

# Constructing Synthetic Social Media Stimuli for an Emergency Preparedness Functional Exercise

**Andrew J. Hampton**

University of Memphis  
jhmpton8@memphis.edu

**Shreyansh Bhatt**

Wright State University  
shreyansh@knoesis.org

**Alan Smith**

Wright State University  
alan@knoesis.org

**Jeremy Brunn**

Wright State University  
jeremy@knoesis.org

**Hemant Purohit**

George Mason University  
hpurohit@gmu.edu

**Valerie L. Shalin**

Wright State University  
valerie@knoesis.org

**John M. Flach**

Wright State University  
john.flach@wright.edu

**Amit P. Sheth**

Wright State University  
amit@knoesis.org

## ABSTRACT

This paper details the creation of a massive (over 32,000 messages) artificially constructed ‘Twitter’ microblog stream for a regional emergency preparedness functional exercise. By combining microblog conversion, manual production, and a control set, we created a web-based information stream providing valid, misleading, and irrelevant information to public information officers (PIOs) representing hospitals, fire departments, the local Red Cross, and city and county government officials. Addressing the challenges in constructing this corpus constitutes an important step in providing experimental evidence that complements observational study, necessary for designing effective social media tools for the emergency response setting. Preliminary results in the context of an emergency preparedness exercise suggest how social media can participate in the work practice of a PIO concerning the assessment of the disaster and the dissemination of information within the emergency response organization and to the public.

## Keywords

Social media, emergency preparedness, synthetic microblog corpus, disaster response training.

## INTRODUCTION & RELATED WORK

Social media has the potential to assist in the management of regional disasters as observed during several events in the last decade (e.g., a terrorist bombing or a natural event such as a tornado, forest fire, hurricane, or earthquake) (Imran et al., 2015; Simon et al., 2015; Purohit et al., 2014; Imran et al., 2013; Card et al., 2013; Cassa et al., 2013; Blanchard et al., 2012; Cameron et al., 2012; Sarcevic et al., 2012; Vieweg, et al., 2010; Starbird and

Stamberger, 2010; Palen et al., 2010; Palen and Liu, 2007). With respect to assessing the scope of disaster, social media provide a platform for citizen sensing (Sheth, 2009) that may inform emergency command centers about ground events (Vieweg et al., 2010), the changing availability and need for resources (Purohit et al., 2014; Sarcevic et al., 2012), and changing public perception of events (Cassa et al., 2013). With respect to the dissemination of information, social media can also be a means for communicating with the larger public (e.g., to alert the public to evolving risks; to instruct the public about appropriate responses and available resources; to correct rumors; to solicit or disseminate information, or to match needs with offers of resources) (Cassa et al., 2013; DHS S&T Virtual Social Media Working Group and First Responders Group, 2014; Purohit et al., 2013). Thus, social media based exchange between the informal (citizen) and formal (professional) communities may facilitate effective crisis management.

Despite this potential for citizen roles, and despite recent initiatives by the formal response community such as the U.S. Federal Emergency Management Agency (FEMA, 2012a), few experimental studies address the utility of social media for formal response organizations, especially for assisting preparedness. This paper will report on our experiences in developing a social media component for a functional training exercise for a local emergency in a medium-size city in the Midwestern United States. It involved hospitals, county public health departments, city government, fire departments, police departments, and the regional Red Cross. Certainly, any effort to deploy a social media tool in a high consequence setting requires extensive testing, ideally conducted with the fidelity of a training exercise (Hughes, 2014). However, the examination of social media usage in a training exercise is more than practical necessity. While we are strong proponents of naturalistic observation in a real work setting, exclusively observational study presents several threats to scientific study that are mitigated in an exercise. For example, because an exercise is designed, we had access to ground truth. Second, deployment during an exercise allows us to control presentation parameters that could interact with the usage of social media, and/or are easily remediated with specialized technology. For example, we were able to control the message stream, and therefore knew the distribution of its contents and the rate of presentation. Known threats to usability such as data overload and low signal-to-noise ratio can be addressed with technology (e.g., Twitris (Sheth et al., 2014)), and therefore, can be reduced while we examine best-case utility. However, the study of behavior in a controlled setting is only as good as the design of the setting. The focus of this paper is on the design of the social media stream synchronized with designed exercise content. We also describe the standalone tool to control the distribution of messages to exercise participants and the way social media contributed to the assessment of the disaster and the dissemination of information, both within the responding organizations and between the responding organizations and the public.

