

# Semantic Geotagging: A Location-Based Hypermedia Approach to Creating Situational Awareness

Ray Bareiss, Martin Griss, Steven Rosenberg, and Yu Zhang

Carnegie Mellon University, Silicon Valley Campus,  
Moffett Field, California 04035, USA

{ray.bareiss,martin.griss,steven.rosenberg,  
ian.zhang}@sv.cmu.edu

**Abstract.** As emergency first responders and commanders increasingly use mobile phones, tablets, and social media to communicate, coordinate, and manage information during disasters, we see a need and opportunity to provide a mobile device-appropriate semantic layer to a geographically-based common operating picture. The challenge is to provide a simple, usable structure for a rapidly growing body of information to simplify the development of situational awareness in an unfolding disaster. We use a hyperlinked structure based on the ASK model to organize information in a readily accessible form. In this paper we describe our initial design and experience with an Android-based prototype, supported by a Ruby on Rails-based repository service. Our prototype allows the incorporation, aggregation, assessment, and redistribution of dynamic human-generated and sensor-derived information.

**Keywords:** mobile applications, emergency response, social media, crowd sourcing, mobile collaboration.

## 1 Introduction

There is an increased urgency among all levels of government in improving disaster management and coordination. Recent disasters have demonstrated the difficulties in mounting fast, coordinated, and successful responses. While many things could be improved, it is clear that there are especially serious problems in acquiring, managing, and disseminating the large body of information required to develop accurate and trustworthy situational awareness.

Disaster Response Teams must share a common situational awareness to prioritize their activities, to work effectively, and to coordinate with other teams. In both the disaster management and other communities, this is ideally provided by a Common Operating Picture (COP), see, e.g., [1]. A natural framework for this COP is a geographical information system in which key elements, such as people, places, resources and events, are situated and tagged with accompanying content, usually status information. We are working with an open architecture, web-based COP, shown in Figure 1. The content of these geospatial elements is then updated as new information flows in.



**Fig. 1.** Web-based Common Operating Picture

Traditionally, disaster response follows a command-and-control model embodying primarily hierarchical information flows, filtered manually by incident commanders and other experts. Situational awareness grows as information flows upward from lower levels and is aggregated and organized by the higher levels; conversely it flows down from commanders to front-line responders on a need-to-know basis. Today the confluence of powerful smart phones and other mobile devices, ambient and mobile sensors, fast and ubiquitous networking, social media, and cloud-based services can and often do provide increasing volumes of on-line information (often “real time”) about breaking events. These new technologies increasingly permit the “crowd sourcing” of information, from volunteers and local residents as well as professional responders, potentially generating broad situational awareness more quickly than hierarchical processing but with increased opportunity for error, disorganization, and overload. Mobile, context-aware applications are increasingly proposed and used for disaster response [2, 3] and for sensing people and things nearby in the environment [4].

The increasing use of smart phones by responders and citizens empowers humans as sensors: They take pictures, annotate, send SMS and Twitter messages, and post Facebook messages; many events are geotagged and can be displayed on maps using appropriate icons. The challenge is to provide a simple, usable structure for a potentially large, diverse, and rapidly growing body of information in order to simplify the development of situational awareness in an unfolding disaster.

Furthermore, any such organizational mechanism must be easy to use in the mobile, hurried, small screen environment, and customizable to the context, such as the roles and tasks performed by different responders at different times during a disaster response.

## 2 The Need for a Structured Record of Emergency Messages

The goal of our work has been to develop a prototype system to support the collaborative construction of situational awareness by distributed, mobile responders, providing a better situational awareness of a complex disaster in less time than traditional hierarchical "processing" would require.

A large-scale disaster, such as the recent pipeline explosion and resulting fire in San Bruno, California, can result in literally thousands of messages among emergency responders. (The number might go up by an order of magnitude were messages from the public -- via phone, text, Twitter, et cetera -- included.) Professional responders currently send messages primarily via voice radio, but are increasing using message-based packet radio, SMS, and even Twitter. Many of these messages may have attached images and can be geotagged. The message stream has a transitory quality. If a responder misses a message or if its content is not relevant at the time it is received, the message is essentially lost to the responder. The only information repository typically available is a linear recording of radio transmissions or the text message stream.

