Annotating Real Time Twitter's Images/Videos Basing on Tweets

Mohamed Kharrat, Anis Jedidi and Faiez Gargouri
MIRACL Laboratory, University of Sfax, Sfax, Tunisia

Keywords: Twitter, Annotation, Image, Video.

Abstract: Nowadays, online social network "Twitter" represents a huge source of unrefined information in various formats (text, video, photo), especially during events and abnormal cases/incidents. New features for Twitter mobile application are now available, allowing user to publish direct photos online. This paper is focusing on photos/videos taken by user and published in real time using only mobile devices. The aim is to find candidates for annotation from Tweet stream, then to annotate them by taking into accounts several features based only on tweets. A preprocessing step is necessary to exclude all useless tweets, we then process textual content of the rest. As a final step, we consider an additional characterization (spatio-temporal and saliency) to get outcome of the annotation as RDF triples.

1 INTRODUCTION

Twitter is a microblogging service which enables people to share between them not only short messages but also multimedia contents called: Tweets. Tweets can include a lot of true and/or false information. Twitter has become an important source for news by reporting real-world events/incidents.

Our concern is about shared pictures/video which represent additional information to understand what user wants to say visually and intuitively.

Our work is motivated by the need of annotating real time and real world image/video. Those can be efficiently used in news and are required in applications that cannot afford the complexity and associated time with current image processing techniques.

On the other hand, Twitter provides unrefined data, in a timely manner so information is spreads incredibly fast and is posted before it makes it into official and suitable resources for knowledge extraction.

In this paper we address the question whether we can exploit or not this social media to extract new facts/news based on shared images or videos. We present approaches for the task of social image/video annotation. The proposed methods are based only on the tweets accompanying shared image, without the use of 'slow levels features.

We primarily use tweet level features and partially user level as authors of (Gupta, 2013) proved the primer performs the best accuracy. Then, only the evaluation make possible to distinguish the best combination of all used features.

Through the process presented in the next section we aim to a comprehension of what is happening in one's environment. Peoples , who find themselves in abnormal circumstances, can describe the current situation in real time with details using on-site information such as what is happening, where, when and who is involved.

Following (Feng, 2008) in this study, we propose that an image/video can be annotated with keywords, visual named entities and semantics interpreted attributes. We hypothesize that tweets containing an image/video are more likely to contain on-the-ground information taken and shared by eyewitnesses. Candidate tweets should be closer to the abnormal situation. In fact on-the-ground information tends to contain highly informative value.

Real-time stream of information provided by Twitter can be accessed via a single API. In addition a rich variety of sources publish information via Twitter like traditional media or citizen journalists (Hermida 2010). Tweets also contain metadata that can be exploited, like location, hashtags and user profile information. But the biggest drawback of Twitter is it noisy and unrefined shared data. Tweets can be real news but also rumors or spam.

Kharrat, M., Jedidi, A. and Gargouri, F..
Annotating Real Time Twitter's Images/Videos Basing on Tweets.
In *Proceedings of the 7th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management (IC3K 2015) - Volume 1: KDIR*, pages 293-300
ISBN: 978-989-758-158-8

Results of this study will help us to state if shared image/video within tweets indicate on-the-ground information and if they are suitable for use as news?

Current approaches exploiting social media, focus primarily on large scale incidents like earthquakes which are characterized by a big number of tweets as well as users interactions. Meanwhile, small-scale incidents/events have usually a small number of tweets and less user's interactions. This is challenging us for detecting image/video candidates for annotation.

Compared with event detection in news texts, Twitter provides more opportunities and challenges. Authors of (Tov, 2011) reported that Twitter can broadcast news faster than traditional media (except those shared and published by websites).

Actually, during any small incidents/events, thousands of microblogs (tweets) are posted in short intervals of time. Typically, only a small fraction of these tweets contribute to share new events and abnormal activities, while usual users simply share their sentiments or opinions. Real-time processing of tweets contributing to supply news is very important before it spreads widely via Internet.

