

Twitter911: A Cautionary Tale

Anis Zaman, Nabil Hossain, Henry Kautz

Department of Computer Science
University of Rochester
Rochester, New York 14610 USA

Abstract

Researchers have argued that social media, and in particular, Twitter, can be searched to improve “situational awareness” in emergency situations; that is, to provide objective, actionable, real-time information to first-responders. Prior studies have examined cases of very rare, catastrophic emergencies that took place over many days, such as the aftermath of Hurricane Sandy. We asked instead if Twitter could provide useful information for first-responders on a more regular basis, by conducting an exhaustive analysis of tweets and fire department data for medium-sized county (population 1 million), and for two larger-scale single-day emergencies in New York City. Our results are resoundingly negative: useful tweets were extraordinarily rare or nonexistent. This study provides a cautionary note as to the potential of Twitter and similar platforms for emergency situational awareness.

Social media posts are routinely analyzed to uncover the beliefs, preferences, and behavior of a large segment of the world’s population, and to reveal how opinions and information percolate through social networks. Some researchers have argued that data mining social media can go beyond gathering subjective and idiosyncratic information, but can also yield objective, actionable, real-time data, for applications such as emergency response and ensuring public health. The existence of a few well-publicized cases of users tweeting or blogging reports from war zones or disaster areas reinforces the notion that social media can be used as a kind of distributed sensor network. Does this impression match reality? We conducted two systematic studies to determine whether Twitter data could be used to enhance emergency response. In the first, we analyzed all posts from a medium-sized metropolitan area to find ones related to fires recorded in the area’s 911 emergency report system over the course of six months. In the second study, we exhaustively searched tweets about two disasters in New York City, the collapse of a construction crane in midtown and a gas explosion in the East Village. Both studies came to the same conclusion: there were essentially no timely, first-hand user posts about any of the events. The vast majority of related posts were retweets of posts that originated from emergency response organizations themselves. A few users tweeted images of the emergency situations long after first-responders

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

were already on the scene. While any limited set of studies cannot prove a negative, we believe ours provide a cautionary note about the potential for social media to be useful for gathering, as opposed to disseminating, objective information. We discuss reasons for this, such as the personal nature of social media and methodological limitations of prior work.

Related Work

While the major purpose of social media analytics is for market research, researchers have sought ways to use or understand Twitter as a force for social good. In public health, for example, tweets have been analyzed to track and predict the spread of contagious disease (Brennan, Sadilek, and Kautz 2013) and to locate sources of foodborne illness (Sadilek et al. 2016). Public safety researchers have studied how Twitter can be an effective tool for conveying information to the public. For example, during the 2007 wildfires in Southern California, Twitter was used to reach out to a wide audience with time-critical information about road closures and community evacuations (Hughes and Palen 2009). During Hurricane Sandy in 2014, police and fire services routinely made use of Twitter to communicate with the public (Hughes et al. 2014), and individuals spread the word by retweeting (Kogan, Palen, and Anderson 2015).

Our focus is on the reverse direction of information flow, from the public to first-responders. Vieweg *et al* (2010) coded tweets from individuals affected by the grass fires and floods in Oklahoma, and Verma *et al* (2011) developed statistical methods to automatically code tweets for situational awareness using tweets about floods and earthquakes. Perhaps the strongest claims of the utility of Twitter for situation awareness were made by Yin *et al* (2015), who described information extraction rules for processing tweets from 27 emergency situations.

Monroe County Fire Calls

We began the Twitter911 project with a practical goal: to augment the stream of 911 dispatches from emergency services with real-time Twitter posts originating from the emergency site. Thus, for example, a fire crew on the way to fire could see pictures of the fire and read about where it started, or a police crew responding to a report of street altercation



Figure 1: Example of a tweet categorized as Fire Alarm.

could see pictures and texts that would give them a sense of the size of the crowd they would encounter. Useful tweets would be ones sent during the early stages of the emergency event, rather than after all the emergency crews were at the scene. We wished to exclude tweets that originated from the emergency dispatch services themselves, as well as retweets of such official posts.

We obtained a complete archive of emergency dispatches from October 2014 through March 2015 for a large upstate New York county from Bryx 911 (<https://www.facebook.com/Bryx911/>), a company that provides 911 alerts and situational awareness services to fire, police, and emergency medical organizations. We decided to begin with fire reports, reasoning that fires would be the kind of common emergency event most likely to be seen by many people and thus most likely to be the subjects of tweets. There were 14,951 fire reports in Monroe County from October 1, 2014 to March 31, 2015. Using the Twitter API, we downloaded all tweets that could be geolocated to a named place or GPS coordinate within Monroe County for the same time period.¹

Matching dispatches to tweets could proceed by first identifying fire-related tweets, and then searching for a dispatch for each, or by starting with dispatches, and searching for tweets for each dispatch. If one wished to build a working system that could find fire tweets before the fires were officially reported, it would be necessary to take the first approach. At this stage of research, however, we simple wanted to determine what percentage of fire reports could in principle be associated with twitter posts, and to create a training set of firsthand tweets about fires, which could then be used to train a tweet classifier for finding such posts.