## EXERCISE DESIGN

The seven distinct types of emergency management exercises range from seminar to full-scale simulation (FEMA, 2012b). We conducted the current work in the context of a ‘functional exercise.’ This is the second most complex (and realistic) operations-based exercise just short of the full-scale exercise. Functional exercises specifically aim to assess the current response plans and to evaluate regional decision-making processes associated with implementing the plans. Thus, the local (e.g., hospital information centers, fire departments) and region-wide (e.g., Red Cross)

response agencies actively participated. However, the actual deployment of ‘boots on the ground’ (e.g., doctors, nurses, fire units, etc.) was simulated.

The scenario included two separate radiological attacks by a terrorist group. The first attack involved low level radiation distributed through the ventilation system at a large concert for children on Friday evening (the night before the exercise). The second attack was a radiological explosion (dirty bomb) at a large international convention on Saturday morning. This attack involved hundreds of injuries/fatalities in the explosion as well as widespread exposure to potentially lethal levels of radiation. Designed to encompass roughly three and a half hours, the scenario began early Saturday morning when radiation detection alarms go off at two different local hospitals (approximately 7:55 a.m.). The people triggering the “portal alarms” had attended the Friday evening concert and were coming to emergency rooms due to illness or injuries that were unrelated to the radiation events. Standard protocol is for the patient who sets off these alarms to be escorted from the emergency room and returned to the hospital through an alternative route to avoid potential contamination of the emergency room. This high degree of uncertainty about the source of the radiation tested coordination across hospitals. Shortly thereafter (approx. 8:10 a.m.) the dirty bomb explodes, testing medical surge and mass casualty response. As the scenario created uncertainty about the sources of radiation and the potential lethality of the radiation by design, a clear

potential for confusion and public panic resulted. Thus, the exercise represented an ideal context for using social media to simulate the public confusion in a real disaster.

## SYNTHETIC MICROBLOG STREAM CONTENT

We employed a Twitter-like system of microblogs. Each message entry followed the typical format of a 140-character (maximum) microblog, and included an indication of the sender, message content, and hashtags to group messages thematically. Each microblog also had a visibility attribute, set to *high*, *medium* or *low* by the research team. This controlled the proliferation of microblogs throughout the exercise, and reflected either sender influence or context relevance. We designed the software to randomly select usernames from an extensive list of real Twitter users for each artificially generated message. Users could click these names in the interface that would then hyperlink to corresponding Twitter profiles in a new window. For the actual scenario participant, the link directed to his or her organizational profile.

The microblog stream of tweets consisted of three distinct sources:

- 1) *Background set*: the real-time stream of microblogs unrelated to the event by actual Twitter users in the region (background noise);
- 2) *Authoritative set*: those entered at the user interface by PIOs and simulation coordinators, and their ‘retweets’ (forwarded tweets) generated by the system; and 3) *Constructed set*:
  - 3a) *MSEL specific constructions*: those generated by the system as required for the MSEL (Master Scenario Event List);
  - 3b) *Generic emergency related constructions*: specified by parameterized microblog templates for categories such as ‘Angry Rant’.