Consequently, automatic differentiating of relevant tweets, to be candidates for annotation from those reflecting opinions/sentiments, is a non-trivial challenge, mainly because of tweets characteristics which include usually emoticons, abbreviations, question marks... We apply in our approach, Natural Language Processing (NLP) techniques to address this challenge to annotate shared image/video by user in their tweets. In addition, we apply several steps to detect candidates for annotation, employing filtering techniques to remove spurious events/incidents. Finally we extract terms and facts from those tweets which best characterize the situation, represent visually the image/video and are most efficacious in retrieving news.

Figure 1 shows the outlines of the baseline system steps described below.

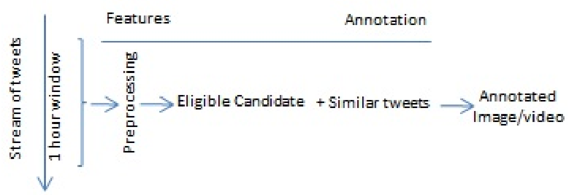

Figure 1: Outlines of the baseline system.

The rest of this paper is organized as following: Section 2 presents background and definitions. In Section 3, we discuss related work in the area of twitter processing and detecting. Section 4 and 5 describe the pre-processing step and used features respectively. We introduce the baseline system in section 6. Finally, section 7 presents our conclusion.

2 BACKGROUND AND DEFINITIONS

Twitter is a social media network, where users follow other users in order to receive information along timeline. Such information could be small text messages called tweets including multimedia resources as images or videos. There are also relationships between users, which mean each user has followers and followees. Tweets can be republished any time throughout the network, this operation is called re-tweeting. A retweeted message usually starts with "RT @username", where the @ sign represents a reference to the user who originally published the message. Users could also use hashtags (#) to identify certain topics. Hashtags are similar to tags. The most used hashtags in the network become trending topics reflecting incidents/ events.

3 RELATED WORK

Through the process presented below, we are able to perform annotation of images/videos from unrestricted domains using only content of tweets.

Most of previous work on Twitter was done on event detection and most of them detect only specific types of events (Qin, 2013) likewise earthquake events detection from Twitter (Sakaki, 2010).

Authors of (Corvey, 2012) analyzed one of the major aspects of applying computational techniques and algorithms to social media data in order to obtain useful information i.e. linguistic and behavioral annotations.

Authors of (Dave, 2010) reported that classification performance is significantly degraded in the more practical cross domain classification. This practical limitation motivated us to avoid using classification, in addition to the fact that we cannot limit the domain, since we are not detecting events, but images/videos about small incidents, events…

Majority of the previous works on event detection using social media has focused on using topic detection methods to identify breaking news stories. Streaming document similarity measures

(Petrovic, 2010), (Osborne, 2014) and online incremental clustering (Becker, 2011) have been shown to be effective for this purpose. Other approaches have aimed to pick up more localized events. These have included searching for spatial clusters in tweets (Osborne, 2013), leveraging the social network structure (Aggarwal, 2012), analyzing the patterns of communication activity (Chierichetti, 2014) and identifying significant keywords by their spatial signature (Abdelhaq, 2013).

Work has been done to extract situational awareness information from the vast amount of data posted on OSM (Online Social Media) during real-world events. (Nicholas, 2015) adapt existing bio-surveillance algorithms to detect localized spikes in Twitter activity which could be classified as real events with a high level of confidence.

Regarding tweets processing, the biggest problems are their contents and their important amount which are considered as an important resource that can play a critical role in crisis (Palen, 2010). However working on tweets, needs also a filter step because many of them are fake or rumors (Palen, 2010). In fact, authors of (Gupta, 2013) proposed a solution to characterize and identify the propagation of fake pictures on online social media during Hurricane Sandy.

(Li, 2011) introduced a system for searching and visualization the tweets related to small scale incidents, based on keyword, spatial, and temporal filtering.

Most existing approaches are focused on analyzing text-based messages from Twitter. Meanwhile multimedia (image/video) based approaches have not been extensively addressed hitherto despite the fact that real time images/video shared on social media can refer to valuable information towards improving news. This work tries to confirm this hypothesis by exploiting this type of media through tweets.

Indeed, we hypothesize that relevant tweets containing an image/video are more likely to contain on-the-ground information – for example photos taken and posted by eyewitnesses. Consequently such tweets should be closer to an incident/event, i.e. affected areas. Therefore, we investigated the following research questions: Does the existence of image/video within tweets indicates on-the-ground information and is thus suited for news?