The challenge is not one of information scarcity but of information overload. For example, of the many, many on-line responses during the recent San Bruno, California disaster, only a small fraction were useful to first responders in dealing with their immediate situations. An easily browsable repository of messages and information would enable emergency responders to quickly find the few most immediately relevant messages among the thousands. The challenge is to define an organizational structure that can be easily applied when new messages are generated and that is at the same time useful for browsing a large corpus of messages.

Our solution, used in our prototype system, is to add semantically linked annotations to the COP. We use geotagging to associate messages (and images, cf. [5]) with locations in the COP's map-based geospatial interface. This provides a situationally relevant context for related messages. A small set of typed links with well-defined semantics then associates messages with incidents, and related messages with each other. For example, an emergency dispatcher might create a new incident based on a citizen's report of a fire; a police patrol might link an initial situation description and upload a geotagged picture; a fire unit might link an assessment to the initial description; still another fire unit might link another assessment which augments the first based on new observations; an incident commander might provide advice for dealing with the situation; et cetera.

It might be informative to compare this approach to two simpler alternatives: GeoTwitter [6] and GeoChat [7]. GeoTwitter provides only location-based indexing of tweets, and a user can easily be overwhelmed by a large number of tweets in a limited geographical space (Figures 2a and 2b). GeoChat indexes online, multi-way conversations (i.e., chats) geographically.

The strictly temporal nature of a chat makes it easy to miss things, and, like a Twitter stream, it is difficult to search, given the likely high similarity of messages within a chat. Ushahidi (<http://www.usahidi.com>), which combines aspects of these two approaches, is an open source system currently gaining traction in the disaster response community.




Fig. 2a. The GeoTwitter Interface (1/2)




Fig. 2b. The GeoTwitter Interface (2/2)

Other researchers have also addressed the use of geotagged information in disaster response. The increasing role of geotagged images and text in spatial decision making is covered in [8], and motivates approaches such as integrating geotagged citizen input from smartphones and desktops into GIS systems [9]; as these authors suggest, it is not a trivial task to automatically merge (incompatible) information into an existing GIS framework. In [10] the implications of user-generated spatial content, especially using GPS-enabled smartphones, and crowdsourcing to add geotagged annotations, extensions and mashups to base maps (such as GeoWeb) are evaluated. While this is an excellent way to quickly generate and distribute dynamic map-based information, the torrent of potentially unreliable information can increase uncertainty and add to confusion, and may in fact inhibit useful communication. Details on the use of geotagged information in the Haiti response and recovery can be found in [11]; issues of information accuracy, reliability and provenance, and incompatible duplication of effort are surfaced. Maiyo [12] stresses the importance of real-time collaborative disaster management tools, to better allow the integration of different types of mapping data from diverse sources.

### 3 A Cognitively Inspired Approach to Organizational Memory

In the 1990's, researchers at Northwestern University's Institute for the Learning Sciences developed an effective approach to organizing large-scale hypermedia systems, called ASK Systems [13, 14, 15]. This family of hypermedia systems uses a theory of conversational coherence [16] to provide a limited variety of typed links between content elements that correspond to the sensible connections between utterances that arise in a coherent conversation. Over 20 such systems have been built in domains as wide-ranging as military planning, management consulting, engineering design, and public water supply issues.

An ASK system must provide three types of user interaction: zooming to provide initial access (beginning a conversation), browsing to explore a sub-area of content (requesting information via a question-and-answer dialog), and responding (contributing new information that will remain a part of the conversational record).

#### 3.1 Zooming into the Body of Information

Given that disasters, such as fires and gas leaks, cause problems in specific locations, a map-based interface provides a natural way for a responder, who is used to reading maps, to indicate an area of interest. When a responder arrives on the scene, he or she sees a map of the local area with icons indicating incidents of various predefined types, e.g., a fire or a gas leak (Figure 3).

