

# Community Incident Chatter: Informing Local Incidents by Aggregating Local News and Social Media Content

**Patrick C. Shih**

The Pennsylvania State University  
patshih@ist.psu.edu

**Kyungsik Han**

The Pennsylvania State University  
kuh178@psu.edu

**John M. Carroll**

The Pennsylvania State University  
jcarroll@ist.psu.edu

## ABSTRACT

The emergence of social media provides an additional channel for broadcasting information to the public and support two-way communication between governmental stakeholders and the public during crisis. Research has focused on large-scale events, and few have investigated how social media can contribute to civic awareness and participation of small-scale incidents in a community-oriented context. Moreover, social media have been criticized because it is overabundant with noisy, inaccurate, and unprofessional information that are often misleading. This presents a serious challenge for community members to identify information that are relevant to a local incident. We introduce Community Incident Chatter (CIC), a smartphone application that is designed to aggregate information reported by formal news agencies and social media surrounding local incidents. Participants in a preliminary user study indicate that the community-oriented information presented in CIC is informative, relevant to the community, and has the potential of empowering community residents for responding to and managing local incidents.

## Keywords

Local incidents and emergencies, Social media, News and tweet aggregation, Civic awareness and participation

## INTRODUCTION

Web 2.0 technologies have been suggested as part of new approaches to creating and disseminating news content (Carroll & Rosson, 2008). Microblogging in particular has garnered much attention as an information dissemination platform in various domains ranging from presidential elections (Tumasjan et al., 2010) to major incidents or disasters (Vieweg et al., 2010; Yardi & boyd, 2010) because of its capability to reach out to more people more quickly. In the context of emergency events or incidents, studies have reported that microblogged content (e.g., texts, photos, videos, etc.) contributes to situational awareness (Vieweg et al., 2010) that facilitates emergency communications among people and organizations. For example, the Pew Research Center reports that news, information, photos, and videos made up more than half of all the Twitter conversations when Hurricane Sandy hit the mid-Atlantic and Northeast U.S. in 2012 (Guskin & Hitlin, 2012). These prior research efforts have primarily focused on understanding social media usage and its implications in national emergency situations, but few have focused on understanding small-scale incidents in a local community context.

There exists an opportunity for microblogs to play a significant role in creating and disseminating local community incidents to people (Kwak et al., 2010). Microblogs allow local residents to access or create diverse community information easily and conveniently. However, a challenge still remains for finding pertinent community information in microblogs, as they tend to contain a lot of noise, which lowers their perceived informativeness and usefulness (Andre et al., 2012). Researchers have raised similar perspective in emergency contexts as well. For example, Tapia et al. (2013) argue that it is necessary to consider appropriate contexts, situations, or conditions when using or applying microblogged data in crisis response and highlight the potential of employing the appropriate microblog users to create and disseminate more reliable, meaningful and

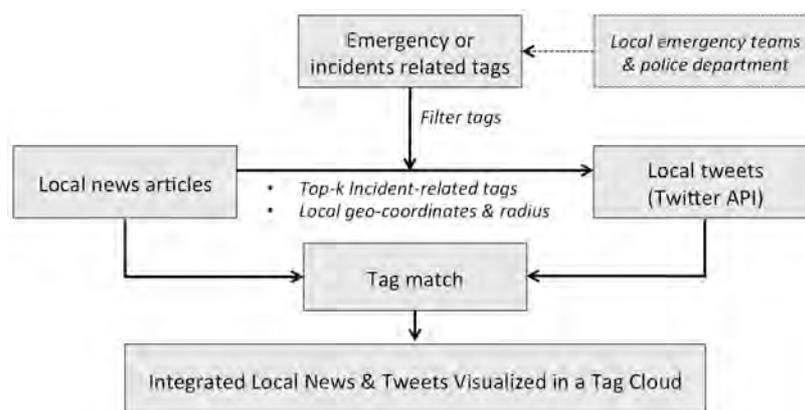
*Proceedings of the 11<sup>th</sup> International ISCRAM Conference – University Park, Pennsylvania, USA, May 2014*  
S.R. Hiltz, M.S. Pfaff, L. Plotnick, and P.C. Shih, eds.

trustworthy information. Therefore, we are interested in how to utilize microblogged content to provide relevant information that helps local residents be aware of or respond to local incidents.

