

Visible Skepticism: Community Vetting after Hurricane Irene

Dharma Dailey

Human Centered Design & Engineering
University of Washington
ddailey@uw.edu

Kate Starbird

Human Centered Design & Engineering
University of Washington
kstarbi@uw.edu

ABSTRACT

Social media enable rapid, peer-to-peer information flow during crisis events, affordances that have both positive and negative consequences. The potential for spreading misinformation is a significant concern. Drawing on an empirical study of information-sharing practices in a crisis-affected community in the Catskill Mountains after Hurricane Irene, this paper describes how an ad hoc group of community members, led by a handful of journalists, employed specific work practices to mitigate misinformation. We illustrate how the group appropriated specific tools and performed *visible skepticism*, among other techniques, to help control the spread of false rumors. These findings suggest implications for the design of tools and the development of best practices for supporting community-led, crowd-powered response efforts during disasters.

Keywords

Crisis informatics; social media; misinformation; information diffusion; rumors; crowdsourcing; journalism

INTRODUCTION: MISINFORMATION AND VETTING ON SOCIAL MEDIA

Recent research has consistently demonstrated social media and other online tools being used during disaster events to share first-hand, eye-witness reports (e.g. Bruns et al., 2012); for improvisational relief efforts (e.g. Sarcevic et al., 2012); and for real-time communication between responders and their affected publics (Denef et al., 2013). Though these tools are clearly providing value and have the potential to improve aspects of response, they also have significant drawbacks.

One issue is the potential for social media to spread misinformation. Some of the same affordances that enable real-time information sharing from the ground also enable misinformation to propagate rapidly—echoing through different social media spaces as re-posts, retweets, and shares. Emergency responders have voiced concerns about this potential, in some cases citing it as contributing to their reluctance to incorporate social media into their work practices (Hughes & Palen, 2012). These concerns have also shown up in media articles after recent events like Hurricane Sandy (e.g. Hill, 2012) and the Boston Marathon Bombings (Madrigal, 2013).

The practices that alternatively either contribute to or discourage the spread of misinformation over social media are only beginning to emerge from current research. Gupta et al. (2013) aim to identify user attributes unique to malicious Twitter accounts. Mendoza et al. (2010) find that the crowd itself challenges misinformation, though corrected information may diffuse differently than the original message (Starbird et al. 2014). The work of identifying, challenging and correcting misinformation as well as coaching others to do the same has been identified as a core task of online volunteer efforts in crises (Starbird & Palen, 2013). Such practices can lead to accurate crowd-derived information (Vieweg et al., 2008).

This research offers a qualitative examination of how an online community responding to a major disaster event adapted available tools and employed specific practices, including *visible skepticism*, a term one community member used to describe a technique she employed to help stem the flow of misinformation.

EVENT BACKGROUND: IRENE'S IMPACT IN THE CATSKILL MOUNTAINS

The Catskill region of upstate New York is a rural, mountainous five-county area approximately 100 miles north of New York City. In 2011, it was struck by two major storm systems, Irene and Lee, which caused record flooding. The two storms were the largest natural disaster in New York state history, and the most expensive prior to Sandy. As floodwaters rose, roads and bridges were washed away, leaving some areas inaccessible and

some people stranded. Information about damage to infrastructure, especially transportation infrastructure, became extremely important. In the days and weeks after the event, locating people, food, shelter, debris drop off points, volunteer opportunities, and donation locations were pressing information needs. Many residents turned to social media to help meet these information needs.

For the northern Catskills, hard-hit by Irene, small local media outlets scaled up to meet the task of providing live information with the support of a groundswell of volunteers. In one salient example, two former print journalists transformed their online regional news operation, the Watershed Post (WP), into a live newswire for the event. The WP was one of two media outlets that distinguished themselves as important real-time information sources for impacted residents. Interestingly, both received regional and national attention for their efforts and both made heavy use of volunteers.

The Watershed Post Hurricane Irene Liveblog

During the flooding event, the Watershed Post deployed and maintained the Hurricane Irene Catskills Liveblog, which became an important information resource for many in the community and a site of collaboration between journalists, community members, and other volunteers.

Liveblogs are event-specific blogs often used to report on sports events or elections. They can be similar to chat rooms, with many participants commenting. The Watershed Post used the CoverItLive liveblog service, which afforded several levels of participation from the crowd. Administrators, or in this case trusted collaborators, could post unmoderated comments, while anyone in the public could submit a comment that would post to the public forum if approved by a moderator. Liveblog administrators could also select Twitter accounts or Twitter threads (via keyword or hashtag searches) to appear automatically in the liveblog stream—advantageous since both Twitter and Facebook are regular venues for WP reporting.

