Social Signal Processing for Real-time Situational Understanding: a Vision and Approach

Kasthuri Jayarajah*, Shuochao Yao†, Raghava Mutharaju§‡, Archan Misra*, Geeth De Mel§, Julie Skipper¶, Tarek Abdelzaher†, and Michael Kolodny‖

*Singapore Management University, †University of Illinois, ‡Wright State University, §IBM T.J. Watson Research Center, ¶Air Force Research Laboratory, ‖US Army Research Laboratory

***Abstract*—The US Army Research Laboratory (ARL) and the Air Force Research Laboratory (AFRL) have established a collaborative research enterprise referred to as the Situational Understanding Research Institute (SURI). The goal is to develop an information processing framework to help the military obtain real-time situational awareness of physical events by harnessing the combined power of multiple sensing sources to obtain insights about events and their evolution. It is envisioned that one could use such information to predict behaviors of groups, be they local transient groups (e.g., protests) or widespread, networked groups, and thus enable proactive prevention of nefarious activities. This paper presents a vision of how social media sources can be exploited in the above context to obtain insights about events, groups, and their evolution.**

I. Introduction

This paper presents a vision and high-level architecture for the exploitation of social media to monitor and understand physical events. The proliferation of social media, such as Twitter, Instagram, and Pinterest begs the question of whether social networks can be leveraged as sensor networks to observe the physical world [22]. Their exploitation is especially beneficial in urban spaces, where a large human population interacts with and reports everyday events. According to the United Nations, presently 54% of the world population live in cities. This percentage will increase to 66% by 2050[1]. Arguably, the most versatile sensor in urban areas is the human observer. Collectively, human observers post over 500 million tweets and over 70 million Instagram photos per day, making social media an excellent *spontaneous* and *distributed* resource for obtaining insights on a variety of events.

Exploitation of social networks as sensor networks has many applications. Human observers could provide us with information regarding suspicious activities with respect to national security scenarios. They are the survivors and first-responders in post-disaster operations. They are the friendly locals in peacekeeping and stabilization missions. They could offer real-time information on unfolding, unpredictable gatherings (e.g., demonstrations and protests), and unusual events that impact safety, such as escalating confrontations or high-speed pursuits. Such information correlated, corroborated, and used in the right context could provide significant clues into both *(a)* the persona and relationships among individuals; and *(b)* the nature and chronology of events in the physical world.

[1]http://www.un.org/en/development/desa/news/population/world-urbanization-prospects-2014.html

Motivated by this observation, we seek to develop an information processing framework to attain real-time situational awareness of physical events and group behaviors by harnessing the combined power of multiple social media platforms. Given the inherently *crowdsourced* nature of social media, this framework must support several key capabilities to truly offer an enhanced situational understanding:

- *Information Relevancy:* Social media carry significant amounts of noise together with useful data. One needs resilient filtering approaches that are aware of the context in which the information is collected, so that the most appropriate information related to the situation at hand is identified.
- *Information Credibility:* Information disseminated on social media is contaminated by human bias, errors, misperceptions, and malicious distortions. Due to the inherent uncertainty regarding trustworthiness of sources and provenance of data, one needs novel algorithms to establish information credibility [23].
- *Multimodal Synthesis:* Social media content has become progressively more diverse and now routinely includes concise text (e.g., Twitter), images (e.g., Instagram) and semantic locations (e.g., Foursquare), among others. One needs appropriate fusion strategies to harness the diversity of multimedia content across such distinct "channels" of social sensing. User interaction with each of these channels has distinct behavioral differences, making them nicely complementary and mutually enhancing.
- *Knowledge Representation and Reasoning:* To automate analysis of physical situations, a machine needs to "understand" the meaning of the social media posts. Knowledge representation and reasoning techniques may be used [4]—especially Semantic Web technologies [9, 18]—to semantically analyze the multimodal content.
- *Information Value:* Collection and analysis of relevant and trustworthy multimodal information on events and sources is not sufficient for deriving value for an application. For example, redundant or outdated information may not be useful. Hence, one needs a mechanism for establishing the true value of information, given questions of interest and current context.