We introduce a smartphone application called Community Incident Chatter (CIC) that has been designed to take advantage of microblogged content in a local community context. Based on a set of tags that pertains to local incidents, CIC aggregates formal local news articles from local news agencies and Twitter feeds in supporting community building and awareness. In this paper, we describe the CIC implementation and usage scenarios, and present results from a preliminary user study of 10 users.

## COMMUNITY INCIDENT CHATTER APPLICATION

Our assumption is that local news sites are the most active community news providers to local residents when compared to those from other organizations and national news media, and that local news articles would be the most accessible resource for getting detailed information about local incidents.



**Figure 1. Overview of Community Incident Chatter (CIC). CIC aggregates locally relevant news articles and tweets and presents them in an integrated fashion based on a set of keywords pertaining to local incidents.**

Figure 1 shows the overview of CIC. CIC collects RSS news article feeds from local news sites and stores them in a database on an hourly basis. We leverage only news feeds that provide locally relevant news content (e.g., ones that end with “.../local”). Each news item contains a set of metadata, including a title, description, source URL, and created date and time. After a typical pre-processing step in NLP (e.g., stop word removal), word tags are extracted from news titles and descriptions, and their TF-IDF (Term Frequency-Inverse Document Frequency) scores are calculated. The TF-IDF score is a standard metric in Information Retrieval to measure the “importance” of a tag (Salton and Buckley, 1988). The TF-IDF value of a tag  $w$  in an article  $a$  is positively correlated with the frequency of  $w$  within  $a$  but is negatively correlated with the frequency of  $w$  in the entire article collection.

Next, we manually defined 50 tags that pertain to local incidents; for example, theft, assault, damage, warning, pain, police, accident, severe weather, fire, vandalism, etc. To increase the credibility and reliability of tags, we plan to obtain tags from official local emergency planning teams or a police department, which we expect could be updated on a regular basis. CIC checks if there are any matches between top- $k$  tags from the local news articles and the set of incident-related tags. Then, the “top- $k$  incident-related” tags along with a geo-coordinate and a radius (we limited our search to a 3-mile radius because it is enough to cover the region of our local community) are used as parameters to the Twitter search API. Twitter returns a set of tweets that are pertinent to the provided tag and location, and the results are stored in the server database. The call for this request occurs immediately after each news article collection cycle. CIC presents both local news articles and tweets that are associated to the specific local incident-related tags in an integrated fashion.

As depicted in Figure 2, CIC displays the local incident tags in a tag cloud, because it supports searching, browsing, and capturing main topics of information in a coherent view. The size of the tag represents the number of articles mentioning the tag. The number of tags displayed changes dynamically depending on whether any related incident occurred recently in this community in a specific timeframe (Figure 2 shows tag clouds generated based on a 48-hour timeframe). If one of the tags is clicked, the corresponding local news article(s) and tweets by local residents will be displayed (see Figure 3 below). The news article is displayed at the top of the page, and tweets are displayed below the article. By default, we show the news article that has the highest TF-IDF score for the selected tag, and the user is able to scroll through news articles associated with the

same tag in order of descending TF-IDF scores. CIC also provides a link to a more detailed news page for those who want to access more detailed information (Figure 2, right). By using a list of tags that best represents local incidents (e.g., vandalism, fire, etc.), we expect that community members will be able to quickly access the most urgent community events using CIC.