Few studies interested in annotation images/video shared on twitter or extracting information from them. Most of them use low levels features which could never be good for real-time processing. Several of them on the contrary, use external photos from other social media to annotate tweets like the approach proposed in (McParlane, 2014) by exploiting Twitter and Wikipedia for the annotation of event images.

In addition, there is a lot of research work analyzing accompanied text on pictures (Srihari, 1994) but, there has been less works on detecting events or analyzing contents of images shared on Twitter.

Authors in (Raad, 2014) proposed an application where events are detected from photos, capturing the Where, When and Who dimensions of an event, and describing (temporal, spatial, and semantic) relationships between events using only image metadata. Another work for Automatic image annotation using auxiliary text information was presented in (Feng, 2008).

In (Phuvipadawat, 2010) authors proposed an approach to find the most frequent image related to a tweet and return it from internet.

In (Leong, 2010) authors introduced extractive approach for automatic image tagging by natural language resources for processing texts surrounding images. However this approach used Flickr repository and Wikipedia which could not be convenient for tweets. Finally, advanced approach in (Chen, 2013) aims to mine salient images related to one specific object by proposing an image clustering and ranking algorithm.

4 PREPROCESSING

We describe here, several steps to apply on tweet stream to eliminate all useless tweets.

To capture duplicated tweets, it is possible to modify the pre-processing also in order to duplication capture, e.g. one could filter out the "RT" string and the user mentioning and repeating the same hashing procedure; or one could detect near duplicates using Jaccard similarity (using also an inverse index for speed).

In fact, we do not exploit redundant/duplicated tweets. If we consider a tweet as relevant candidate for annotation, then if the same information is shared and published after a while (out of window time), it will not be considered as new anymore but old one. Otherwise, we take in consideration redundant tweets if only they were shared in the window interval, which suppose that users are in the same location and watching the same incident/event.

To avoid spam tweets from bots, and basing on the tweeting interval (Chu, 2010), we can detect automated users who tend to have a periodic or regular timing of tweets. Though, it is possible that bots overcome such detection but it still an additional filter for us.

In the same context, non-spammers users spend more time interacting with other users (Benevenuto, 2010). In order to exclude tweets of spammers users, we looked at a several of features of them, like the number of times the user was mentioned by other users and number of times the user was replied.

All tweets not containing real images or video taken from mobiles will be eliminated.

All tweets containing URLs or shared image or video from Internet or other mobile apps will also be eliminated.

Further, we remove tweets automatically generated by check-in services such as Swarm by detecting the patterns “I’m at” and “mayor”.

In fact, Twitter applications designed specifically for mobile devices (e.g., twitter for iPhone/Android) are frequently used in author tweets and used by individual persons. Organizations, unlike, primarily use the Twitter web version and content management software applications to publish and manage content on Twitter (De Silva , 2014)

Terms with less than three characters: trigrams will be eliminated. stopwords and performing POS tagging will be removed as well as tweets containing smileys (mdr, ptdr…).

We consider only English tweets reducing thus the number of tweets that need to be processed in further steps.

Last process is to convert slang words to real words using slang dictionary, i.e convert “abwt” to “about”.

5 FEATURES

This section describes the linguistic features used for distinguishing candidate tweets or to filter non eligible tweets.

The parts-of-speech (POS) tags are identified using a probabilistic tokenizer and POS tagger designed explicitly for tweets (Owoputi, 2013), which can also identify emoticons and exclamations.

The required feature extraction step, is based on the content and the user information of tweets. Those are summarized in Table 1 and described in the following paragraph.

Table1: Content and user features.

Content features
Length of tweet
Contains question marks
Number of retweets
Number of hashtags
Number of Favorite
User features
Number of followers
Number of tweets
Number of following friends
Follower-friend ratio

5.1 Content Features

These features are solely based on the content of tweets. We rely on the features used by Gupta et al. (Corvey, 2012), to which we add the number of retweets for each tweet. The features are listed in Table 1. Starting from the tweet characteristics, we compute features such as the length of the tweet and the number of words it contains. Also, we include features such as the number of question marks (as part of expressing shock, disgust…) and exclamation marks and exclamatory words (e.g., :(, ‘omg!’, ‘oh no!’). We expected incident tweets to contain exclamations much more frequently than ordinary tweets.