The geographical extent of the event, damage to infrastructure, and official response policy made it difficult for journalists to do their own reporting and fact-checking from the ground as had been done in previous flooding events. Regional reporters responded by collaborating remotely. Encouraged by the WP editors, several local reporters treated the liveblog as a “newswire,” both taking information from it and contributing to it. When flooding began in earnest, the liveblog grew to be dominated by eyewitness accounts from community members, those seeking information and help, and those wishing to volunteer. This swell in participation meant that by the second day of the event, the liveblog needed its own volunteer force to help shape the information space, with some acting as moderators and others reporting and fact checking.

DATA COLLECTION AND ANALYSIS

This paper draws on a multi-sited ethnographic study (see Marcus, 1995) examining the information-sharing practices in crisis-affected communities after Hurricane Irene. It includes participant observation during the event; interviews with 11 individuals who shared information publicly during the event and were subsequently named as relied-upon information sources (via chain sampling); and content analysis of the digital record.

FINDINGS

Employing and Modeling Visible Skepticism

During the flooding event, the liveblog became a primary source of damage reports, help requests, information about evacuations, etc. Because much of this information was coming directly from the public, the WP editors, who were also the administrators of the liveblog, needed a strategy that on one hand encouraged on-the-ground accounts from across the impacted region but that also addressed the potential for rumors and misinformation as WP Editor Lissa Harris describes:

Harris: We knew that we were not going to be able to vet everything. We just knew. Because of the nature of us being two people and there being thousands on the blog sharing things. And because it became very apparent to us that official sources were no better than unofficial sources in terms of accuracy. We got lots of stuff from “official” people that turned out to be wrong. And so, eyewitness accounts were probably where most of our news was coming from. And we didn’t want to shut that down, but ... we were very aware that this is a situation where rumors spread and we didn’t want to be contributing to misinformation.

Though the editors required each guest comment to be approved by a moderator before posting to the liveblog, they were liberal in what they approved. They employed a strategy that intentionally allowed misinformation to rise up to the surface—so it could be addressed publicly.

Harris: *We made a decision early on not to disallow things like that, when something came in that we couldn't verify, but to be **skeptical visibly**. To say, "Ok, how do you know this?" For days we were fielding rumors that the Gilboa dam had failed, and that was not true, and there were major news outlets reporting that it had failed... We decided to let people say it... have them say it and then say, "Ok, no. This is the information that I have and this is not happening."*

Here Harris described how she employed a stance of *visible skepticism*, a technique that was particularly important for dealing with the persistent and widely circulated rumor about a dam failure.

2:17 Comment From JP: Can someone confirm we just got word the Gilboa dam broke. Please god tell me that was a mistake.

2:18 Julia Reischel: JP: The Gilboa dam has not broken. The water behind it is at record levels, however. Schoharie County scanner traffic is all responders urging residents to get to shelters or to higher ground. Can you get to on of the shelters we've listed?

The first comment in this example, a guest comment that contains the false rumor about the Gilboa dam, posted at 2:17pm on August 28. Shortly after it posted, WP Publisher Julia Reischel addressed the rumor, which she and others would do again dozens of times in the ensuing days.

This model of visible skepticism perhaps encouraged other crowd members to participate in the work of vetting information, including this guest who commented on the dam situation at 3:07pm, bringing in news from an official source on another platform (local TV):

3:07 Comment From Mark Warwick: Greene County Administrator live on Fox 5 - Dams did not fail but were breached by high water. All dams are intact.

Repeatedly allowing a rumor to surface and be corrected is different from correcting it once and then blocking it from surfacing again. The latter stops the rumor's spread within the current platform, but the former may do more to quell its spread through other channels, by challenging it as it comes up again and again.

The practice of visible skepticism is one that builds on and extends traditional journalistic practices. Andy Carvin, a forerunner in crowdsourced news reporting, employed a similar strategy in his efforts to report political turmoil during the Arab Spring (Carvin, 2012). In that situation—as the journalists were in this event—Carvin was separated from his sources, and he relied upon the crowd in two ways: both to provide and to verify information. This practice is, in some ways, a response to affordances of linear online tools where records of previous "corrections" disappear, as newly posted information pushes the old out of view.

Vetting Pop-Up Charities

In the aftermath of Irene, as immediate needs shifted to long-term recovery efforts, several community-based volunteer efforts took shape. Some of these groups established an online presence to enhance their visibility, to help them recruit and coordinate with others, and to collect donations. For example, locally organized groups *Neighbors Helping Neighbors Greene County* and *Hurricane Irene Save the Catskills* both created FaceBook pages to help them with their work. Many of these community-situated efforts became known through social media. Greene County Neighbors Helping Neighbors was started by a volunteer from the impacted area:

I don't pretend this is anything more than me and my friends and my family volunteering. The viralness of the [FaceBook] site took off really quickly. Within the first week we had 1000 fans making donations and updates from Canada to New York City. I knew I could leverage social media to help people.