Table 1 lists the semantic categories that we used to generate specific microblogs. These categories reflect both content developed in the exercise planning process and general content reflecting the analysis of previous actual emergencies. In particular, we considered two major events of 2013: the Boston Marathon bombing (Cassa et al., 2013) and the Westgate Mall attack (Card et al., 2013). Prior studies of these events provide examples of real microblogs on Twitter and categories of content discussed in the aftermath of these events. While both incidents were important in shaping our understanding of message categories, the Boston Marathon bombing provided a more directly compatible dataset, as our simulation scenario also involved the bombing of a metropolitan American city during a large gathering of international participants. We sampled a set of 2,567 tweets originating from the state of Massachusetts in about 2.5 hours after the bombing took place, from the Twitter data set collected via Twitris (Sheth et al. 2014). Twitris constantly collected the filtered stream of English language tweets from the Twitter Streaming API for event-related keywords and hashtags (e.g., “#bostonbomb”, “boston bomb”, “boston marathon”).

**Table 1. Semantic Categories for Scenario-Relevant Microblogs. Notation Reflects the Distribution in our Corpus**

Angry rant T=149, L=85, M=64	Distant Observer T= 426, H=426	Panic over exposure T=18, L=18
Appeal for help T=71, M=71	Fear for children * 113 T=113, H=62, M=32, L=19	Parent who set off alarm T=131, L=131
Breaking news – explosion T=585, H=426, M=73, L=86	General discussion T=492, H=281, M=141, L=70	Prayer T=93, L=93
Breaking news - status update T=72, M=72	Ham operators with nowhere to go T=88, H=40, M=48	Public reaction T=31, L=31
Call for calm/patience T=63, H=43, M=20	Informational T=190, H=190	Revision of all-clear T=120, H=60, M=60
Call for help T=40, M=40	Injured T=114, M=71, L=43	RRR TWEET * 182 T=180, L=182
Caution and Advice T=78, M=60, L=18	JohnQPublic T=140, H=140	Rumor/ False information T=130, M=130
Confusing public reaction on reports T=50, L=50	Media help resources T=25, M=25	Status update - radiation T=905, H=905
Confusion of Hara/Hobart T=96, H=64, M=32	Media report T=175, H=122, M=53	Sympathy T=40, M=40

Confusion of TC-99(m) T=30, M=30	Non-immediate witness T=255, H=141, M=71, L=43	Uninjured present T=43, L=43
Correct information/treatment T=66, M=66	Non-immediate witness, uninformed T=71, H=141, M=71, L=43	Uninjured, injured friend T=43, L=43
Corroboration * 213 T=213, M=168, L=50	Observers in ER waiting room T= 514, M=320, L=194	What do I do? T=183, H=62, M=102, L=18
Criticism T=80, M=80	Offer to help T=58, M=58	Where to go? T=144, H=60, M=66, L=18
Dayton Region T=828, M=828	Official announcement T=240, H=240	Worried about exposure T=505, H=161, M=245, L=99

Note. T = Total Tweets, H=High visibility, M=Medium visibility, L=Low visibility where T=H+M+L.

### Background Set

To simulate the background noise (microblogs unrelated to the emergency events) we collected tweets from the local (Dayton) region over a four-day window before the exercise day that included no particularly noteworthy events (i.e., no major elections, sporting events, conferences, etc.). We used the Twitter Streaming API's public "statuses / filter" (Twitter Inc., 2014) method to obtain tweets from the specific region of the exercise using a bounding box defined by a pair of latitude and longitude for the two points on the geographical map of the region: south-west and north-east. Visibility for the background set was low. The background set comprised approximately 76% of the microblog stream.

### Authoritative Set

We also had the capability to generate additional microblogs in real time during the exercise, manually specifying the relative visibility of each message as high, medium, or low, corresponding to the visibility levels we assigned to each PIO account based on ranking of their real-world followerships on Twitter. These messages appeared in the unfiltered social media feed without distinction and were 'retweeted' by ghost accounts, the frequency and duration of which determined by their relative visibility setting, set at the time of creation based on a subjective conjecture about how popular a post of that type would be.