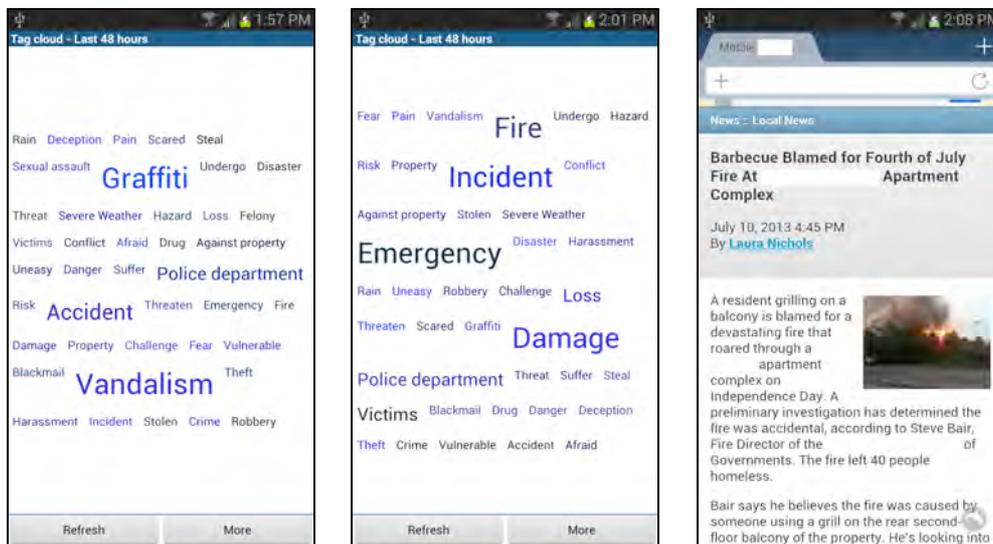


Figure 2. Left and Middle: A Tag cloud that presents current local incidents in different tag sizes depending on frequency of mention in local new articles and tweets; Right: local news article (names anonymized).

COMMUNITY INCIDENT CHATTER USAGE EXAMPLES

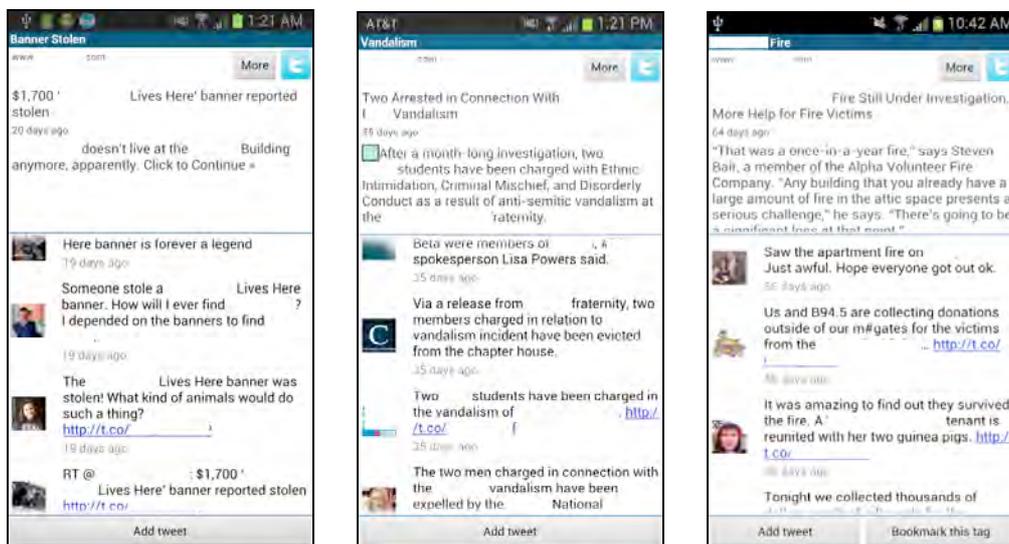


Figure 3. Local news articles and tweets displayed in CIC presented during the user study. Left: news about the missing banner; Middle: a vandalism incident; Right: a fire incident (names anonymized).