This standard geospatial layout, also used in a COP, provides the first level of context for relevant information. Semantic Geotagging provides additional layers that organize related relevant information in conversational threads associated with the incident. The following system description shows how a responder can quickly find and drill down on aspects of an incident of immediate interest to him or her.

Touching an icon displays a pop-up summary of the incident (e.g., "Interior fire, Building 23, Moffett Field). Touching the summary displays a more detailed description and categories of information about the incident (Figure 4):

- Situation Description: A living document that provides an overview of the situation via a series of messages to which a responder contributes when he or she learns more about the incident
- Problem Report: A message that reports a responder has been unable to complete an assigned task because of unforeseen circumstances, lack of necessary personnel or equipment, et cetera.

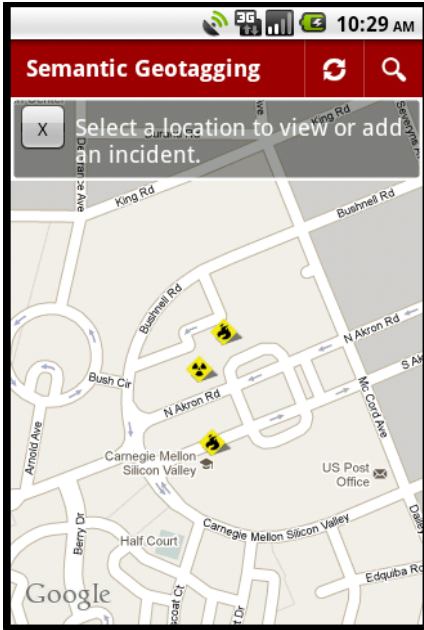


Fig. 3. The Top-Level Interface



Fig. 4. Top-Level Information about an Incident

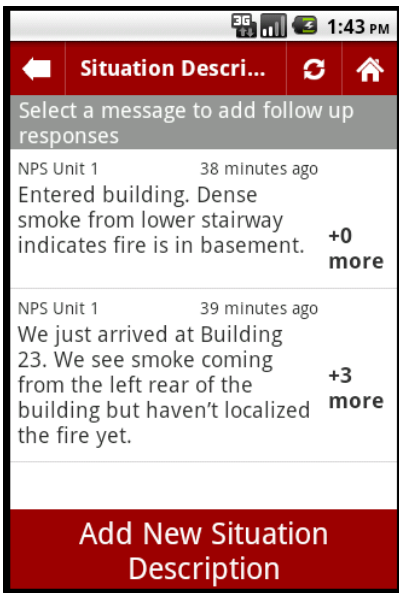


Fig. 5. Situation Description Messages

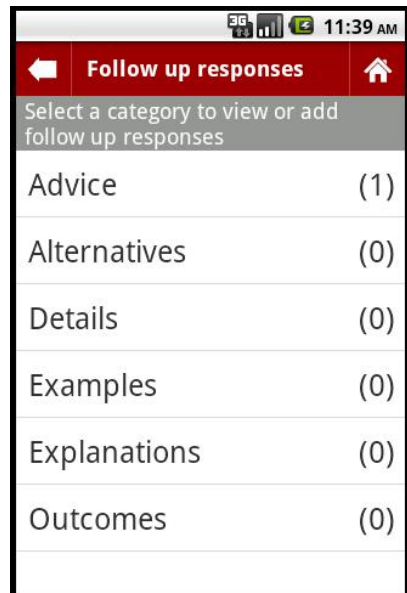


Fig. 6. Categories of Follow-Up Information

- **Something Else:** A catch-all category included in the current prototype for testing purposes to provide a place for test subjects to enter information that they do not believe belongs in one of the other categories. (We hope to remove this category and possibly to extend the list of categories as a result of testing.)

The screen also provides an Alert button, which enables a responder to broadcast a high-priority message to all responders in the area (e.g., “man down” or “evacuate”), and an All Clear button, which enables a responder to notify others that the incident has been resolved.