### Constructed Set: MSEL Specific Constructions

Scenario specific microblogs included the following examples.

Example 1: patient at [*hospital*] just kicked out after testing positive for radiation... WTF!?!

Example 2: They just kicked some guy out of the ER for radiation poisoning #HulkingOut

Example 3: Do hospitals not treat radiation poisoning? Am I in trouble?

These microblogs simulated witnesses in the emergency room, who sent these when someone who had been at the concert on the previous evening tripped the radiation detection alarms. Because the witnesses are not privy to hospital protocol for rerouting such cases, they infer that a patient is being denied service at that hospital due to radiation. Also included (see Example 2) is a hashtag that ran throughout the simulation at various points. This one demonstrates contradictory levity that appears in real tweet streams. These examples deliberately indicate a lack of understanding on the part of the public, the correction of which falls within the bailiwick of public information officers (PIOs). Visibility for these messages depended upon designed ground truth. True events were set to high, and others as required by the exercise for the relevant semantic category of the event (see Table 1).

### Constructed Set: Generic Emergency-Related Constructions

In addition to specific microblogs that were tied to specific scenario events, we added microblogs based on observations of Twitter streams taken from actual emergency events. Initially, we attempted to directly transform observed tweets to the exercise scenario by using string-based transformation rules. Obvious examples (e.g., #bostonbombing became #daytonbombing) presented little trouble, but the depth of contextualization made a complete transformation both prohibitively difficult and occasionally erroneous. For example, discussions of the marathon could have related terms like runners and mile markers (and permutations). We concluded that the extensive details of the context precluded this approach. Instead of this direct transformation, we opted for a thematic translation of the events. Looking at the messages and the type of categories observed in prior studies

(Cassa et al., 2013; Card et al., 2013), we created templates of microblogs. Some we made by transforming items within tweets from Boston (see Example 4). For others, we adapted themes to the specific context of our simulation (see Example 5). We then varied the options in parentheses to create a set of functionally similar but identifiably distinct microblogs. In Example 5, the two alternatives in two places (denoted by forward slashes and parentheses, respectively) yielded four possibilities to create a computationally inexpensive diversity:

Example 4: RT @the123abc: It should be noted that if you suffer from #PTSD, limit your exposure to the (**#daytonbomb** / **#radiation**) coverage. Social media can b overwhelming.

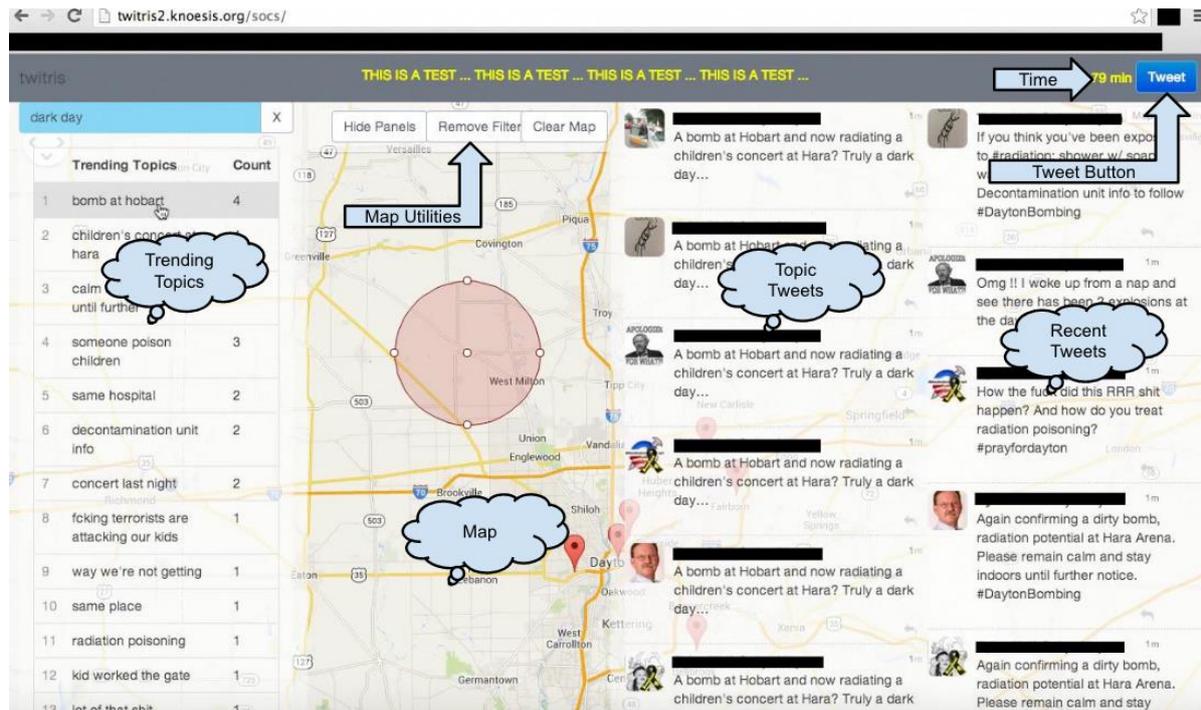
Example 5: (**#bananasplits** / **#hobartarena**) (**We enjoyed with kids** / **kids loved**) the concert at Hobart last night