The architecture presented in this paper is inspired by the above needs.

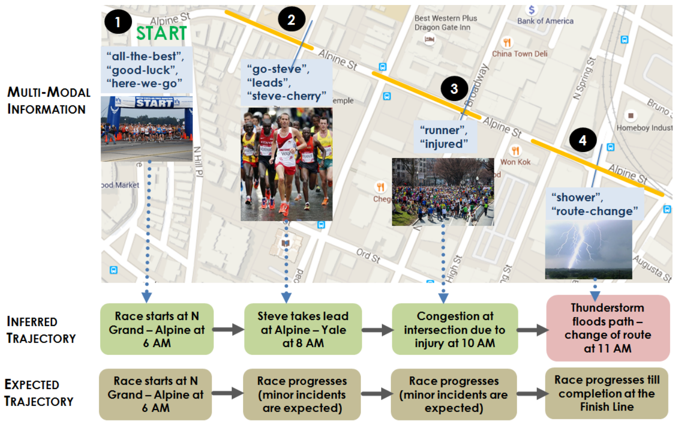

Figure 1. Multimodal social network information for situational understanding—a marathon example.

The rest of the paper is organized as follows. In Section II, we provide an illustrative example of social sensing to offer a context for later discussion. Motivated by this example, in Section III, we present an information processing architecture that addresses the requirements mentioned above. We then discuss its key components and underlying technologies. In Section IV, we highlight research challenges and open avenues for investigation, brought about by this architecture. We present related work in Section V, and conclude the paper in Section VI.

II. An Illustrative Example

In this section, we provide an example of situational understanding to illustrate the main goals of our proposed research. To this end, consider a marathon event, illustrated in Figure 1.

Imagine that the marathon will take place in the city of Los Angeles, today. Twitter activity related to the marathon will probably start a few days before the event, peak during the event and persist for a few days after the event. From such tweets, one can map the community of sports-enthusiasts, as well as their internal influence structure, by observing who forwards whose posts. We can thus establish a perceived authority level of the individuals commenting on the marathon. During the actual event, real-time tweets can inform us of the general progress of the race, as well as any notable incidents that might be occurring in the race. Location references within tweets can offer a first level of event localization.

Coincidentally, spectators as well as the runners themselves may use other social media platforms, such as Instagram and Vimeo, for posting pictures and videos from the event, respectively. Exploiting these additional media may offer a next level of event understanding. For example, the appearance of an "ambulance" object in a certain subset of an uploaded geo-tagged image may indicate the occurrence of injuries to runners at a specific location. Alternatively, a combined spatial correlation of the intensity of tweets and Instagram images may provide a timely picture of the flow of runners as the event progresses. In Figure 1, Stages 1, 2 and 3 illustrate three sub-events at the marathon – start of the race, one of the runners taking a clear lead, and a runner getting injured. The shaded blue box contains keywords used by Twitter users to describe these sub-events and the images below are examples of those shared on Instagram. The geotags available with Instagram posts help localize the sub-events.

Hence, situational understanding could be significantly enhanced if we were able to accurately represent the relationships between the multimodal content across diverse social media platforms. Concepts such as *runner injured* or *congestion at the intersection* (from Twitter feeds) might be correlated and corroborated to provide a common understanding of a possibly minor incident and its (perhaps short term) effects (Stage 3 in Fig. 1). Moreover, relationship models could exploit relationships across different forms of social media to further improve situational understanding (e.g., extract *crowd levels* from Instagram pictures through simple image processing, and relate it to the *congestion near intersection* concept).

As the marathon progresses, the event evolves both spatiotemporally and semantically, allowing us to conceptualize events as *trajectories* in these spaces.

A *normal* marathon's trajectory would typically begin with runners warming up near the start line, runners spreading along the course after the race begins, and finally reaching the finish line with the crowd cheering the winners. Let us now assume that due to a heavy downpour, part of the marathon pathway is flooded. The organizers of the event have anticipated this and the path of the marathon is changed at the last minute (Stage 4 in Figure 1).

