Citizen-Sourced Event Reporting with Veracity Assessment

Christodoulos Demetriou
Open University of Cyprus
christodoulos.demetriou1@st.ouc.ac.cy

Loizos Michael
Open University of Cyprus &
CYENS Center of Excellence
loizos@ouc.ac.cy

ABSTRACT

The overwhelming number of events of potential interest to local or national authorities is often in stark contrast to the severely limited resources that authorities can avail for identifying such events. The active engagement of citizens in the identification task is, arguably, a potential solution to this problem, but brings with it a number of other issues that need to be addressed: How are citizens incentivized to participate? How is the veracity of reported events assessed? To what extent can relevant processes be automated? This position paper puts forward an architecture for an integrated platform that couples citizen-sourced reporting with the semi-autonomous use of drones for assessing the veracity of sourced reports, and discusses the multiple facets of such an architecture in relation to typical research questions for autonomous agents and multi-agent systems.

KEYWORDS

Human-Machine Collaboration, Citizen-Sourced Information, Information Aggregation and Validation, Cognitive Offloading

ACM Reference Format:
Christodoulos Demetriou and Loizos Michael. 2021. Citizen-Sourced Event Reporting with Veracity Assessment. In *Proc. 2nd International AAMAS Workshop on Autonomous Agents for Social Good (AASG 2021), London, UK, May 3, 2021*, IFAAMAS, 5 pages.

1 INTRODUCTION

A large number of incidents arise on a daily basis that require handling or intervention by government agencies, ranging from serious natural disasters such as forest fires or flooding, to car accidents or law violations, and to state equipment malfunctions such as broken pipes or road and sign damages. Due to the great volume and diversity of these events, both in their type and origin, it is widely accepted that public agencies responsible for their handling are overburdened and unable not only to investigate all relevant events, but often to even perceive that such events have occurred.

One would expect that technological advances on mobile devices and wearable gadgets would alleviate the problem of event reporting, having effectively transformed people into a network of "mobile wireless sensors" able to reliably detect events and efficiently transmit relevant information to the intended recipients. This prospect has not, however, fully materialized for at least two reasons. First is what is called the "diffusion of responsibility" phenomenon [4] or the "bystander effect": the more people witness an event, the less likely it is for someone to report it. Second is that even when humans do report events, there remain big challenges in gathering, aggregating, and validating that information [13].

Proc. 2nd International AAMAS Workshop on Autonomous Agents for Social Good (AASG 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3, 2021, London, UK.

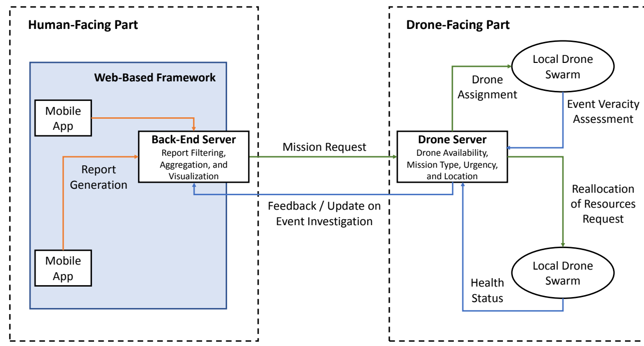

Figure 1: Architecture of the proposed integrated platform.

In this paper we put forward a proposal for the design and development of an integrated platform to alleviate the two aforementioned issues by incentivizing people to report events, and by semi-autonomously aggregating those reports and assessing their veracity. In particular, we focus on the broad case of external open-space events occurring both at land and in the sea, from forest fires to maritime search and rescue, which can be observed by ordinary citizens and can also be aerially monitored by drone teams.

From an abstract high-level point of view, the platform is viewed as a cyber-physical system that tracks and validates information generated from many and unreliable sources. Aligned with its diverse set of goals, our proposed platform comprises several modules: mobile applications in the role of a wireless sensor network supporting the citizen-sourced reporting of events, a back-end server for recording, aggregating, and visualizing the incoming reports and creating missions, a drone server for the resource allocation and management of the drone teams that act as first responders to those events, and strategically-placed local stations of drones that are responsible for the mission execution and event verification.