Touching a category, such as Situation Description, displays a list of previously entered messages, providing top-level information about an incident time-stamped and in reverse chronological order (Figure 5). The annotation +N more to the right of a message indicates the number of follow-up messages that responders have entered. Touching the message displays it full-screen and also provides access to the categorized list of follow-ups.

### 3.2 Browsing Follow Up Responses

A top-level message, such as the initial description of a building fire, can result in tens if not hundreds of follow-up messages from a number of responders. The ASK approach to organizing a hypermedia organizational memory uniquely balances power and simplicity. Follow-up messages are categorized into one of six categories, each of which might be thought of as corresponding to a general type of follow-up question that a listener in a conversation might ask in response to presentation of a piece of information. Several variant category schemes have been used in previous systems; the ones chosen for use in this prototype are (Figure 6):

- **Advice:** What do you recommend doing (or not doing) in this situation?
- **Alternatives:** Are there alternative explanations, pieces of advice, et cetera?
- **Details:** Can you provide more details about the information I’ve just received?
- **Examples:** Can you provide a specific example of the generality I’ve just read (or heard)?
- **Explanations:** Can you explain what you’ve reported, recommended, et cetera?
- **Outcomes:** What happened (or might happen) as a result of the situation I’ve just read (or heard)?

Note that although the categories remain consistent (and consistently placed within the interface), the meanings of the categories may shift slightly according to the contextualizing piece of information with which they are associated. For example, as noted above, Alternatives might be thought of as a request for an alternative explanation, alternative advice, an alternative example, or an alternative possible outcome.

To view a particular category of follow-up information, the user selects the category by touching it. (Note that the number in parentheses indicates the number of follow-up pieces of information of that type that have been entered.) Follow-ups can have follow-ups which can have deeper follow-ups, et cetera, enabling a user to explore as deeply as he or she feels to be necessary. Note that all users may not take the same path through the information. Instead, the questions raised in a particular user's mind in response to a piece of information determines the next category of information requested and viewed. Back and Home buttons enable the user to "back out" of a particular thread of conversation.

### 3.3 Responding to a Message

A user is not restricted to viewing information entered by others. At any time, he or she can contribute information to the conversation. When responding, the categories enumerated above take on the semantics of conversational connectives, the bridges that people often use when responding to utterances by others:

- Advice: Here's some advice for dealing with that.
- Alternatives: I disagree with your explanation, advice, et cetera, and here's my alternative.
- Details: I can tell you more about that.
- Examples: Here's an example of what you (or I) just said.
- Explanations: I can explain the phenomenon just reported.
- Outcomes: Here's what happened (or is likely to happen).

A user begins his or her response by selecting one of these categories and then goes on to enter a response. The current prototype requires the user to type the response, but we recognize that this is likely to be unworkable in an emergency situation (Figure 7). Our ultimate goal is to enable the user to dictate a response (much like emergency responders use their radios today) and to use speech-to-text technology to produce a text-based version for default display to future users.

### 3.4 Summary

The features described in the preceding section show how first responders can quickly home in on exactly the information of relevance to them at that time and location. Further, they have the capability to augment it, on the spot, if desired. This additional information becomes immediately available to everyone – responders in the field and chain-of-command at the emergency operations center. Of course, if it were simply broadcast to everyone, we would have recreated a situation of information overload. By providing a simple system for placing new information in the appropriate geographical context and conversational thread, we leverage the benefits of today's always connected, mobile environment for gathering information with the advantage of the semantic geotagging system for ensuring responders can easily find relevant information.