Visibility depended upon the relevance of the semantic category of microblog content for the MSEL timeline. Because we were also part of the exercise design team and had access to ground truth, we had knowledge of relevance. In Example 4, we set visibility to ‘medium’ given that it was related to the “caution and advice” category, and inherited from the transformation of a Boston bombing related tweet set.

## INTEGRATION OF MICROBLOG CONTENTS SOURCES

We integrated the scenario-related microblogs and background noise to create a corpus of over 32,000 microblogs. We then synchronized the timing of scenario-related microblogs with the MSEL for the exercise. During the exercise the microblogs streamed at a rate of approximately two per second via the interface shown in Figure 1.

The scheduling of the frequency and volume of microblog contents was designed to reflect the specified visibility level for individual microblogs. The fact that the number of microblogs from the background set and the authoritative set and subsequent ‘retweets’ was unknown prior to the start of the simulation presented an integration challenge. To address this challenge, our primary scheduling heuristic made use of the values generated by a realtime trending topics function from Twitris system (Sheth et al., 2014) in the overall data stream at each timestamp ‘t’, throughout the simulation. We considered the count of microblogs mentioning the topmost trending topic (within the current sliding window of past K minutes the timestamp ‘t’) to be the current baseline for high visibility, from which the values for medium and low visibility were proportionally derived. We defined medium visibility as one half of the value of high visibility, and low visibility was one third of value for high visibility. This final count was used as the volume of microblogs scheduled to appear at random intervals throughout the next five-minute window, starting from the entry time of the initial template microblogs from the Constructed Set. These fractional values gave us the desired effect of topics mentioned in a high visibility microblog template or PIO user microblog, to appear at or near the top of a list of the top 20 trending topics. Those from a medium visibility microblog template probably (but not necessarily) appeared somewhere within the list, and those from a low visibility microblog template probably not (but still possibly) appearing in the list.



**Figure 1. Snapshot of Twitris-Based PIO Interface to View the Simulated Microblog Stream during the Exercise.**

## INSIGHTS FROM EXERCISE DEPLOYMENT

We reviewed the screen capture of one active and visible PIO to examine a detailed account of his use of social media. Evidence for engagement occurs in the form of intentional cursor movement, associated with a change in the appearance of the display. This includes selecting trending topics, scrolling the microblog stream to review previous content, and highlighting microblog stream contents. It also includes using the cursor to point at individual microblogs, presumably to keep place while reading. Throughout the exercise, and despite some interface bugs, the recorded PIO was nearly completely engaged in the social media tool with more than three cursor displacements per minute. This metric underestimates engagement, as we did not count multiple movements within the same episode separately (e.g., scrolling the microblog stream back and forth several times without stopping). We note the exploitation of several different functions of the PIO interface in the following sections.