The processing of information that happens across these modules can be split into two phases, as depicted in Figure 1, roughly corresponding to the human-facing part and the drone-facing part of the platform, dealing, respectively, with the incoming reports and the generated missions. The purpose of the first phase is to incentivize people to generate information at the front end, and to record and aggregate the generated information, at the back end, in order to generate mission requests based on the incoming reports. The purpose of the second phase is to investigate the veracity of information contained in the reported events, while also processing the mission-specific information having to do with all the parameters of the mission operation (e.g., drone availability, number of active missions, location and status of the drones, etc.).

In the sections that follow, we describe the various aspects of the proposed platform and information processing in more detail, and discuss issues that arise from the involvement of humans in the platform, both as participants in decision making and task execution, but also as potential subjects of monitoring by the platform.

2 HUMAN-FACING REPORT MANAGEMENT

Our platform envisions the use of mobile applications, easily deployed on the existing mobile phones of ordinary citizens, as the source of reported events and the interface via which the citizens will communicate with the rest of the platform. The core concept of the design is user-friendliness to motivate people to interact.

The whole process of report generation is expected to be completed within six "click" steps: *(i)* log in or use anonymously, *(ii)* select event type, *(iii)* provide extra information, *(iv)* provide event location, *(v)* upload optional photo, and *(vi)* review and submit.

Given the principled availability of this distributed network of event sensors, two are the key considerations: how to incentivize citizens to download and use, responsibly, the mobile application, and how to aggregate, visualize, and act on the incoming reports.

2.1 Citizen Incentivization

A citizen incentivization scheme is central in the adoption of our platform, as it is what eventually facilitates the transformation of an event perceived by a citizen into an actually-reported event that can be further processed and acted-upon by the platform.

The first problem that needs to be addressed is the information generation by the citizens. Drawing information from the citizen-sourcing literature [1, 7, 8, 22], one can see that there are two major categories of motives, of equal importance, that push people to participate in citizen reporting: those that satisfy the desire for serving / helping other people, and those that originate from the need for self-protection and self-improvement.

To provide the aforementioned motives for the participation of citizens, the development of a sole mobile application for the reporting of the events is not enough; a broader social framework that would embrace / encapsulate the mobile application is needed.

We propose the development of a social framework as a web-based application that will contain the means for people to interact with each other and with the authorities, and to promote and advertise the results of the platform in terms of the safety and well-being of people and the community, in order to satisfy the motives for self-protection and self-improvement and serving / helping others.

To cope with the idiosyncrasies of individual "sensors" in the platform, who might perceive an event differently from others and introduce complexity into the process of data aggregation, we propose a simple partial solution by having pre-determined types of events to select from on the mobile application. These can be chosen from a larger list of such events according to what the local authorities or the particular citizen believe are more pertinent in the citizen's particular context (e.g., location of their residence).

Another key consideration is the feature of user login in the mobile application. On the one hand, this feature could be helpful for keeping track of a user's quality of reports, and, at a later stage, using that as an indicator of a report's veracity. Also, the login feature and the verification of a user's credentials will act as a measure to discourage the creation of false or malicious reports. On the other hand, several citizens may hesitate to enter their credentials in order to use the mobile application to report events, especially if an event is of a criminal nature, e.g., a perceived robbery.

Ultimately, one could support the use of both eponymous and anonymous reports, leaving it to each individual citizen to weight the benefits (e.g., the opportunity to publicly share their contribution for social good) or the perceived risks (e.g., the possibility of being asked to testify in a court of law on a reported crime) from each option on each particular event that they choose to report.

2.2 Back-End Server Module

The role of the back-end server is twofold: to store in databases and manage all relevant information regarding the platform description (number and types of events, user veracity ratings, etc.), and to generate mission requests from the incoming event reports, through a process of aggregation and validation (i.e., veracity filtering).

Although there exists a large amount of work in the literature regarding crisis and emergency informatics and utilization of citizen reporting via social network platforms like Twitter, researchers agree that there is a lack of a solid methodology for information validation [6, 13, 23]. The most prominent method for validation is the use of indicators like user credibility, user proximity to the event, number of tweets by a user, etc., along with the use, in some cases, of machine learning methods for the assessment of the reports [6].

In our particular case, information aggregation includes the grouping of incoming reports of each type into event clusters, based on some specific and predetermined radius that may depend on the type of the event itself. All reports within a cluster are effectively treated as referring to a single event for further processing.