We take into account the sentiment of the tweet which usually contains a higher fraction of ‘subjective’ words, relying on a predefined list of strongly subjective words and subjectivity lexicons (Volkova, 2013) specifically developed for tweets. We compute the number of subjective words it contains.

Unlike headline verbs which tend to be more formal, personal verbs tend to represent personal activities, communications, and emotions. We use a verbs list identified by authors of (De Silva , 2014) containing 2221 personal English verbs.

After detecting the text language of each of the tweets using an open language detection library, a list of tweet features for each Tweet of the stream is produced.

5.2 User Features

We also extract features from the Twitter user, the author of the post. It includes user’s number of friends and followers (In fact there are opportunistic users that follow important number of people in order to be followed back), as well as frequency of tweeting.

6 BASELINE SYSTEM

We try to find the names of persons who are both visible in the image/video and described in the tweet. A baseline system will be started when all persons found in the text are assumed present in the shared image/video. This assumption depends of the precision of the used NER (Named Entity Recognition) and percentage of tweets that discuss people not present in the image. In some cases this problem can be solved by using a pattern of syntactic formulas but developing a system that could extract this information is not trivial, and even so only a very small percentage of the texts in our test corpus contain this kind of information.

6.1 Mechanism

Most of tweets are not real world stories, but rather talks about personal life, conversations, expressing humor, or spam. Running a first story detection or event detection system on this huge data would yield an important amount of new stories every day, most of which would be interesting only for few people. However, when something significant happens (e.g., a minister seen in a restaurant), a lot of users write about this, either to share their opinion or just to inform others of the situation. Our goal here is to automatically detect these significant candidates (i.e Figure 2) for annotation, with a minimal of information.

Figure 2: Example of candidate Tweet.

We considered bursts of tweets what appear after an event happens. We exploited this fact and tried to measure the number of times the given entities refers to the main tweet on Twitter recently.

Over 1 hour of fixed window time, and after a tweet have been trigged the process by analyzing the sudden peaks (number of retweet and of replies reached), we crawled the follower list of all the unique users that had tweeted the original tweet to compute the overall of followers users and those who retweet without following the original user.

We did not consider all retweets of the original one, because according to research in (Gupta, 2013) 86% tweets spreading the fake images were retweets, hence very few were original tweets. In fact, at crisis time, users retweet anything from each other regardless of the fact whether they follow them or not.

Thus, the semantic representation of a tweet "t" consists of keywords (KW(t)) and keyphrases (KT(t)).

In fact segments can be advantageous for tweet processing as they have much smaller quantity than tweets itself and are more semantically, more meaningful than keywords (Li, 2011).

We extracted the longest sequence of nouns as well as proper nouns/Named entities and keyphrases using POS tagger designed explicitly for tweets (Owoputi, 2013).

We utilized lemmatized terms instead of raw terms through Wordnet

We considered nouns, hashtags and proper nouns as keywords.

Next, a graph based tweets was created consisting of nodes which represents tweets while the similarity between two tweets is represented by the edge as weight. For similarity between two tweets, t1 and t2 we adopt Eq. 1 proposed by (Panem, 2014) authors with a small adaptation.

$$\begin{aligned}Sim(t1,t2) = &w \times sim(KT(t1), KT(t2)) + (1-w) \\ &\times sim(KW(t1), KW(t2))\end{aligned} \quad (1)$$

Where

$$Sim(KT(t1), KT(t2)) = |KT(t1) \cap KT(t2)| \quad (2)$$

And

$$Sim(KW(t1), KW(t2)) = |KW(t1) \cap KW(t2)| \quad (3)$$

Here, "w" denotes the weight given to the keyphrases and $(1 - w)$ denotes the weight given to the keywords. In our experiments, "w" will take at the beginning many values in order to determine the best one.

Stream of tweets is unbounded, for this reason we do not store all the previous data in main memory nor compare the new document to all the previous tweets.

Automatic distinguish between a candidate tweet or not, should be made in bounded time (preferably constant time per document), and using bounded space (also constant per document). We have chosen 1 hour as property of our approach.