Figure 3 depicts three representative examples that best demonstrate the CIC usage scenarios. These examples refer to the local news and citizen tweets about a case of missing banner, a vandalism incident, and an apartment fire accident.

In our local community, a small college town in the Northeastern US, there are large banners displayed in the local neighborhoods with the slogan “[anonymized] Lives Here.” It is an initiative to provide an opportunity for students, alumni, faculty, and staff to connect with each other. Each banner costs about \$1,700 each, and one of them was suspected to be stolen during the winter break in 2013. CIC presents a number of responses from local Twitter users. Some of them simply retweeted the existing posts to spread the information to more people, while others expressed personal thoughts or opinions on this incident.

*“The [anonymized] banner was stolen! What kind of animals would do such a thing?”*

*“Someone stole a [anonymized] Lives Here banner. How will I ever find [anonymized]? I depended on the banners to find [anonymized]”*

The second example describes the incident of a case of vandalism near the fraternity housing, of which two undergraduate students were accused and implicated in November 2013. It was a serious incident because the graffiti, which the two painted on the street, the buildings, and cars, consisted of abusive languages, sexual images, initials such as “KKK”, and other offensive scribbling, which greatly undermined the reputation of the university. There were a number of news articles and tweets posted over a month-long period, and we found that most of the tweets were posted to provide related news articles and updates. The following comment shows an update of this incident tweeted by a local resident. The resident posted the tweet before the eviction was even reported in any official news articles.

*“Via a release from [anonymized] fraternity, two members charged in relation to vandalism incident have been evicted from the chapter house”*

The last example depicts a fire incident in a local apartment complex occurred in July 2013. The cause of the fire was due to an unattended charcoal grill, which destroyed the building and displaced more than 40 residents; many of them included young children. This incident was covered extensively by all local news media and heavily discussed by local residents in social media. Many residents tweeted what they saw at the scene.

*“Saw the apartment fire on [anonymized]. Just awful. Hope everyone got out ok”*

*“Difficult to count how many but handful of [anonymized] residents are outside holding infants and young children”*

We also found that the local residents mobilized the social media by tweeting resourceful information that contained directions about how to make donations of cash, food, clothing, toys, and other goods to the victims or expressing personal sentiment, which are not easily accessible through typical online or offline news channels.

*“[anonymized] are collecting donations outside of [anonymized] for the victims from the [anonymized] Drive...”*

*“One day after that raging apartment fire on [anonymized] Drive, residents return home to assess the damage”*

These examples show that CIC is able to utilize tags to successfully extract community incidents reported in formal local news articles as well as acting a bridge to connect these official reports to tweets related to these incidents that are generated by local community members.

## PRELIMINARY USER STUDY

To evaluate how people might perceive the community incident information presented in the CIC interface, we conducted a preliminary user study of 10 participants. Participants are residents living in the local area who are generally interested in local community information. They were provided with a smartphone installed with CIC, and were presented with local incident information similar to those described earlier in this paper. They were asked to explore CIC and provide feedback regarding the information relevance and usefulness.

Overall, 8 out of 10 participants found that the local news articles contained community incident information that are relevant to the tag, and the aggregated tweets provided additional information beyond what were reported by official news sources.

*“It gives me a single place to see people’s tweets as well as the news articles. It is a one stop shop to emergency information in this area” (P5).*

Some participants mentioned that reading other people’s opinions and perspectives regarding community incidents.

*“I would have more opinions because I would be better educated on local crisis and would feel that I could tweet educationally about things that are happening in the area and maybe even participate and help out!” (P9).*

The ability to identify people’s inclinations further afforded a way to bridge the community resources and bring people together.

*“I would like to use this to maybe find people also interested in helping and pool our resources together” (P3).*

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

CIC shows the potential benefits helping local residents become more aware of and be able to more actively participate in helping with local crisis. The tweets sometimes contained more meaningful or informative content than news articles. There are many other examples that describe local incidents such as police investigations, severe weather (heavy rain or snow) reports, car accidents, and other smaller community incidents with a number of related news articles and tweets that show the advantage of CIC in helping local residents connect to the local incidents and corresponding to them.