Such deviations in the event *trajectory* can be detected and corroborated in real-time by leveraging multimodal information available in social media posts by defining appropriate notions of normalcy (i.e., expected trajectories) to eventually determine anomalies. Note that, the same information processing capabilities can help understand other types of events that may last longer or be more geographically dispersed. In the next section, we describe the architecture we envision to realize situational understanding.

III. System Architecture

In this section, we describe the key components of an architecture that offers some of the capabilities motivated by the above example. In Figure 2, the components shaded in blue represent those that are necessary for filtering information that is *relevant* to the context and events of interest (e.g., protests around the world), while the components shaded in green represent those that are necessary for gaining an *understanding* of such events. The components in pink represent our a priori knowledge of the events. In the following subsections, we briefly describe these components and their collective goals.

A. Multimodal Social Sensing

We consider incoming streams of user-generated text (e.g., tweets), images (e.g., Instagram) and video (e.g., Vimeo) data. The **relevancy filters** remove out-of-context data. The filtering criteria include keywords, geographical regions, and other

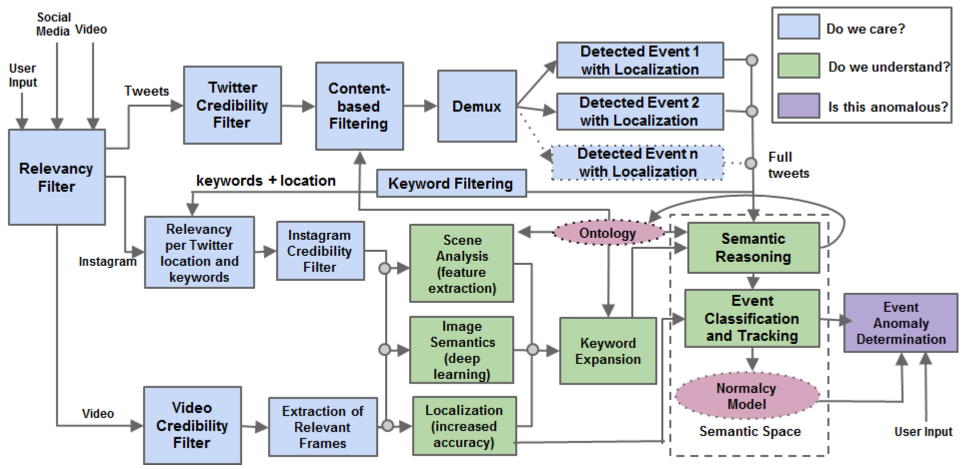

Figure 2. Proposed architecture of the real-time situational understanding framework.

information provided by the user. In addition, in the case of video feeds, consecutive frames with no or minimal change in content are removed as this reduces the computational cost. Relevant feeds are then passed through **credibility filters**. Since humans are biased by their experiences and viewpoints, establishing credibility of the reported observations is an important part of the architecture. In particular, Twitter feeds, which are used to detect events (Section III-B), are first vetted for credibility using algorithms that jointly estimate source trustwrthiness and information veracity. In contrast, an image feed, which is additionally filtered by keywords pertaining to detected events, is vetted by assessing whether an image pertains to the physical event of interest, or is reused from another one. The **content-based filter** and the **keyword expansion** blocks use prior knowledge from the ontology to exploit keywords pertaining to events.

It is reasonable to assume that while users could write about an event from anywhere in the world (as in the case of tweets), they are more likely to post an image about an event only when they are on location (as in the case of Instagram). Hence, images offer a better opportunity for event localization compared to Twitter-based analysis of location references (such as the work of Giridhar *et. al* [7]). The **localization** block utilizes geotags that accompany Instagram posts to localize physical events. Note that, most events are composed of many sub-events of interest. For example, in the case of a marathon, a runner getting injured during the run is a sub-event of the main marathon event. Such sub-events typically involve many users, roughly within the same geographical area, posting images that may share common points of interest that are likely to be oriented differently. Based on the location of the source and the relative orientation of select key points in the images, accurate sub-event location information can be computed.