Using strategies of data aggregation from sensor networks [21], an estimation of the probability of an event's veracity will be calculated taking into consideration the number of reports in the cluster, the reporting users' veracity ratings, and the type of the event. Taking into account the expected veracity and the level of urgency in investigating that event, the back-end server decides regarding the further investigation of the event by creating a mission request.

For the database management system, we have to determine the relevant information to be kept for each event: its description, the concerned authority, its level of urgency, the type of mission needed for its verification (e.g., surveillance, simple verification, area monitoring). In case an event is reported eponymously, the user's veracity rating is stored and updated (based on feedback from other users and from the drone verification process).

3 DRONE-FACING MISSION MANAGEMENT

Once a mission request is issued, the second phase of information processing is initiated to assess the veracity of a reported event.

3.1 Drone Server Module

The purpose of the drone server is to distribute the mission requests to the proper local drone stations, and to monitor and visualize all the information regarding active missions and drone status.

Gathering information from the various local drone stations, and combining that information with an incoming mission request, the drone server decides which local drone station will be allocated the particular mission, and which subset of the drones in that local station will be involved. The decision is based, in particular, on the current status of the various drones (e.g., whether they are actively engaged in another mission), the prognosis on their health (e.g., the charge level of their batteries), the location of the event, and the urgency in investigating it compared to other ongoing missions.

More specifically, the process of allocating drones to missions should be able to re-assign drones currently engaged in less critical missions, and to pre-plan the re-allocation of drones between local stations to anticipate future needs in the vicinity of those stations. The priorities of pending (or projected) missions, which depend on their urgency as dynamically updated with the lapse of time since their generation, should be taken into account in these decisions.

On a slightly more technical level, all local drone stations have a radius of action determined by the endurance and the maximum range of the drones they currently host. As the prognostic health management system (PHM) of each station will have access to the availability of each drone, an optimization algorithm will assign to each mission the appropriate swarm of drones. Work on the dynamic vehicle routing problem [14] seems relevant for this task.

3.2 Local Drone Stations

Upon the assignment of a mission to a local drone station, the latter undertakes the task of investigating the associated reported event. We envision these local stations not as being fully autonomous and independent of human-staffed stations, but rather as being part of existing human-staffed stations such as fire department stations, police department stations, coast guard stations, etc.

The human staff of a local station is responsible to monitor the mission evolution, having as their primary task the validation of the corresponding investigated event, and the offering of feedback to the other modules of the platform. If human intervention is found to be necessary for the investigated event, then the staff of the station will be the first human responders for that event.

Two phases need to be considered: the navigation loop, which represents the steps that must be taken in order to navigate the drone to a certain location, e.g., the need for a path planner; and the mission loop, which represents the management of the specific mission, including the sensor monitoring and decision making based on the incoming information flow from the drone sensors [3].

We divide the navigation loop into two parts. First, navigating a drone from its local station to the event location. Here the navigation path follows a simple approach by dividing the shortest distance between the two points into equally-spaced way-points and flying through them to the location of the event. For the obstacle avoidance problem, a local planner will be searching for obstacles in the proximity of the current flying area and it will deviate the drone from its current route until the obstacle is avoided.

When the drone navigates to the event location, it enters the mission loop, which comprises the execution of a search pattern, flying over the event cluster area with a mission-specific search pattern according to the camera coverage problem of the specific mission requirements [18]. This continues until the objective of the mission request is validated, rejected, or aborted in case of drone malfunction or fuel / battery shortage. If the mission is aborted prior to completion, then a new mission request is created.

The drone follows the search pattern until candidate targets are identified either through an image recognition algorithm or directly by the drone operator. Information on the type of evidence that is relevant for each particular event will be made available from the back-end server (e.g., “vehicles” for a car accident event, “people” for an injury event, “smoke” for a forest fire event). When a newly-created report reaches the back-end server, user-uploaded photos will be searched for the specific pieces of evidence for validation purposes, and will also be shared with the drone operator for their consideration and cross-validation with the mission-gathered data.

In effect, the drone operator at each local drone station serves not only as another level of checks-and-balances that the event needs to be investigated, but also as the ultimate assessor on whether the mission-gathered data validates the presence of the investigated event and establishes the need for its handling by a human agent.