All of the tweets we collected were geotagged, since only by using a geographic criteria could we capture tweets from

¹Up through April 2015, the end of our study, approximately 3.5% of tweets available through the Twitter API included geolocation tags (Weidemann and Swift 2013). This percentage decreased significantly after April 2015, for reasons that Twitter has not revealed.




Figure 2: Breakdown of types of tweets found by searching in a 0.25 mile radius and 30 minute +/- time window of fire dispatches in Monroe County.

our target area, Monroe County. Each fire dispatch also included a geographic tag. We therefore proceeded as follows: For each fire dispatch, define the set of potentially relevant tweets as those tagged within 0.25 miles of the fire and within 30 minutes plus or minus of the dispatch timestamp. We manually inspected each potentially relevant tweet and classified it as one of the following types:

- **Firsthand Encounter:** Original post by a user who is on the scene of the fire, and made while the fire is active; *e.g.:*
You there's a house on fire the next street over what is this
- **Post Fire:** Original post by a user from the scene of the fire, but made after the fire was over.
- **Fire Alarm:** Original tweets by a user about hearing a fire alarm or seeing fire trucks, but not actually witnessing a fire; see Figure 1 for an example.
- **EMS:** Original tweets created by emergency management services or EMS dispatch aggregators such as Sigalert.com about the fire; *e.g.:*
Accident in #Rochester on Lexington Ave Both EB/WB at Mt Read Blvd #traffic http://t.co/0wsFrmqFZG
- **Public Info Sharing:** Retweets of any of the above, as well as tweets from news organizations, such as:
On scene at a fire at the La-Tea-Da Tea Room on Alexander St. #NEWS5
- **Other:** Tweet not related to the fire.

The results of this are summarized in Figure 2. The number of potentially relevant tweets across all dispatches totaled 14,032. 93% were in fact not relevant to any of the fires (Other). The second largest category at 4% were official posts from emergency service organizations (EMS), followed by news reports and retweets (Public Info Sharing) at 2%. At less than 1% were tweets about hearing fire alarms, firsthand encounters, and post fire tweets. There were just 29 tweets that would have been useful in our envisioned application of providing firsthand reports to emergency responders, or if we more generously included tweets that were made after the fire was extinguished, 35 tweets total. In summary, only 35 of the 14,951 fire dispatches, or 0.25% of the



Figure 3: One of the two Firsthand Encounter tweets of the February 6, 2016 New York City crane collapse.

dispatches, could in principle have been augmented with a twitter post.

We experimented with increasing the time window to 60 minutes, and to account for errors in location tags, with increasing the geographic radio to 0.5, 1, 1.5, and 2 miles. None of the changes increased the number of Firsthand Encounter tweets. The marginal usefulness of being able to augment at most a quarter of one percent of 911 fire dispatches discouraged us from pursuing our original goal of building a practical system for day to day use by fire departments and other first-responders.

Still, we wondered if the reason that we found so few relevant tweets was not they did not exist, but that they were not among the 3.5% of tweets that were geotagged at the time of the study. Even if geotagged tweets could not be found using the public Twitter API, it might well be possible for the company Twitter to localize tweets coming from cell phones using metadata that it did not republish, such as the ID of cell tower to which the phone was talking, or the series of IP address through which the packet passed on its way to Twitter. If so, then it would still be possible to a system for first-responders of the kind envisioned in collaboration with Twitter and/or communication company such as AT&T.

To test this hypothesis, we investigated two recent high-profile emergency incidents in New York City, for which we could search the Twitter stream by text and hashtag content, without limiting ourselves to geotagged tweets. If we could not find useful tweets for these incidents, then we could be confident that they did not exist.

NYC Crane Collapse and Gas Explosion

On February 6, 2016, at about 8 o'clock in the morning, a construction crane collapsed in the streets of lower Manhattan, killing one bystander and injuring three. Glenn Zito, an electrician working on the 30th floor of a nearby building, noticed the crane starting to tip, and captured a video of the collapse on his cell phone. He did not post the video on Twitter or any other social media platforms; instead, he sold the

video to NBC and other news organizations, which headlined it in their evening news reports. Thousands of people were in the immediate vicinity of the accident. Did some of them tweet about it?

GNIP is a paid service from Twitter that allows users to search the historic Twitter firehose. We used GNIP to collect all tweets from 7 a.m. to 9 a.m. that day that originated from a mobile phone, and that contained the words “crane” and “collapse” or the hashtag #cranecollapse. This yielded 4,899 tweets, of which 227 were original and 4,672 were retweets. The vast majority of the original tweets were people’s emotional responses to news of the event, for example:

#cranecollapse rip to the man who lost his life and a speedy recovery to the others injured..

Many users also wrote tweets with pointers to online news articles, such as:

Huge construction crane collapses in Manhattan, kills one <http://t.co/3RRrAGIdYI>

Note that by restricting the data to tweets from mobile phones and eliminating retweets, we had already removed official EMS and news tweets and their retweets. We identified 2 tweets as definite firsthand encounters, and 3 others as possible firsthand encounters. The 3 possible firsthand encounters mentioned the crane collapse shortly after it occurred, but did not include enough detail to determine whether the poster learned of it first or secondhand. One of the two firsthand encounters is shown in Figure 3. However, by the time either of the firsthand tweets was made, emergency responders were already on the scene, because they appear in the background. The utility of these tweets for informing first-responders about the situation, therefore, would be low.