This property allows us to avoid limiting the number of documents inside a single bucket to a constant. If the bucket is full, the oldest document in the bucket is removed. Note that the document is removed only when it exceed 1 hour. Note that this way of limiting the number of kept documents is in a way topic-specific.

To achieve that, each tweet should be processed as it arrives with building up dynamically tweet clusters representing events. In particular, for each incoming tweet, it should be compared against the stream of previously seen tweets using a fast hashing strategy

If the current tweet is sufficiently (textually) dissimilar from its nearest neighbor, it is flagged. The system attempts to reduce false positives by waiting for a very short deferral period of 1 hour, thus it can collect all follow-up posts and produce clusters of closely related tweets. For an event cluster, the tweet closest to the centroid of the cluster (using a standard vector space) is emitted.

We boost the terms that correspond to named entities and hashtags by some constant factor

6.2 Additional Characterization

6.2.1 Saliency

We use salience to determine what terms will be used for annotations and included in triples.

We define here, salience measure, which represents a value between 0 and 1 that refers the importance of an entity in a tweet. Usually it is calculated simply as tf (term frequency) x idf (inverse document frequency) of terms that represent the entity in the text.

However, for tweets which are short texts, another measure is necessary, because almost all entities are only mentioned once so we need a reliable way to discern their salience.

To determine which terms are salient and describing well the shared image/video, we calculate the co-occurrence of terms and its variants, for candidate tweets and its replies. We estimate the term co-occurrence statistics with the user frequency of a term: the number of people using that term in a given location. In this case, the user frequency is not significant, so term co-occurrence is computed within the term frequency of all similar tweets.

Term co-occurrence is traditionally computed as the number of times term ‘t’ and term ‘w’ appear in the same tweet TW, divided by the number of times term ‘t’ and ‘w’ appear in any tweet in the same window time and in close location.

Some users of twitter are very active, so they may generate a lot of tweets in a small interval of time. Term frequency may produce an estimate of the term distribution biased toward a particular user or set of users. To avoid dominating messages from one user, we estimate the term co-occurrence with the user frequency. This is proved to be efficient by author in (O’Hare, 2013)

For better understanding, let’s consider the following scenario. Assume that a user took with his smart phone a photo of Michael Schumacher leaving the Grenoble hospital after 3 years. There are two cases, conditional by enabling the GPS on this phone: Or the user tags the tweet with GPS location/coordinates, either the GPS receiver of the phone calculates the latitude and the longitude of the location and the data are stored in the Exif descriptor of the shared photo/video.

6.2.2 Geolocation

Then, the system applies a geobased search on following tweets to find user-tagged tweets taken by other users in the surroundings of this hospital by applying similarity between tweets.

We identify the tweets about the same topics by looking for same terms. We compute the term co-occurrence between terms in the main tweet, and the terms that occur in following tweets.

From the retrieved set, the system begins processing to select the tags of the visually matching image/video which can be used to produce annotation.

There is difference in case of tweets which are not geotagged. Indeed, in such case, we identify locations using Stanford NER. Secondly, to relate the location mention to a point where the incident happened, we geocode the location strings. In this case, we create a set of word unigrams, bigrams, and trigrams. These are sent to the geographical database GeoNames to identify city names in each of the n-grams and to extract geocoordinates. As city names are ambiguous around the world, we choose the longest n-gram as the most probable city.

6.2.3 Temporal

To identify temporal dimension of a tweet, we adopt the HeidelTime (Jannik, 2010) framework for temporal extraction. HeidelTime is a rule-based approach that extracts temporal expressions from text documents and normalizes them according to the TIMEX3 annotation standard.

6.2.4 Facts Extraction Triples

A metadata, describing multimedia resource image/video, consists of a set of attributes. We formally represent a metadata of this resource as: res: (G, T, O, F) where:

- G represents the geo-location of the resource when captured (i.e., GPS coordinates or Geoname),
- T represents the creation date/time of the resource,
- O represents a set of objects of interest identified in the resource (i.e. person, monument….).
- F represents facts

Let 'res ∈ I' be an image or video shared by a user. Formally, we represent a resource as follows: res: (resid, meta)

Where:

- resId: is the identifier of the resource,
- meta: is the metadata describing the resource.