Ultimately, the use of drones allows the offloading of some of the investigation workload and cognitive burden that typically falls on humans, and allows them to prioritize the reported events that require their attention. Even in cases where human agents choose to respond to a reported event (e.g., because of its high potential risk) prior to its verification through the drones, the drones are expected to reach the event location in advance of the human responders, and can potentially offer critical information that can be utilized by the human responders upon their arrival (e.g., the drones can identify the location of a drowning incident, so that the human responders know exactly where to concentrate their efforts.

4 HUMANS IN THE DECISION-MAKING LOOP

Our proposed platform embraces the core idea behind human-machine teaming [10]. Rather than seeking a supervisory role for humans in a human-machine interaction, our platform seeks for a synergistic cooperation where human agents and machine agents team up and complement each other’s abilities and characteristics.

On the one hand, machine agents are excellent at performing repetitive tasks with great consistency, but lack performance in terms of improvisation when dealing with unseen and novel situations. On the other hand, human agents have variability in their behavior and performance in repetitive tasks, but stand out when dealing with unpredictable conditions [9]. Bringing machines and humans together in a synergistic and complementary manner, allows hybrid teams of machines and humans to achieve maximum performance over the whole spectrum of operating scenarios.

We suggest, in particular, that human agents will be associated with each module of our proposed integrated platform, symbiotically operating with the rest of the cyber-physical system for shared decision-making and role-switching during mission execution.

A key challenge in this human-machine teaming is the requirement to cope with the lack of ability of humans to quickly adjust to a dynamically changing context, which is inherent in our proposed platform. Typically, the human agent will have the role of an observer, and the machine will be undertaking all decision-making. When disruptions or abnormal conditions occur, however, the human agent will need to immediately change their role to that of a decision-maker or a task-performer [9, 12], including potentially taking control of a drone during its navigation or mission loops.

In order to perform well in such changing contexts, human agents must be fully aware of the conditions of the system, meaning that they have to be sufficiently engaged with the system to be able to identify when to change roles, and to be able to undertake their new role effectively when a role-change does take place. There is a fine balance to strike here. Insufficient engagement with the system or engagement with low arousal activities such as monitoring for low probability events can cause poor human performance. Over-exposure to information, and over-engagement or arousal can cause cognitive overload and lead, again, to poor human performance.

The key question, then, is what information, at what rate, and how it is to be presented to a human agent in order to provide the latter with the necessary knowledge of the system's status, before the human agent engages with the performance of a task [3, 5, 12].

A potential answer is to adopt into the platform features from an adaptive cyber-physical-human (ACPH) framework, which incorporates adaptive user-interfaces and automation levels varying according to the dynamic changes of a human agent's cognitive state, and according to the system and environmental disruptions. ACPH deduces the cognitive state of a human by means of measuring appropriate physiological behaviors and then using machine learning models to predict the human's current cognitive load [2, 9, 15].

Taking into account effectiveness and ergonomy, suitable physiological behaviors could be heart rate, heart rate variability, pupil dilation, and blinking rate. The effectiveness of these indicators has been demonstrated in the literature [2, 15, 16], and increased ergonomy is achieved due to the minimally-intrusive devices used for their measurement, such as wearable smartwatches for heart measurements and cameras or eye-tracking glasses for eye movements; and in this in contrast to the use of electrodes for other popular measurement methods like electroencephalograms (EEG) [2].

A machine learning model, trained on such measurements of the physiological data, can then be used to classify the cognitive state of a human agent [2]. The platform can, subsequently, utilize this classification to adjust the rate that information is presented to a human agent, the modality that is used to communicate that information (e.g., visual, auditory, tactile), and the degree of autonomy of the platform in making decisions and taking actions.

Complementary to the above, human agents might undertake the role of coaches to machine agents towards improving the latter's decision-making capabilities when the latter err [11]. In this context, human agents need to be more actively engaged during the early phase of their teaming with machine agents, and can gradually reduce their engagement level over time as the machine agents become more competent and gain the trust of the human agents.

5 ETHICAL CONSIDERATIONS

Our proposal for the platform put forward in this work is motivated by the goal of promoting social good through the use of technology. In particular, our suggested approach is geared towards closing the gap between authorities and citizens and promoting their cooperation in maintaining a safe and efficient living environment.

The design and implementation of our proposed platform might raise certain ethical considerations in relation to the handling of the personal data of citizens who choose to file eponymous reports, and in relation to the use of drones to monitor public areas.