## DISCUSSIONS AND CONCLUSION

We argue that CIC is especially helpful in the context of a local community. To demonstrate the utilities and usage scenarios of CIC, we described three representative examples and conducted a preliminary user study that highlighted CIC's potential of increasing civic awareness and participation. Our preliminary user study with local residents showed that CIC presents richer and more dynamic local community incident information. By integrating the scattered voices of the individuals from local tweets with local news reports about community incidents, CIC demonstrates that tweets can be leveraged to support community awareness to small-scale local incidents and emergency responses. Furthermore, our intention is not only to enable a more directed and accessible way for local residents to be aware of community incident information, but also to encourage them to actively interact with these local information and helping with the community.

Another salient perspective is that CIC leverages existing content because tweets already posted on Twitter and local news articles are readily available public information. This is especially helpful for real world tool adoption because CIC is built to leverage existing news and social media infrastructure, and it is not just a novel online platform with no starting content. In their preliminary observation, Hughes and Palen report (2009) that Twitter users who were involved in broadcasting information of emergency events tended to become long-term adopters of the technology, and we hope to see similar impacts and consequences from CIC in the future.

We are currently refining the CIC interface and adding additional features based on the findings of our preliminary user study. We are in the process of releasing CIC in two versions (Android and iOS) to the public and recruiting participants for a long-term deployment study. We plan to investigate the affordances of smartphones, namely mobility and accessibility, on reading, creating, and sharing urgent local incident information at the community level. We also plan to utilize local emergency planning resources to provide more reliable and trustworthy local incidents information to local residents.

## REFERENCES

1. Guskin, E. & Hitlin, P. (2012). Hurricane Sandy and Twitter, *Pew Research Journalism Project*. <http://www.journalism.org/2012/11/06/hurricane-sandy-and-twitter/>, accessed on January 15, 2014.
2. Andre, P., Bernstein, M. S., & Luther, K. (2012). Who Gives A Tweet? Evaluating Microblog Content Value. In *Proceedings of the International Conference on Computer Supported Cooperative Work and Social Computing (CSCW '12)*, ACM Press, 471-474.
3. Carroll, J. M. & Rosson M. B. (2008). Theorizing mobility in community networks. *Journal of Human-Computer Studies*, 66 (12), 944-962.
4. Hughes, A. L. & Palen, L. (2009). Twitter Adoption and Use in Mass Convergence and Emergency Events. In *Proceedings of the 10th International ISCRAM Conference (ISCRAM '09)*.
5. Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a Social Network or a News Media? In *Proceedings of the International Conference on World Wide Web (WWW '10)*, ACM Press, 591-600.
6. Salton, G. & Buckley, C. (1988). Term-weighting approaches in automatic text retrieval. *Journal of Information Processing & Management*, 24 (5), 513-523.
7. Tapia, A. H., Moore, K. A., & Johnson, N. J. (2013). Beyond the Trustworthy Tweet: A Deeper Understanding of Microblogged Data Use by Disaster Response and Humanitarian Relief Organizations. In *Proceedings of the 10th International ISCRAM Conference (ISCRAM '13)*.
8. Tumasjan, A., Sprenger, T. O., Sandner, P. G., & Welpe, I. M. (2010). Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM '10)*.
9. Vieweg, S., Hughes, A. L., Starbird, K., & Palen, L. (2010). Microblogging During Two Natural Hazard Events: What Twitter May Contribute to Situational Awareness. In *Proceedings of the International Conference on Human Factors in Computing Systems (CHI '10)*, ACM Press, 1079-1088.
10. Yardi, S. & boyd, D. (2010). Tweeting from the Town Square: Measuring Geographic Local Networks. In *Proceedings of International Conference on Weblogs and Social Media (ICWSM '10)*.