known *NLP* package. First, the input text is segmented into words using a tokenizer. Then, each word is tagged with part-of-speech tag (*POS*), which provides the lexical categories for words. With the *POS* tag, we use the *Regexp-Parser chunker* in *NLTK* to create the chunker tree and identify each sub noun phrases in the input string. The process is depicted in Figure 3

Then, to determine the semantic meaning of each noun phrase, different techniques are employed to extract the location, time, and topic information from the parsed noun phrases. The location is obtained through a query of the DBpedia[3] database, which provides structured information extracted from Wikipedia. If geographical metadata, such as "geo:point", "geo:area", are found for a noun phrase, the corresponding location is kept as reference for the events to collect. For determining the temporal information, we develop a script which parses and converts the human readable string, such as "tomorrow","last week", "Monday" to a time structure. We use the DBpedia API and our script to parse the sub noun phrases in order to obtain the location and time information addressed by the query. Since it is hard to model the topic in sophisticated way, we assign the

[2] `http://nltk.org/`
[3] `http://dbpedia.org`

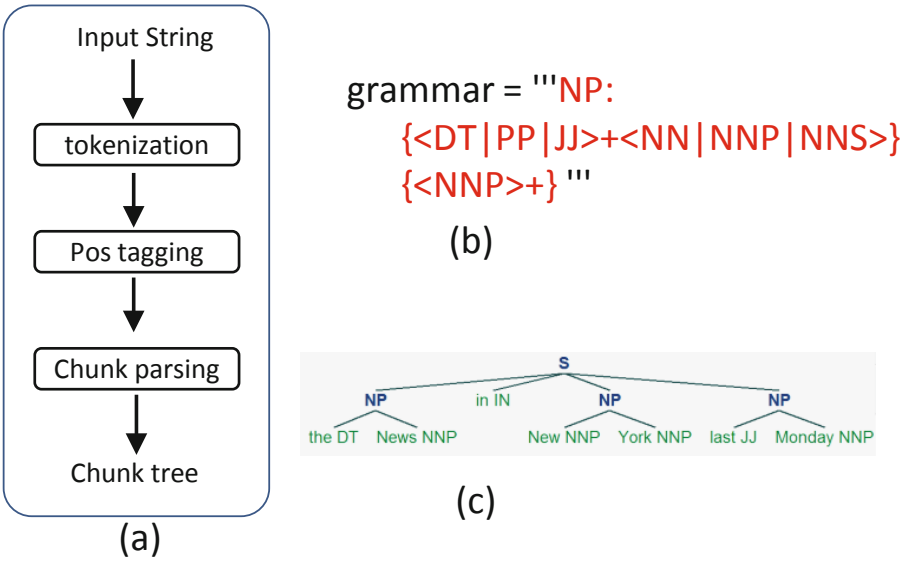

Fig. 3. Semantic parsing using NLP. (a) the flowchart; (b) the grammar used for chunker parsing; (c) an example for input query "the News in New York last Monday"

nouns as the topic keywords of the event to search for, if neither time or location can be determined from it.

3.2 Events Extraction

Recently, there has been some research focusing on detecting events from social media data, such as detecting events from Twitter stream [15], or Flickr [4,9]. Indeed, useful information can be mined from community-contributed data. However, in these methods, huge amount of data have to be downloaded from the web service. Therefore, it takes lots of resources to store the data and lot of time to process it.

Currently, the world's happenings are collected as News by communication services and broadcasted publicly to the mass audience through various channels. Some web services, such as Google News[4] or Digg[5] have been developed to organize the news data in a structured data. In this paper, rather than detecting events from social streams, we aim at querying the events from the social news web service Digg. Digg features user-posted news that are ranked based on the user popularity and comments. With the ranking, popular news can easily be identifyed, leveraging from users' collective intelligence. We take the time,

[4] `http://news.google.com`
[5] `http://digg.com/`

location and topic keywords as the query parameters to retrieve popular events, ranked according to their popularity. In the returned results, many useful metadata can be found for each event, such as "title","submit_date" to describe the property of events. We also extract the "link" metadata that refers the original content of the news, so that the original content could be retrieved and be used in the text description in the final result.

Although the events is extracted from Digg in this paper, the whole framework is flexible and can be extended to handle more events directories, such as Google News or any such public API easily.