We use (REVERB), open information extraction system (Fader, 2011), to filter out relevant facts as RDF triples. We evaluate facts in terms of their well-formedness, their correctness, and their relevance.

RDF Subjects describe agents, causers, or experiencers, while RDF objects describe entities undergoing a state of change or being affected by an action.

Triples are characterized by (i) the first argument of the extracted triple is one of the named entities in candidate tweet, (ii) the most frequent sense of the verb has the super sense stative, possession, competition or creation according to WordNet, and (iii) t none of the arguments are stop words. We then extract triples made up of a verb and the head words of the two arguments.

After creating RDF subject and objects corresponding to RDF triples, not of the main tweet but of related tweets (replies), we select only candidates named entity that have high salience.

7 CONCLUSIONS

Social media networks, particularly Twitter, become strong and fast broadcaster of news. In this paper, we propose an approach to annotate automatically multimedia shared documents (videos or images) in Twitter social media. It corresponds to real world resources taken and shared only by mobile devices, not using web interface nether sharing applications. We aim to enhance news with unshared and unpublished incidents and events. Before the annotation we detect candidates for annotation, and then we apply several steps processing contents of tweets in order to produce annotation as RDF triples. After crawling tweets using Twitter API, which return about 1% of all current tweets, next step will be the evaluation of this approach.

REFERENCES

Abdelhaq, H., Sengstock, C., and Gertz, M., 2013. Eventweet: Online localized event detection from twitter. In *Proceedings of the VLDB Endowment 6(12):1326–1329.*

Aggarwal, C. C., and Subbian, K., 2012. Event detection in social streams. In *SDM, volume 12, 624–635. SIAM.*

Benevenuto, F., Magno, G., Rodrigues, T., Almeida, V., 2010. Detecting Spammers on Twitter. In *Collaboration, Electronic messaging, Anti-Abuse and Spam Conference.*

Becker, H., Naaman, M., Gravano, L., 2011. Beyond trending topics: Real-world event identification on twitter. In *Proceedings of the Fifth International Conference on Weblogs and Social Media, Barcelona, Catalonia, Spain. ICWSM 11:438–441.*

Chen, X., Chen, M., Kim, EY., 2013. Mining salient images from a large-scale blogosphere. In *Proceeding of 8th International Conference for Internet Technology and Secured Transactions (ICITST).*

Chu, Z., Gianvecchio, S., Wang, H., and Jajodia, S., 2010. Who is tweeting on Twitter: human, bot, or cyborg?. In *Proceedings of the 26th Annual Computer Security Applications Conference Austin, Texas: ACM, 2010.*

Chierichetti, F., Kleinberg, J., Kumar, R., Mahdian, M., and Pandey, S. 2014. Event detection via communication pattern analysis. In *Eighth International AAAI Conference on Weblogs and Social Media.*

Corvey, WJ., Verma, S., Vieweg, S., Palmer, M., and Martin, JH., 2012. Foundations of a multilayer annotation framework for twitter communications during crisis events*. In Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12), Istanbul, Turkey.*

Dave, KS., Varma, V., 2010. Pattern Based Keyword Extraction for Contextual Advertising. In *Proc. of the 19th ACM Intl. Conf. on Information and Knowledge Management (CIKM), pages 1885–1888.*

De Silva, L., Riloff, E., 2014. User Type Classification of Tweets with Implications for Event Recognition. In *Proceedings of the Joint Workshop on Social Dynamics and Personal Attributes in Social Media.*

Fader, A., Soderland, S., Etzioni, O., 2011. Identifying Relations for Open Information Extraction. In *Proceedings of the Conference of Empirical Methods in Natural Language Processing, EMNLP.*

Feng, Y., and Lapata, M., 2008. Automatic image annotation using auxiliary text information. In *Proceedings of the Association for Computational Linguistics.*

Hermida, A., 2010. From TV to Twitter: How Ambient News Became Ambient Journalism. In *Media/Culture Journal, Vol. 13, No. 2, May 2010.*

Jannik, S., Gertz, M., 2010. HeidelTime: High Qualitiy Rule-based Extraction and Normalization of Temporal Expressions. In *SemEval'10.*