A key requirement of the architecture is that of understanding the meaning of the images (and frames in videos) to complement the Twitter-inferred insights. Thus, a mechanism to transcribe images into textual information becomes necessary; the process to achieve this is two-fold. First, the **image semantics** block generates sentence descriptions of the images based on a pre-trained convolutional neural network model similar to NeuralTalk [5]. Secondly, the images are passed through a series of classifiers to detect the presence (or absence) of particular objects in **scene analysis** (e.g., objects including *fire*, *police car*, and *ambulance* are reliable indicators of the level of violence in a scene and better aids in gauging the situation).

B. Event Identification and Understanding

A key needed capability is to distinguish between multiple distinct events of the same type from the incoming stream of credible tweets that passed the relevanc filters. Towards this end, the **demultiplexer** separates out the single incoming data stream into multiple event-specific streams along with their locations. For example, a feed reporting all protest-related tweets can be demultiplexed into multiple feeds, each on a specific protest. New events may be generated over time, and old events are eventually removed. Each event has a finite lifespan during which the event is said to be ongoing. Additionally, the **keyword filtering** block extracts meaningful and discriminating keywords for each demultiplexed event. These event-specific keywords can then be used for filtering image and video feeds.

Event understanding includes classifying, tracking and predicting the stages of evolving events in the spatiotemporal and semantic spaces. In the general domain of sensing, this is analogous to object classification, object tracking, and object trajectory prediction. In the **event classification and tracking** component, we regard detection of certain words on social data as an act of semantic sensing and map "semantic sensors" (word detectors) onto a high-dimensional semantic space. The geometric positions of these word "sensors" allow for *localization of objects* in the semantic space, much like a thoughtful sensor deployment allows physical signal detection and object localization in physical space. Further, tracking and prediction algorithms in this semantic space facilitate the process of understanding the evolution of event trajectories. Knowledge representation techniques (described next) help reduce the dimensionality of this semantic space, thereby improving the fidelity and confidence of the event understanding process.

C. Knowledge Representation and Reasoning

In order to make sense of user generated content—be it tweets or annotations associated with images—in a particular domain, we need a formal model for it. An **ontology** enables us to capture important domain concepts, their properties (or features), and relationships among them. In addition, ontologies allow us to implement mechanisms such that computational procedures can automatically make use of the captured (or inferred) knowledge.

The **content-based filtering** block utilizes ontology services to make sense of available information by consuming the relationships and the class hierarchies defined within the ontology, whereas the blocks within the **semantic space** utilize the exposed capabilities to perform further inferences and support **event classification and tracking** by exposing possible **normalcy models**. Another benefit of having a service-based model to access the ontology is that we can always augment the ontology—or introduce a new ontology model—without having to change any of the blocks within the architecture programmatically. Moreover, we use **reasoners**[2] to infer logical consequences from facts in an ontology, thus making implicit facts explicit. After converting social media data (e.g., tweets) into axioms, reasoning on these facts helps us to determine the correlations (e.g., a riot is a kind of a protest) or conflicts (e.g., based on a set of features, an event $\mathcal{E}$ seems both violent and non-violent protest) among the data items.

A synonym service is also implemented so that it can be coupled with the querying services such that similar keywords are taken into consideration while answering queries. The service is utilized in the **keyword expansion** block to correlate keywords from tweets with concepts and relationships from the ontology. There is a multitude of synonym databases, especially for social media domains: Freebase[3] and Tweet NLP[4] to name a few; we have utilized them in our work.

D. Normalcy Models and Anomaly Determination

An important goal of the framework is to be able to determine anomalies as they occur based on our developed understanding of the events. To determine anomalies, it is essential to define what is normal. The **normalcy model** comprises prior knowledge of events and provides guidance on the expected trajectories for each event class. Prior knowledge could be either user-specified or based on analysis of large-scale historical data. The **event anomaly determination** block compares the observed trajectory of the event to the normalcy model in order to determine anomalous behavior.