While online fora extend the reach of community-situated charities, such efforts are easy to spoof in an online environment, creating the opportunity for exploitation, a known feature of the disaster relief landscape (Fritz & Mathewson, 1957). To meet the demand for information on local charities, the two WP editors collaborated with Ivan Lajara, Life Editor of the Kingston Daily Freeman to create a crowd-supported system for verifying their authenticity.

Harris: *We fielded hundreds of calls and emails from people in NYC saying I want to give to Catskills Irene relief. We're like, "Ok, what town?" "I don't know," you know? [Laughs] There was no central coordinated thing for the Catskills or for upstate New York, or for whatever. There was nothing on a regional level. So it was all this town-by-town area-by-area, some of them springing up overnight, places that were scrambling to get IIC3 status, and, there weren't a lot of scammers, but there were things we were leery of.*

The aim of the three was to link those willing to help to direct aid taking place at the community level:

Proceedings of the 11th International ISCRAM Conference – University Park, Pennsylvania, USA, May 2014
S.R. Hiltz, M.S. Pfaff, L. Plotnick, and P.C. Shih, eds.

Harris: But from our perspective, the people that we really wanted to support in this were the local organizations. Because ... there's a limited span of time at which the Red Cross has responses and then they pack up and go to the next disaster. And on the ground, there is no Catskills-wide organization to donate to.

This good intention, though, posed strong logistical challenges in terms of fact checking:

Harris: It was as much of a logistical challenge as getting the flood information was getting information about where to donate, and what they should donate to, and what were the legitimate non-profits, and now what were people collecting for.

Using crowdsourced information, the three editors ultimately identified 96 local donation and relief efforts that formed in the first two weeks of the storm, most of which were hyperlocal in nature.

Early in the crisis, Lajara, Reischel and Harris experimented with using a Google Spreadsheet to keep track of persons reported missing or stranded. That spreadsheet was editable by anyone, but they found that completely open access had several drawbacks. First, due to technical limitations of the platform, too many editors at once made the spreadsheet hard to edit. Second, people often added information about missing persons, but tended not to update the spreadsheet when those missing were found. Lajara, Reischel and Harris later incorporated these lessons into how they vetted charities. For this activity, Lajara created a form-based submission system using Google Survey that anyone could fill out to report a donation or volunteering opportunity. The survey pushed crowd-generated reports to a spreadsheet that only Harris, Reischel and Lajara could edit. As they vetted information, the editors updated items with text stating that it was "verified." Vetted information found its way into stories posted by the WP and the Freeman. But the spreadsheet itself was also posted—publically visible, though not publically editable.

| County | Relief for: | Relief resource name | Money | Volunteers | Food | Clothing | Misc. | Checks payable to: | Address | To |
|--------|--|---|-------------------------------------|------------|------|----------|-------|--------------------|--|----------|
| | Phoenicia | Shandaken Phoenicia Rotary Civil Emergency Response Team - C.E.R.T. | Yes -- see "Phoenicia rotary" above | Yes | No | No | | | | |
| | Verified flood relief organizations in this spreadsheet in Ulster, Schoharie, Greene, and Delaware counties. | "After the deluge" Benefit concert September 28th 5:30-10p.m. | Verified | No | No | No | No | | Keegan Ales, 20 St. James St., Kingston NY 12401 | Kingston |
| | | Free Hurricane Irene Holistic | | | | | | | | |

Figure 1. A Google form created by the Daily Freeman's Life Editor Ivan Lajara outputted to this published spreadsheet last titled "Catskill Flooding: Donation and relief centers." The form generated entries that were then vetted by Lajara and the Watershed Post editors with help from trusted sources. The public helped populate entries.

Harris and Lajara reported that charity vetting took place through a combination of traditional reporting techniques such as phoning those involved in an effort as well as through their network of collaborators and trusted sources, a process that took place largely outside of public view. While promoting the spreadsheet, the three made no general request of the crowd to help verify the charities.



Our Catskill Flooding spreadsheet: Donation and relief centers (click here to open):

This spreadsheet, which you can open in a separate window by clicking here, is a list of volunteer opportunities and places to give money.