Gupta ,A., Lamba, H., Kumaraguru, P., Joshi, A., 2013. Faking Sandy: Characterizing and Identifying Fake Images on Twitter during Hurricane Sandy. *In Proceedings of the 22nd international conference on World Wide Web companion.*

Li, R., Lei, K.H., Khadiwala, R., Chang, K.C.C., 2011. Tedas: A twitter-based event detection and analysis system. In *11th International Conference on ITS Telecommunications (ITST).*

McParlane, PJ., Jose, J., 2014. Exploiting Twitter and Wikipedia for the Annotation of Event Images. In *Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval.*

Nicholas A. Thapen, Donal Stephen Simmie, Chris Hankin. 2015. The Early Bird Catches The Term: Combining Twitter and News Data For Event Detection and Situational Awareness. In *Journal: CoRR Vol. abs/1504.02335.*

Leong, CW., Mihalcea, R., and Hassan, S., 2010. Text Mining for Automatic Image Tagging. In *Proceedings of the 23rd International Conference on Computational Linguistics.*

Osborne, M., Moran, S., McCreadie, R., Von Lunen, A., Sykora, M. D., Cano, E., Ireson, N., Macdonald, C., Walther, M., and Kaisser, M. 2013. Geo-spatial event detection in the twitter stream. In *Advances in Information Retrieval. Springer. 356–367.*

Ounis, I., He, Y., et al., 2014. Real-time detection, tracking, and monitoring of automatically discovered events in social media. In *Proceeding of the 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations, Baltimore, MD, USA, 23-24 Jun 2014, pp. 37-42.*

Owoputi, O., O'Connor, B., Dyer, C., Gimpel, K., Schneider, N., and Smith, NA., 2013. Improved part-of-speech tagging for online conversational text with word clusters. In *Proceedings of NAACL-HLT, 2013, pp. 380–390.*

O'Hare, N., Murdock, V., 2013. Modeling locations with social media. *Journal of Information Retrieval, 16(1).*

Palen, L., Anderson, KM., Mark, G., Martin, J., Sicker, D., Palmer, M., and Grunwald, D., 2010. A vision for technology-mediated support for public participation & assistance in mass emergencies & disasters. In *Proceedings of the 2010 ACM-BCS Visions of Computer Science Conference, ACM-BCS '10.*

Panem, S., Bansal, R., Gupta, M., Varma, V., 2014. Entity Tracking in Real-Time using Sub-Topic Detection on Twitter. In *Proceedings of 36th European Conference on IR Research, ECIR 2014, Amsterdam.*

Petrovic, S., Osborne, M., Lavrenko, V., 2010. Streaming first story detection with application to twitter. In *Proceeding of HLT '10 Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics.*

Phuvipadawat, S., Murata, T., 2010. Breaking news detection and tracking in twitter. In *Web Intelligence and Intelligent Agent Technology, IEEE/WIC/ACM International Conference on, 3:120–123, 2010.*

Qin, Y., Zhang, Y., Zhang, M., Zheng, D., 2013. Feature-Rich Segment-Based News Event Detection on Twitter. In *International Joint Conference on Natural Language Processing IJCNLP 2013.*

Raad, EJ., Chbeir. R., 2014. Foto2Events: From Photos to Event Discovery and Linking. In *Online Social Networks. SocialCom2014, Dec 2014, Sydney, Australia.*

Sakaki, T., Toriumi, F., and Matsuo, Y., 2011. Tweet trend analysis in an emergency situation. In *Proceedings of the Special Workshop on Internet and Disasters, SWID '11, pages 3:1–3:8, New York, NY, USA, 2011. ACM.*

Srihari, R., Burhans, D., 1994. Visual Semantics: Extracting Visual Information from Text Accompanying Pictures. In *Proceedings of AAAI-94.*

Tov, EY., Diaz, F., 2011. Location and timeliness of information sources during news events. In *Proceedings of SIGIR, pages 1105–1106, Beijing, China.*

Volkova, Q., Wilson, T., and Yarowsky, D., 2013. Exploring Sentiment in Social Media: Bootstrapping Subjectivity Clues from Multilingual Twitter Streams. In *Proc. ACL Vol2.*