Having in mind the purpose of the platform, the American Civil Liberties Union suggestions about protecting privacy from aerial surveillance [19], the features of certain prior work [17], and the incentivization of the citizens for participation, we suggest the adoption of the following guidelines in the development and deployment of the platform to help alleviate the raised ethical concerns.

Although personalization is a useful (or necessary) feature in several modules of the platform (e.g., in accessing the mobile and web applications, in measuring an operator's cognitive load, in coaching machines to adopt an operator's decision-making policy), personally-identifiable data need not be kept. Instead, a "soft" login can be used, identifying each individual only by a unique number.

Regarding aerial surveillance, drones will be used for monitoring only when the reported events are such that necessitate their use. No patrolling or preventive flights will be executed, and image or video footage containing information that could lead to the identification of people will not be kept unless to the extent allowed and required by law (e.g., as evidence for a reported criminal activity).

In terms of the platform's web-based application, this can incorporate publicly-available information specifying the policies and procedures regarding the use of the platform, as well as analytics for its deployment and effectiveness. Policies regarding the deployment and operation of the local drone stations could be reviewed and decided jointly by the authorities and the citizens. These policies could specify conditions on which areas can or cannot be monitored, and conditions under which areas actively being monitored need to be reported as such through the mobile or web-based applications.

6 CONCLUDING REMARKS

This paper has put forward a proposal for an integrated platform aimed to promote the more active involvement of citizens in reporting events of interest, identifying several considerations that such a platform needs to address. Not the least of these considerations are mechanisms to incentivize the participation of citizens, and to assess the veracity of the reported events in a semi-autonomous manner by utilizing drones as first responders to the event locations.

Arguably, the design of the platform requires a multi-disciplinary perspective, and its development touches upon a multitude of technologies, spanning game theory, mechanism design, multi-agent systems, machine vision, human-computer interaction, and cognitive systems, just to mention a few. Our own take on the matter is based on adopting an agile design and development methodology, by reusing existing technologies to the extent possible (including algorithms for cognitive swarms of drones [20]), towards a quick first prototype deployment. Having the platform operate in a real-world setting, we then plan to focus our research effort on the synergistic interaction of humans and machines in reducing human errors and offloading the cognitive burden for tactical decision-making, while also maintaining humans in the loop for strategic decision-making.

ACKNOWLEDGMENTS

This work was supported by funding from the European Regional Development Fund and the Government of the Republic of Cyprus through the Research and Innovation Foundation under grant agreement no. INTEGRATED/0918/0032, from the EU's Horizon 2020 Research and Innovation Programme under grant agreement no. 739578, and from the Government of the Republic of Cyprus through the Directorate General for European Programmes, Coordination, and Development.

REFERENCES

[1] G. Abu-Tayeh, O. Neumann, and M. Stuermer. 2018. Exploring the Motives of Citizen Reporting Engagement: Self-Concern and Other-Orientation. *Business & Information Systems Engineering* 60, 3 (June 2018), 215–226.

[2] M. I. Ahmad, I. Keller, D. A. Robb, and K. S. Lohan. 2020. A Framework to Estimate Cognitive Load Using Physiological Data. *Personal and Ubiquitous Computing* 24 (September 2020).

[3] M. L. Cummings. 2015. Operator Interaction with Centralized Versus Decentralized UAV Architectures. In *Handbook of Unmanned Aerial Vehicles*, Kimon P. Valavanis and George J. Vachtsevanos (Eds.). Springer Netherlands, Dordrecht, Netherlands, 977–992.

[4] J. Darley and B. Latané. 1968. Bystander Intervention in Emergencies: Diffusion of Responsibility. *Journal of Personality and Social Psychology* 8 4 (1968), 377–383.

[5] J. Franke, V. Moffitt, T. Spura, E. Alves, and L. Martin. 2005. Inverting the Operator / Vehicle Ratio: Approaches to Next Generation UAV Command and Control.

[6] M. García Lozano, J. Brynielsson, U. Franke, M. Rosell, E. Tjörnhammar, S. Varga, and V. Vlassov. 2020. Veracity Assessment of Online Data. *Decision Support Systems* 129 (February 2020), 113–132.