3.3 Media Illustration

Years of multimedia research have shown that it is easier and more accurate for the computer to identify specific pattern compared with abstract concept. To find media illustrating events, a query is specifically tailored for each event. To illustrate events using text and images, we retrieve the original news content, related tweets from Twitter and images from Google image search, and show them in the proper format.

The event is fully described in the original news web page. For each event, we extract the "link" metadata in Digg and crawl the original web page. To extract the the main body of the web page, the method proposed in [5] is employed. It recognizes the main content based on character number under the assumption that characters accumulate more in the main content than in other part of web page.

For each event, we would like to collect other textual description and comments from different users. The textual data is retrieved from Twitter, the famous online microblog service that enables the users to send instance posts known as "tweets". As there are hundreds of thousands of tweets posted every minute in a city, tweets are also studied as the data source to mine events/trends [15]. In this paper, we retrieve related tweets to illustrate events. The event title is often the most useful information for describing/retrieving an event. We perform full text search queries on Twitter based on the event title extracted from Digg. The tweets retrieved are also filtered using a time interval from the time of posting the event in Digg to the end of querying period.

To provide a nice and meaningful visualization, we use a tag cloud to organize the textual data. Tag cloud is a form of histogram which can represent the frequency of over a hundred items. In tag cloud, the importance of each tag is shown with different format. This format is useful for quickly perceiving the most prominent terms. For each event, we segment each tweets into tags, count the frequency of each tag and generate the tag cloud with tags in different font size and color.

Visual summerization is fundamental tool to provide vivid description on image/video collections [16]. Besides tag cloud, we also employ such techniques

to illustrate event. Photo collage [11] is a visual clustering technique that can depict the event in different points of view. To generate the photo collage, we retrieve photos from the Google image search engine with the event title as parameter and filter out those photos for which the cosine distance between its textual metadata and the event title is below a given threshold. We decided, after experimentation, not use Flickr as photo source because most of photos are taken during tours/vacations and the dominant concepts are scene, landmark, and building. With the selected photos, the methods proposed in [11] is used to create the photo collage for each events.

Last but not least, to assist the users quickly grasp preponderant information, we would like to show the events according their importance. In details, we measure the importance of each event by the entropy of its tag cloud [1]. After ranking, the events with more information will be shown first and events with less information shown later.

Table 1. Events found for query "New York in the last 3 days", performed on June 2rd, 2012

Date	Event Title
01/06/2012	Companies fire back at proposed NYC big soda ban
01/06/2012	Motorcyclist clocked at 193 mph? in the rain
01/06/2012	New York Bill Proposes Mandatory Wearable ID Tags @ skewnews.com
31/05/2012	Carmelo Anthony Becomes VP Of PowerCoco Energy Drinks
01/06/2012	Cheap Dresses Online Australia - On Sale From $10
02/06/2012	A Powerful Interview with Former Guant?namo Prisoner Lakhdar Boumediene
01/06/2012	NYPD vs. CPD: How Police Deal With Press and Protesters - New York - Slideshows
31/05/2012	Obama Campaign May Be Fooling Itself
03/06/2012	6 New York cops claim they were demoted for being Republicans
02/06/2012	Mets Pitch First No-Hitter in 50-Year History
31/05/2012	2013 Mercedes-Benz GL63 AMG Makes Debut
01/06/2012	Shadows Meet The Clouds, Gray On Gray, Like Dusty Charcoal On An Ashen Brow
31/05/2012	Scientific Proof That Men Have the Dirtiest Desks
03/06/2012	Space shuttle hardware is on the move in Houston and NYC
31/05/2012	Obama Presides Over Secret 'Kill List'
01/06/2012	Facebook Forced To Let You Vote On Privacy Changes
02/06/2012	Paul Krugman Pompously Insults Ron Paul And His Supporters Mediaite
02/06/2012	NBA Arrested For Marijuana Possession
31/05/2012	Bar Refaeli gets close to Olympic star