However, the Watershed Post and the Daily Freeman did promote the survey that enabled the crowd to readily, though not directly, add information to the spreadsheet. By posting the spreadsheet the three journalists were able to scale up their ability to vet information because trusted collaborators and sources could help monitor the status of the vetting process. For example, another interviewee reported vouching for specific community efforts visible on the spreadsheet to the WP editors via private communication. In addition, audience members who were seeking volunteer and donation information did not have to wait until the information became a news item. They could elect to review the spreadsheet and conduct their own vetting process.

CONCLUSION

This research, which examines online collaboration during a crisis event, foregrounds work practices that developed in the effort to help curb misinformation. One of these practices, *visible skepticism*, addresses misinformation as it occurs without discouraging the crowd from sharing information. In this case, the practice was intentionally cultivated by those leading the effort and taken up by other participants in the crowd who acted to curtail rumors as they surfaced. These leaders also made visible both verified and unverified information about charity efforts. They made it easy to distinguish what information was vetted and what was not, thus striking a balance between timeliness and accuracy.

The appropriation of familiar technologies such as a liveblog, Google survey and Google spreadsheet enabled the group to include contributions from the larger crowd. The leaders of this effort carefully adapted these tools and structured their work practices around them. Ultimately, they were able to channel crowd participation through an artful combination of collaboration with trusted sources such as other regional reporters (sometimes taking place behind the scenes) and collaboration with the crowd through appropriately moderated activities. Importantly, these practices and adaptations were able to raise the visibility of dozens of community relief efforts thereby making a significant contribution to community resiliency.

ACKNOWLEDGMENTS

We would like to thank all participants and participating organizations. In particular, project sponsor WGXC 90.7 FM.

REFERENCES

1. Bruns, A., Burgess, J. E., Crawford, K., & Shaw, F. (2012). #qldfloods and @QPSMedia: Crisis Communication on Twitter in the 2011 South East Queensland Floods. White Paper. ARC Centre of Excellence for Creative Industries & Innovation (CCI).
2. Carvin, A. (2013). Distant Witness. CUNY Journalism Press.
3. Fritz, C. E., & Mathewson, J. H. (1957). Convergence Behavior in Disasters: A Problem in Social Control: a Special Report Prepared for the Committee on Disaster Studies. National Academy of Sciences National Research Council.
4. Gupta, A., Lamba, H., Kumaraguru, P., & Joshi, A. (2013). Faking Sandy: characterizing and identifying fake images on Twitter during Hurricane Sandy. In *Proc. of WWW 2013 companion* (pp. 729-736).
5. Hill, K. (2012). Hurricane Sandy, @ComfortablySmug, and the Flood of Social Media Misinformation. *Forbes.com*. (October 30, 2012). Available at: <http://www.forbes.com/sites/kashmirhill/2012/10/30/hurricane-sandy-and-the-flood-of-social-media-misinformation/>
6. Hughes, A. L. & Palen, L. (2012). The Evolving Role of the Public Information Officer: An Examination of Social Media in Emergency Management, *J. of Homeland Security & Emerg. Management*, 9(1), article 22.
7. Madrigal, A. (2013). It wasn't Sunil Tripathi: The Anatomy of a Misinformation Disaster. *The Atlantic*. (April 19). Available at: <http://www.theatlantic.com/technology/archive/2013/04/it-wasnt-sunil-tripathi-the-anatomy-of-a-misinformation-disaster/275155/>
8. Marcus, G. (1995) Ethnography In/Of the World System: the Emergence of Multi-sited Ethnography. *Annual Review of Anthropology*, V. 24, 95-117.
9. Mendoza, M., Poblete, B., & Castillo, C. (2010). Twitter under crisis: can we trust what we RT? *SOMA '10 Proceedings of the First Workshop on Social Media Analytics*. ACM, NY. Pp 71-79.
10. Sarcevic, A. Palen, L., White, J., Starbird, K., Bagdouri, M. & Anderson, K. (2012). "Beacons of hope" in Decentralized Coordination: Learning from On-the-Ground Medical Twitterers during the 2010 Haiti Earthquake. In *Proc. of CSCW, 2012*. New York: ACM, 47-56.
11. Starbird, K., & Palen, L. (2013). Working and sustaining the virtual Disaster Desk. In *Proc. of CSCW 2013*, ACM, 491-502.
12. Starbird, K., Maddock, J., Orand, M., Achterman, P., & Mason, B. (2014). Rumors, False Flags, and Digital Vigilantes: Misinformation on Twitter after the 2013 Boston Marathon Bombings *iConference 2014*.
13. Vieweg, S., Palen, L., Liu, S., Hughes, A., & Sutton, J. (2008). Collective Intelligence in Disaster: An Examination of the Phenomenon in the Aftermath of the 2007 Virginia Tech Shootings. Proceedings of the *Information Systems for Crisis Response and Management Conference (ISCRAM 2008)*.