[2] e.g., Pellet—https://github.com/complexible/pellet

[3] http://www.freebase.com

[4] http://www.ark.cs.cmu.edu/TweetNLP

IV. Research Challenges

Our architecture combines several components, some of which exploit the existing state of the art, whereas others offer avenues for future research. In this section, we discuss open technical challenges that arise from the aformentioned components.

A. Multimodal Social Sensing Under Uncertainty

As with any social sensing application, a key challenge to be addressed arises from the lack of information reliability. Unlike sensors that are objective and unbiased, humans are biased by their experiences, viewpoints and opinions, as well as their social ties. It is non-trivial to quantify reliability of sources. Research on algorithms for establishing truthfulness of sources and credibility of observations remains an important active area in social sensing. Another important aspect of situational understanding is to determine locations of events. Often, this is achieved by inspecting location references in text, since posts on social media are often not geo-tagged. Although the resulting localization can be reasonably accurate over large spaces, it may generate unacceptable inaccuracy when one needs to establish fine-grained locations of sub-events and entities involved. Unlike infrastructural sensors such as CCTV cameras, whose exact location and orientation are known, human observers may generally have unknown location, mobility and orientations. These uncertainties warrant probabilistic inference models that perform robustly even in the presence of unknown errors. An unexplored and exciting possibility is the notion of *semantics-aided localization*, where event locations are determined not solely on the basis of direct references to landmarks or available location data (e.g., GPS tags associated with Instagram images), but refined in a more cascaded or coupled fashion, based on relationships among events and sub-events (only some of which may have explicit location indicators).

B. Formal Knowledge Representation

To perform reasoning over tweets, one could map them to specific concepts and relations in an ontology. Mapping involves identifying specific terms in the tweets that can be considered as an *instance* of a concept in the ontology (e.g., *Mazda* is an instance of the *Car* concept). This allows one to map the informal language of tweets into ontological concepts and relations that are clean and crisp. It is possible to use parts-of-speech tagging on tweets to determine potential candidates for mapping. Unfortunately, these tools are not perfect. Hence, probabilistic analysis techniques are needed that properly quantify uncertainty and inaccuracy.

Furthermore, an ontology evolves over time due to evolving application requirements and acquisition of new knowledge that is of interest. It would be time-consuming and slow if the only way to expand an ontology were through a manual effort. An alternative would be to look for semi-automatic approaches for *ontology evolution*. Accurately identifying ontological concepts and relations from micro-posts is a hard task. A combination of natural language processing, statistical and machine learning techniques may be needed.

C. *Probabilistic Situational Understanding*

Deductive reasoning is typically used to derive logical conclusions from a crisp set of premises. However, given the uncertainty inherent in tweets, this approach is not directly applicable to Twitter content. One should measure the performance of reasoning (standard reasoning over ontologies is deductive in nature) and depending on the results, consider alternatives that accommodate uncertainty, such as probabilistic ontology modeling.

Our present architecture uses simplified solutions to the above problems, essentially offering placeholders that allow for future extensions. The approach allows us to investigate new algorithms in the context of a complete and functional system for multimodal analysis of social sensing feeds. Hence, the practical impact of any proposed extensions can be assessed using real applications and data.

V. Related Work

Next, we discuss some of the recent work related to the aforementioned research challenges.

A. *Event Classification and Tracking using Social Data*

Event detection from user-generated social media content (e.g., Twitter feeds) is a widely studied topic. Events are detected primarily by identifying changes in the frequency distribution of usage of hashtags or keywords (i.e., volume of usage) using a combination of machine learning techniques [17]. In Walther [21], spatially localized events such as house fires or parties are detected by first setting up spatial filters to identify clusters of tweets and then using a combination of topic and semantic analyses to identify events. In contrast, Twitcident [2] focuses on monitoring events; here, it is assumed that events to be monitored are known a priori and Twitter data are then mined to monitor those events. Some work, such as [15], focused on the classification and tracking of such events. While many solutions rely on pre-computed vocabularies for such classification, practical systems should develop and update their vocabulary autonomously, which remains an important research challenge.