The second emergency event we explored was a gas explosion in New York City’s East Village on March 26, 2015, which destroyed a building and tore open a street. The explosion occurred at 3:17 p.m.. We employed GNIP as before, searching for tweets from the time of the event to an hour afterwards that contained the words “village” and either “explosion” or “fire” or corresponding hashtags, and that originated from a mobile platform. From 3:17 to 3:40, there were no such tweets. After 3:40 p.m., long after first-responders were at the site, a few firsthand encounters started to appear; e.g.:

Huge fire in east village as seen from 7/a #nyc #fire #EastVillage <http://t.co/wVx3TMu1iI>

The number of tweets from spectators at the scene or viewing the plume of black smoke from a distance continued increase until tapering off around 5:30 p.m.. Twitter would be a useful source of data for researchers studying the aftermath of the explosion, and how it affected the surrounding community. It would not, however, been useful in providing information to first-responders.

Discussion

Why do our results on the utility of Twitter data for improving situational awareness differ so dramatically with the claims of prior research? First, it is important to note that

there is no disagreement about the utility of Twitter for information sharing from emergency and news organizations to the general public, as demonstrated by the high volume of such posts and retweets in all of our studies. Our results differ with previous studies regarding how useful Twitter data is in providing “eyes and ears on the ground” for first-responders, that is, the reverse flow of information.

We believe there are two major reasons for this difference. One reason is methodological, in that some studies, such as Yin *et al* (2015), began with a relatively small hand-curated set of tweets about a relatively large number of disasters, and did not address the problem of actually finding useful tweets. The second reason arises from the fact that prior studies considered multiday emergencies over large regions, rather than local single day emergencies. In addition to simply providing more opportunities for tweeting, the large scale emergencies *directly impacted* the lives of millions of people. It is reasonable to hypothesize that people are more likely to create an original tweet about something that personally affects them, than about something they merely observe. In Verma *et al* (2011) for example, the most convincing examples of tweets with situational content were about the movements of the users and their loved ones to escape a flood. In other words, Twitter may be best considered to be an introspective sensor network, capturing personal feelings and experiences, rather than a possible source of objective data.

Acknowledgements

We thank Bryx 911 for access to historic first-responder data. This work was supported by NSF award 1319378, NSF NRT award 1449828, the Intel Science and Technology Center for Pervasive Computing, and the New York State Center of Excellence in Data Science.

References

- Brennan, S.; Sadilek, A.; and Kautz, H. 2013. Towards understanding global spread of disease from everyday interpersonal interactions. In *Proceedings of the 23rd International Conference on Artificial Intelligence (IJCAI 2013)*.
- Hughes, A., and Palen, L. 2009. Twitter adoption and use in mass convergence and emergency events. In *Proceedings of the 2009 Information Systems for Crisis Response and Management Conference (ISCRAM 2009)*.
- Hughes, A. L.; St. Denis, L. A. A.; Palen, L.; and Anderson, K. M. 2014. Online public communications by police & fire services during the 2012 Hurricane Sandy. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2014)*, 1505–1514. New York, NY, USA: ACM.
- Kogan, M.; Palen, L.; and Anderson, K. M. 2015. Think local, retweet global: Retweeting by the geographically-vulnerable during Hurricane Sandy. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work (CSCW 2015)*, 981–993. New York, NY, USA: ACM.
- Sadilek, A.; Kautz, H.; DiPrete, L.; Labus, B.; Portman, E.; Teitel, J.; and Silenzio, V. 2016. Deploying nEmesis: Preventing foodborne illness by data mining social media. In *Proceedings of the 28th Annual Conference on Innovative Applications of Artificial Intelligence (IAAI 2016)*.
- Verma, S.; Vieweg, S.; Corvey, W. J.; Palen, L.; Martin, J. H.; Palmer, M.; Schram, A.; and Anderson, K. M. 2011. Natural language processing to the rescue? Extracting “situational awareness” tweets during mass emergency. In *Proceedings of the 5th International Conference on Weblogs and Social Media (ICWSM 2011)*.
- Vieweg, S.; Hughes, A. L.; Starbird, K.; and Palen, L. 2010. Microblogging during two natural hazards events: What Twitter may contribute to situational awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010)*, 1079–1088. New York, NY, USA: ACM.
- Weidemann, C., and Swift, J. 2013. Social media location intelligence: The next privacy battle — an ArcGIS add-in and analysis of geospatial data collected from Twitter.com. *International Journal of Geoinformatics* 9(2).
- Yin, J.; Karimi, S.; Lampert, A.; Cameron, M.; Robinson, B.; and Power, R. 2015. Using social media to enhance emergency situation awareness. In *Proceedings of the 24th International Conference on Artificial Intelligence (IJCAI 2015)*, 4234–4238. AAAI Press.