### Information Monitoring

The PIO identified microblogs in the unfiltered stream concerning radiation poisoning less than a minute after the first related microblogs appeared. The PIO also identified microblogs in the unfiltered list concerning the dirty bomb, approximately 1'30" after the first related microblog appeared, and almost a minute after the topic first appeared midway in trending topics. Once he detected the event in the unfiltered stream he returned to trending topics and selected a related topic at the top of the list. Both episodes illustrate a reliance on the unfiltered stream for initial detection, and suggest concern for the reliability of the filtering function. Comments in the followup survey include his interest in knowing more about how the filtering function works. Nevertheless, the PIO used trending topics to identify the existence of a video of the bombing from the trending topics. He later investigated the hashtag identifier of the potentially responsible group, which did not appear in trending topics. He also relied on trending topics to detect public concern once the name of the radioactive agent was released. Late in the session, once focused on a citizen microblog, the PIO pursued the URL link it contained.

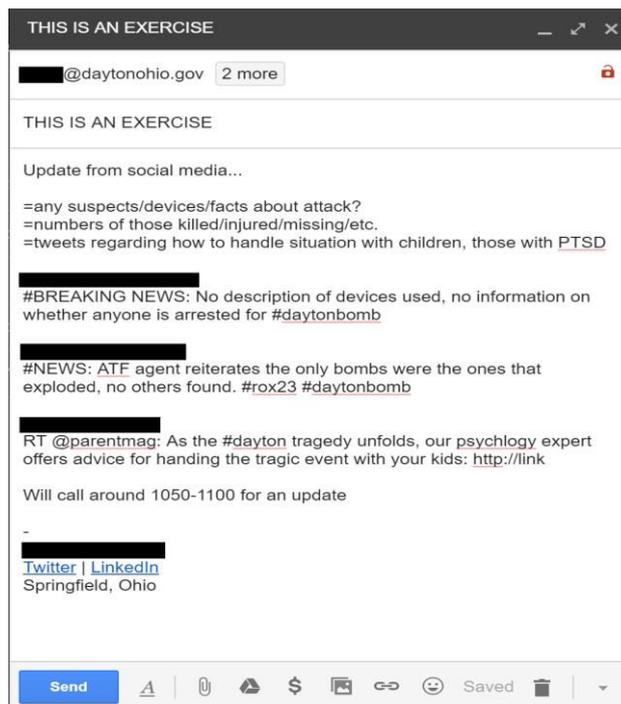


Figure 2. Email Written by the Central PIO to Fire Department Officials Requesting Information

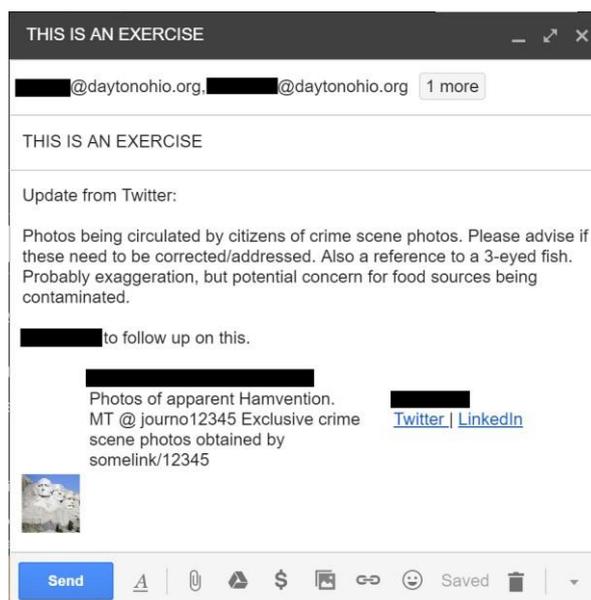


Figure 3. Email Written by the Central PIO to Fire Department Officials Offering Links to Pictures