B. *Multimodal Social Sensing*

Recently, work in social sensing has emerged that addresses reliability challenges, such as establishing the credibility of observations and the trustworthiness of sources. In [23], the authors take a maximum-likelihood approach to addressing both quantities simultaneously. Their work focuses only on textual data (e.g., Twitter feed). In a multimodal setting, we may also consider image content posted by users of social media platforms, such as Instagram, to corroborate text. Another related social sensing problem is the localization of detected events. Prior work [7] jointly localized events and sources (the users who generated the content) based on tweet content and user profile information. Integrtaing additional modalities of data, such as images and video, can improve localization accuracy. A challenge is to enable finer-grained localization of sub-events, in the presence of location uncertainty, unknown orientation of points of view and unknown use mobility patterns [10, 11].

C. *Semantic Web in Social Sensing*

A number of papers evaluated the possibility of harnessing the power of semantic web for solving problems in the social data mining domain. In [8], the author provides a comprehensive narration of how collective intelligence emerges from the exploitation of semantic web content at scale. Further, in [13], the authors describe an architecture for a distributed semantic web microblogging system. More recent work leverages semantic web tools for sentiment analysis [16], and for enriched user profiling [1]. To attain a situational understanding, semantic web content could be used to reconstruct events from disparate data sources. As the system evolves, we envision an evolution of the ontology and probabilistic event-related knowledge. Automated techniques that enable this evolution are an active research area at the present time.

D. *Anomaly Determination*

Outlier detection is a well researched topic across diverse research areas and application domains. In [3], a comprehensive survey is presented of the key techniques used in anomaly detection. For example, in data mining, such techniques were used to detect popular or trending topics [20] and identify outliers [19]. More recently, in [6, 12, 14], outlier detection was used to detect unusual events from changes in mobility traces of individuals and knowledge of their inferred social ties. The recognition that an observation is anomalous, however, goes beyond outlier detection. It requires an understanding of context in order to properly interpret the observed outliers. For example, a gathering for "Eid El-Fitr" prayer is not an anomalous event, although may seem like a statistical outlier. An interesting question is therefore to determine whether or not an observed outlier constitutes an anomaly, as opposed to a rarely-occurring normal pattern. This is especially challenging to automate when local context is not accurately known.

VI. Conclusion

In this paper, we articulated a vision for real-time, situational understanding through the processing of multi-modal social data. We presented the architecture and discussed the key challenges to be addressed in realizing this vision. The corroboration of information from multiple sources and the use of semantic web technologies would significantly enhance the understanding of physical events as they unfold. For a true situation understanding, a key take-away point is the need for collaboration among multiple disciplines including sensor networks, data mining, natural language processing and semantic web. Such a collaboration offers interesting avenues and directions for future work.

VII. Acknowledgments

This material is based on research sponsored in part by the Air Force Research Laboratory, under agreement number FA2386-141-002, and by the Army Research Laboratory under agreement numbers W911NF-06-3-0001, W911NF-06-3-0002, and W911NF09-2-0053, as well as partially supported by the Singapore National Research Foundation under its International Research Centre@Singapore Funding initiative. This research was also supported in part by an appointment to the Postgraduate Research Participation Program at the U.S. Air Force Research Laboratory, 711th Human Performance Wing, administered by the Oak Ridge Institute for Science and Education, through an interagency agreement between the U.S. Department of Energy and the U.S. AFRL. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Army Research Laboratory, the Air Force Research Laboratory or the U.S. Government.

References

[1] F. Abel, Q. Gao, G.-J. Houben, and K. Tao. Semantic enrichment of twitter posts for user profile construction on the social web. In *The Semanic Web: Research and Applications*, pages 375–389. Springer, 2011.

[2] F. Abel, C. Hauff, G.-J. Houben, R. Stronkman, and K. Tao. Semantics + Filtering + Search = Twitcident. Exploring Information in Social Web Streams. In *In Proc. of HT'12*.

[3] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Comput. Surv.*, 41(3):15:1–15:58, July 2009.