4 Experimental Results

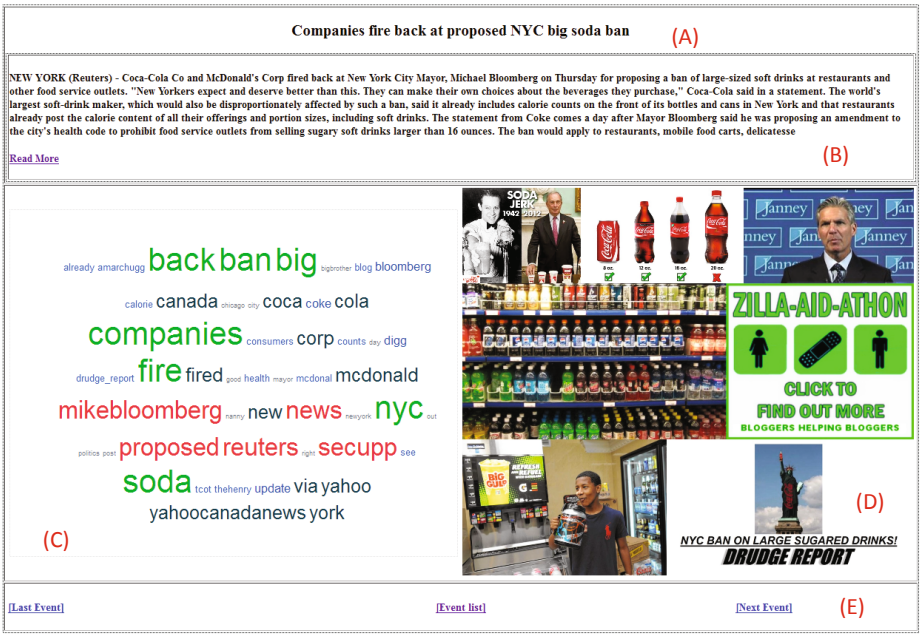

Fig. 4. An example of event visualization, to illustrate the event "Companies fire back at proposed NYC big soda ban", (A) Event title; (B) Event abstract, a link to the original news; (C) Tag cloud; (D) Photo collage; (E) Navigation

Our system can be used to extract and illustrate events from a query input without limitation. For a given query, we first parse it to find out the topics, location, and time information. It should be noticed that for some queries, one or more semantic dimension is missing, and a default parameter will be set. For example, in query " the Olympics in London", no time information is provided. The default value are "in the last 7 days" for time, to ensure the timeliness of events, and "worldwide" for location, and "" for topic, to reduce the limitation of event querying. The motivation for using "in the last 7 days" as default value for time, is that recent news are more likely to be the topic of interest when not specified. We then use the parsed topic, time, and location to query events from Digg. The top 20 events are kept as hot candidates according their ranking promoted by the users.

For example, from the query "New York in the last 3 days", the following events in Table 1 are found.

For all of the queried events, we provide an indexing interface as shown in Figure 1, that provides the thumbnails for all of the events with the combination of photo and event title. All of the events are ranked by their importance

as described in Section 3.3. When navigating to a specific event, the event illustration as shown in Figure 4 is provided, giving the user information related on that event solely. Although currently limited to a simple layout, which can easily be improved by a graphic designer, it is composed of five individual parts to help people understand the event well.

The event title is shown in part A, which is the highest level description for the events.

To provide the users with detailed information about the event, we also parse the original news web page, and mine the main textual content part, as shown in part B.

We parse the title and time metadata from the obtained events, use the parameters to query from Twitter the comments from different users, then use the tag clouds to show the results. The tag cloud is presented in part C. We can clearly see that the larger size content such as "back", "ban", "companies", "fire", matches the event topic "companies fire back at the big soda ban" very well.

Besides text visualization, we also retrieve photos with Google image search, and filter the ones that can not be matched based on their textual metadata. All of the matched photos are collaged in the same layout. The photo collage is shown in part D. From the figure, it could be found that most of photos are relevant to the event; The ban of large-sized soft drinks at food service outlets.

To assist users navigating between events, the link to the previous/next event and the index page is located at the bottom of the webpage, as shown in part E.

To conclude, while tag cloud and photo collage in part C and D provide attractive interface with abstract content, the textual content in part B gives concrete description that would assist the users to understand the event well.