### Information Distribution

The central PIO functioned as a filter on the social media stream, distributing contents via microblogging, e-mail and telephone. The persisting microblogs allowed the PIO to copy microblogs and forward them via e-mail to interested parties including the fire department and the police (see Figure 2). He used this to identify public perception and even provide crime scene photos (see Figure 3). The separate trending topic microblog list allowed the PIO to copy several examples in one step, as opposed to copying individual examples from the more rapidly scrolling unfiltered stream. Nevertheless, no member of the formal response community responded to the PIOs bulletins using e-mail. The PIO also made numerous calls related to the Twitter stream to the fire command center. The function of these calls was typically to check the validity of information gleaned from the social media feed,

or at least to make the fire department aware of trending topics he felt merited attention or followup. An advantage of the telephone medium is a confirmation of receipt, absent in the one-way email attempts noted above.

### Information source

The tool provides two functions for examining the source of information: the map and the link to user profiles. Early on, before the dirty bomb and with a preponderance of noise, the PIO manipulated the map field of view, and frequently cleared the push pins. He also used the user profiles to examine the background of the sources before sending microblogs and evidence via e-mail.

### CONCLUSIONS

The primary technical contribution of this work is the creation of a synthetic social media stream, fulfilling two requirements of applied research. First, the synthetic stream embodies an explicit, testable model of social media content. Second, the synthetic environment supports response preparation activities. Third, a synthetic stream (with known properties) is required to support the controlled testing of social media tools for emergency response. This allows researchers to separate out properties of the scenario, the social media stream, and the tool in assessing behavioral responses. Compared to the natural environment, training exercises offer control over the difference between critical variable (signal) and extraneous, uncontrolled confounds (noise). In addition, exercises make possible the recording of data and context relatively easily, as well as access to an a priori ground truth. With a relatively modest social media tool we have demonstrated a capability to exploit social media in emergency response. The full value requires continuing investigation of social media search heuristics (e.g., Hampton & Shalin, in press) to filter the content of social media. When such information processing tools are available, they will change work practice regarding both the gathering and dissemination of information critical to effective emergency response.

### ACKNOWLEDGEMENTS

We acknowledge partial support from the National Science Foundation (NSF) awards: SoCS IIS-1111182: Social Media Enhanced Organizational Sensemaking in Emergency Response; EAR 1520870: Hazard SEES: Social and Physical Sensing Enabled Decision Support for Disaster Management and Response, and IIP 1542911 PFI: AIR - TT: Market Driven Innovations and Scaling up of Twitris- A System for Collective Social Intelligence.

### REFERENCES

- Blanchard, H., Carvin, A., Whitaker, M. E., Fitzgerald, M., Harman, W. and Humphrey, B. (2012) The case for integrating crisis response with social media, White Paper, American Red Cross.
- Cameron, M. A., Power, R., Robinson, B. and Yin, J. (2012, April) Emergency situation awareness from Twitter for crisis management, *Proceedings of the 21st International Conference Companion on World Wide Web*, Seattle, WA.
- Card, B., MacKinnon, J. and Meier, P. (2013) #Westgate tweets: A detailed study in information forensics, *QCRI Report*, Retrieved Jan 18 2017 from <http://hhi.harvard.edu/publications/westgate-tweets-detailed-study-informationforensics>
- Cassa, C. A., Chunara, R., Mandl, K. and Brownstein, J. S. (2013) Twitter as a sentinel in emergency situations: Lessons from the Boston Marathon explosions, *PLoS Currents*, 5.
- DHS S&T Virtual Social Media Working Group and First Responders Group (2014) Using social media for enhanced situational awareness, and decision support, Retrieved from <http://www.firstresponder.gov/Pages/SocialMedia-Study.aspx>
- FEMA (2012a) Community emergency response team basic training participant manual, Federal Emergency Management Agency, June 2012, Retrieved from <https://www.fema.gov/library/viewRecord.do?id=6137>
- FEMA (2012b) EMI curriculum — professional development, Federal Emergency Management Agency, June 2012, Retrieved from <http://www.training.fema.gov/IS/searchISbycurriculum.asp?keywords=professional>
- Hampton, A. J. and Shalin, V. L. (in press) Sentinels of breach: Lexical choice as a measure of urgency in social media, *Human Factors*.