[4] R. Davis, H. E. Shrobe, and P. Szolovits. What Is a Knowledge Representation? *AI Magazine*, 14(1):17–33, 1993.

[5] A. Frome, G. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, and T. Mikolov. Devise: A deep visual-semantic embedding model. In *Advances In Neural Information Processing Systems, NIPS*, 2013.

[6] Z. Fu, W. Hu, and T. Tan. Similarity based vehicle trajectory clustering and anomaly detection. In *IEEE International Conference on Image Processing (ICIP)*, volume 2, pages II–602. IEEE, 2005.

[7] P. Giridhar, S. Wang, T. F. Abdelzaher, J. George, L. M. Kaplan, and R. K. Ganti. Joint localization of events and sources in social networks. In *2015 International Conference on Distributed Computing in Sensor Systems, DCOSS 2015, Fortaleza, Brazil, June 10-12, 2015*, pages 179–188, 2015.

[8] T. Gruber. Where the social web meets the semantic web. In *The Semantic Web-ISWC 2006*, pages 994–994. Springer, 2006.

[9] P. Hitzler, M. Krötzsch, B. Parsia, P. F. Patel-Schneider, and S. Rudolph. OWL 2 Primer. http://www.w3.org/TR/owl2-primer, 2012.

[10] P. Jain, J. Manweiler, A. Acharya, and K. Beaty. Focus: Clustering crowdsourced videos by line-of-sight. In *Sensys*, pages 8:1–8:14, New York, NY, USA, 2013. ACM.

[11] P. Jain, J. Manweiler, and R. R. Choudhury. Overlay: Practical mobile augmented reality. In *MobiSys*, pages 331–344, 2015.

[12] K. Jayarajah, A. Misra, X.-W. Ruan, and E.-P. Lim. Event detection: Exploiting socio-physical interactions in physical spaces. In *To appear in 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*.

[13] A. Passant, T. Hastrup, U. Bojars, and J. Breslin. Microblogging: A semantic web and distributed approach. 2008.

[14] C. Piciarelli, C. Micheloni, and G. Foresti. Trajectory-based anomalous event detection. *Circuits and Systems for Video Technology, IEEE Transactions on*, 18(11):1544–1554, Nov 2008.

[15] A. Ritter, Mausam, O. Etzioni, and S. Clark. Open domain event extraction from twitter. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '12, pages 1104–1112, New York, NY, USA, 2012. ACM.

[16] H. Saif, Y. He, and H. Alani. Semantic sentiment analysis of twitter. In *The Semantic Web–ISWC 2012*, pages 508–524. Springer, 2012.

[17] T. Sakaki, M. Okazaki, and Y. Matsuo. Earthquake shakes twitter users: Real-time event detection by social sensors. In *Proceedings of the 19th International Conference on World Wide Web*, WWW '10, pages 851–860, New York, NY, USA, 2010. ACM.

[18] G. Schreiber and Y. Raimond. RDF Primer. http://www.w3.org/TR/rdf11-primer/, 2014.

[19] T. Takahashi, R. Tomioka, and K. Yamanishi. Discovering emerging topics in social streams via link anomaly detection. In *11th IEEE International Conference on Data Mining (ICDM)*, pages 1230–1235. IEEE, 2011.

[20] D. Thom, H. Bosch, S. Koch, M. Wörner, and T. Ertl. Spatiotemporal anomaly detection through visual analysis of geolocated twitter messages. In *Pacific visualization symposium (PacificVis), 2012 IEEE*, pages 41–48. IEEE, 2012.

[21] M. Walther and M. Kaisser. Geo-spatial event detection in the twitter stream. In *In Proc. of ECIR'13*.

[22] D. Wang, T. Abdelzaher, and L. Kaplan. *Social Sensing: Building Reliable Systems on Unreliable Data*. Morgan Kaufmann, 1st edition, 2015.

[23] D. Wang, L. Kaplan, H. Le, and T. Abdelzaher. On truth discovery in social sensing: A maximum likelihood estimation approach. In *Proceedings of the 11th International Conference on Information Processing in Sensor Networks*, pages 233–244. ACM, 2012.