[7] C. B. Jackson, C. Østerlund, G. Mugar, K. D. Hassman, and K. Crowston. 2015. Motivations for Sustained Participation in Crowdsourcing: Case Studies of Citizen Science on the Role of Talk. In *2015 48th Hawaii International Conference on System Sciences*. IEEE, 1624–1634.

[8] D. Linders. 2012. From e-Government to we-Government: Defining a Typology for Citizen Co-production in the Age of Social Media. *Government Information Quarterly* 29, 4 (October 2012), 446–454.

[9] A. M. Madni and C. C. Madni. 2018. Architectural Framework for Exploring Adaptive Human-Machine Teaming Options in Simulated Dynamic Environments. *Systems* 6, 4 (December 2018), 44.

[10] A. M. Madni, M. Sievers, and C. C. Madni. 2018. Adaptive Cyber-Physical-Human Systems: Exploiting Cognitive Modeling and Machine Learning in the Control Loop. *INSIGHT* 21, 3 (2018), 87–93.

[11] L. Michael. 2019. Machine Coaching. In *IJCAI 2019 Workshop on Explainable Artificial Intelligence (XAI)*. Macao SAR, P.R. China, 80–86.

[12] E. Ordoukhanian and A. M. Madni. 2018. Human-Systems Integration Challenges in Resilient Multi-UAV Operation. In *Advances in Human Factors in Robots and Unmanned Systems (Advances in Intelligent Systems and Computing)*, Jessie Chen (Ed.). Springer International Publishing, Los Angeles, California, USA, 131–138.

[13] L. Palen, S. Vieweg, J. Sutton, S. Liu, and A. Hughes. 2007. Crisis Informatics: Studying Crisis in a Networked World.

[14] M. Pavone and E. Frazzoli. 2010. Dynamic Vehicle Routing with Stochastic Time Constraints. In *2010 IEEE International Conference on Robotics and Automation*. IEEE, 1460–1467.

[15] L. Planke, Y. Lim, A. Gardi, R. Sabatini, T. Kistan, and N. Ezer. 2020. A Cyber-Physical-Human System for One-to-Many UAS Operations: Cognitive Load Analysis. *Sensors* 20 (September 2020), 5467.

[16] D. W. Rowe, J. Sibert, and D. Irwin. 1998. Heart Rate Variability: Indicator of User State as an Aid to Human-Computer Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '98)*. ACM Press/Addison-Wesley Publishing Co., USA, 480–487.

[17] S. Shah, F. Bao, C. Lu, and I. Chen. 2011. CROWDSAFE: Crowd Sourcing of Crime Incidents and Safe Routing on Mobile Devices. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '11)*. Association for Computing Machinery, New York, NY, USA, 521–524.

[18] R. Shakeri, M. A. Al-Garadi, A. Badawy, A. Mohamed, T. Khattab, A. K. Al-Ali, K. A. Harras, and M. Guizani. 2019. Design Challenges of Multi-UAV Systems in Cyber-Physical Applications: A Comprehensive Survey and Future Directions. *IEEE Communications Surveys Tutorials* 21, 4 (2019), 3340–3385.

[19] J. Stanley and C. Crump. 2011. *Protecting Privacy From Aerial Surveillance: Recommendations for Government Use of Drone Aircraft*. Technical Report. American Civil Liberties Union. 15 pages. https://www.aclu.org/report/protecting-privacy-aerial-surveillance-recommendations-government-use-drone-aircraft

[20] G. Voirin and L. Michael. 2020. Cognitive Swarm of UAVs for Search and Rescue. In *13th Cyprus Workshop on Signal Processing and Informatics (CWSPI)*. Nicosia, Cyprus.

[21] X. Wang, A. Walden, M. C. Weigle, and S. Olariu. 2014. Strategies for Sensor Data Aggregation in Support of Emergency Response. In *2014 IEEE Military Communications Conference*. IEEE, 1120–1126.

[22] F. Wijnhoven, M. Ehrenhard, and J. Kuhn. 2015. Open Government Objectives and Participation Motivations. *Government Information Quarterly* 32, 1 (January 2015), 30–42.

[23] J. Zhu, F. Xiong, D. Piao, Y. Liu, and Y. Zhang. 2011. Statistically Modeling the Effectiveness of Disaster Information in Social Media. In *2011 IEEE Global Humanitarian Technology Conference*. IEEE, Seattle, USA, 431–436.