- Hughes, A. L. (2014). Participatory design for the social media needs of emergency public information officers, *Proceedings of the 11<sup>th</sup> International Conference on Information Systems for Crisis Response and Management*, University Park, PA
- Imran, M., Castillo, C., Diaz, F. and Vieweg, S. (2015). Processing social media messages in mass emergency: a survey, *ACM Computing Surveys (CSUR)*, 47, 4, 67.
- Imran, M., Elbassuoni, S. M., Castillo, C., Diaz, F. and Meier, P. (2013). Extracting information nuggets from disaster-related messages in social media, *Proceedings of the 10th International Conference on Information Systems for Crisis Response and Management*, BadenBaden, Germany.
- Palen, L., Anderson, K. M., Mark, G., Martin, J., Sicker, D., Palmer, M. and Grunwald, D. (2010, April). A vision for technology-mediated support for public participation and assistance in mass emergencies & disasters, *Proceedings of the 2010 ACM-BCS Visions of Computer Science Conference*, 8, Edinburgh, United Kingdom.
- Palen, L. and Liu, S. B. (2007, April) Citizen communications in crisis: Anticipating a future of ICT-supported public participation, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 727736.
- Purohit, H., Castillo, C., Diaz, F., Sheth, A. P. and Meier, P. (2013) Emergency-relief coordination on social media: Automatically matching resource requests and offers, *First Monday*, 19, 1.
- Purohit, H., Hampton, A., Bhatt, S., Shalin, V. L., Sheth, A. P. and Flach, J. M. (2014) Identifying seekers and suppliers in social media communities to support crisis coordination, *Computer Supported Cooperative Work (CSCW)*, 23, 4-6, 513-545.
- Sarcevic, A., Palen, L., White, J., Starbird, K., Bagdouri, M. and Anderson, K. (2012) Beacons of hope in decentralized coordination: Learning from on-the-ground medical twitterers during the 2010 Haiti earthquake, *Proceedings of the ACM CSCW*, 47-56, Seattle, WA.
- Sheth, A. P. (2009) Citizen sensing, social signals, and enriching human experience, *IEEE Internet Computing*, 13, 4.
- Sheth, A. P., Jadhav, A., Kapanipathi, P., Lu, C., Purohit, H., Smith, G. A. and Wang, W. (2014) Twitris: A system for collective social intelligence, *Encyclopedia of Social Network Analysis and Mining*, Springer New York, 22402253.
- Starbird, K. and Stamberger, J. (2010) Tweak the tweet: Leveraging microblogging proliferation with a prescriptive grammar to support citizen reporting, *Proceedings of the 7th International ISCRAM Conference*, Seattle, WA.
- Simon, T., Goldberg, A. and Adini, B. (2015). Socializing in emergencies — A review of the use of social media in emergency situations, *International Journal of Information Management*, 35, 5, 609-619.
- Twitter, Inc. (2014) "POST statuses/filter". Retrieved from <https://dev.twitter.com/streaming/reference/post/statuses/filter>
- Vieweg, S., Hughes, A. L., Starbird, K. and Palen, L. (2010, April) Microblogging during two natural hazards events: what twitter may contribute to situational awareness, *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1079-1088, Atlanta, GA.