5 Conclusion

We proposed an original framework leveraging on social media data (News, Media Sharing Platform and Microblog) to extract and to depict public events. The process is done without any human assistance on a given textual query. At first, natural language processing is employed to parse and make sense of the input query. Then, a social news web service is employed to extract relevant events (with respect to the query) and the metadata from selected events is used to query textual and visual content from different online sources (Twitter and Google). We present the results with an attractive visual format combining tag clouds and photo collage. We demonstrated how our novel approach addresses the event retrieval and illustrating problem and showed its effectiveness on a real world example.

In future work, we will study how to model the relation between multi-modal content by exploiting multiple features, such as text, tag, visual content to help remove noisy data and therefore improve the quality of the presentation of events.

Acknowledgments. This work was supported by the European Commission under contracts FP7-287911 LinkedTV and FP7-318101 MediaMixer.

References

1. Aouiche, K., Lemire, D., Godin, R.: Web 2.0 OLAP: From data cubes to tag clouds. In: Cordeiro, J., Hammoudi, S., Filipe, J. (eds.) WEBIST 2008. LNBIP, vol. 18, pp. 51–64. Springer, Heidelberg (2009)
2. Bao, B.-K., Min, W., Sang, J., Xu, C.: Multimedia news digger on emerging topics from social streams. In: Proceedings of the 20th ACM International Conference on Multimedia, MM 2012, pp. 1357–1358 (2012)
3. Becker, H., Iter, D., Naaman, M., Gravano, L.: Identifying content for planned events across social media sites. In: ACM Conference on WSDM (2012)
4. Chen, L., Roy, A.: Event detection from flickr data through wavelet-based spatial analysis. In: ACM Conference on CIKM (2009)
5. Chengjie, S., Yi, G.: A statistical approach for content extraction from web page. Journal of Chinese Information Processing 18(5), 17–22 (2004)
6. Delgado, D., Magalhães, J., Correia, N.: Assisted News Reading with Automated Illustrations. In: ACM Conference on Multimedia, pp. 1647–1650 (2010)
7. Firan, C.S., Georgescu, M., Nejdl, W., Paiu, R.: Bringing order to your photos: Event-Driven Classication of Flickr Images Based on Social Knowledge. In: Proceedings of the 19th ACM International Conference on Information and Knowledge Management, New York, USA, p. 189 (October 2010)
8. Joshi, D., Wang, J.Z., Li, J.: The Story Picturing Engine—a system for automatic text illustration. ACM Transactions on Multimedia Computing Communications and Applications 2(1), 68–89 (2006)
9. Liu, X., Huet, B., Troncy, R.: Eurecom@ MediaEval 2011 social event detection task. In: Proceedings of the MediaEval 2011 Workshop (2011)
10. Liu, X., Troncy, R., Huet, B.: Finding Media Illustrating Events. In: ACM International Conference on ICMR, Trento, Italy (2011)
11. Mei, T., Yang, B., Yang, S.-Q., Hua, X.-S.: Video Collage: Presenting a Video Sequence Using a Single Image. The Visual Computer, 39–51 (2009)
12. Pan, C.-C., Mitra, P.: Event detection with spatial latent Dirichlet allocation. In: Proceeding of the 11th Annual International ACM/IEEE Joint Conference on Digital Libraries, New York, USA, p. 349 (June 2011)
13. Quack, T., Leibe, B., Van Gool, L.: World-scale mining of objects and events from community photo collections. In: Proceedings of the 2008 International Conference on Content-Based Image and Video Retrieval, New York, USA, p. 47 (July 2008)
14. Sakaki, T., Okazaki, M., Matsuo, Y.: Earthquake shakes twitter users: real-time event detection by social sensors. In: International Conference on World Wide Web, Raleigh, North Carolina, USA (2010)
15. Weng, J., Lee, F.: Event detection in twitter. In: AAAI Conference on Weblogs and Social Media, Barcelona, Spain (2011)
16. Yahiaoui, I., Merialdo, B., Huet, B.: Comparison of multiepisode video summarization algorithms. EURASIP J. Appl. Signal Process. 2003, 48–55 (2003)
17. Zhu, X., Goldberg, A.B., Eldawy, M., Dyer, C.R., Strock, B.: A Text-to-Picture Synthesis System for Augmenting Communication. In: Proceedings of the 22nd National Conference on Artificial Intelligence, vol. (2), pp. 1590–1595 (2007)