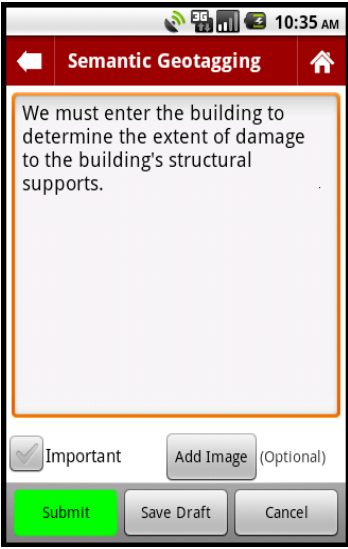


Fig. 7. The Current Information Entry Screen

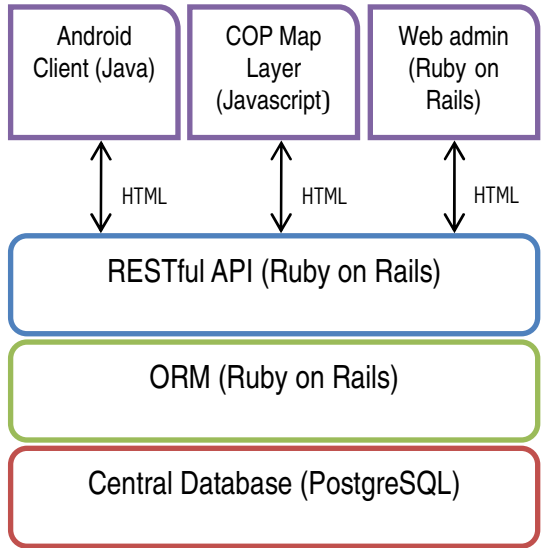


Fig. 8. The System Architecture

#### 4 An Example of the System in Action

Mike, a fire department battalion chief, arrives at Building 23 of Moffett Field with his engine company. They see smoke but no visible fire. Mike pulls out his Android phone running the Semantic Geotagging application. A fire icon, superimposed on Building 23, is centered in the map interface of his phone. (His location has been determined via his phone's GPS.) Mike touches the fire icon to select it and then selects Situation Description from the list of top-level categories of information about the fire. Two messages have already been entered by NASA Protective Services (NPS); after skimming them Mike does his own size-up (see, e.g., [17]).

Mike then returns to the two messages that have been input. He scans the first one, an initial posting that lacks detail, and then the second. He sees that someone has entered follow-up messages to the second situation description message, so he navigates to view them. In particular, he sees that someone has entered Advice on dealing with the situation, and he reads that. The advice doesn't jibe with his experience, so Mike enters Alternative advice, linked to the first advice message.

Mike and his crew go into the building. He sees that the structural integrity of the main staircase is potentially a problem, so he adds a message to that effect to the Building 23 fire Situation Description (at the top level, not as a follow-up to an existing message).

Later, while taking a break and walking around the Moffett Field parade ground, he smells gas near Building 19. He sees that no one has entered information about this incident, so he enters this new event into the system. He begins by touching the

location on the screen to create the new incident. He is prompted to choose a type of incident that generates a corresponding map icon, and he enters basic information – the building number and the nature of the problem. He also enters a sparse initial Situation Description and takes a picture to aid others in identifying the building (cf. [5]); the picture is automatically associated with his description. Finally, he enters a Problem Report because his crew is not well-equipped to handle this sort of emergency. This new situation will be immediately visible both at the emergency operations center and to any other responders interested in the status of this location (i.e., immediately updated situation awareness is available). Of course, determining an appropriate response is still the responsibility of the chain of command.

## 5 Prototype Implementation

The system employs a multi-client/server architecture (Figure 8). The primary client is a native Android application written in Java, using the Android framework. There are two versions: One with an interface optimized for Motorola Droid phones and the other with an interface optimized for the Samsung Galaxy tablet. The latter interface is resized, and its buttons are enlarged, but the text size is the same as the phone interface to display more content while minimizing the need for scrolling during use. Other clients include a web map-based application that we developed for administration and testing and the Golden Gate Safety Network’s Common Operating Picture (also web based). All clients employ a Google map layer. The clients contain all of the “business logic” of the system and handle all interactions with users.

The server is implemented in Ruby on Rails. Use of the model-view-controller design pattern facilitates support for diverse clients. The server provides an object relational mapping (ORM) layer to interact with a PostgreSQL database. A set of representational state transfer (RESTful) API’s are implemented on top of the ORM to expose different functionalities to the various clients (see, e.g., [http://en.wikipedia.org/wiki/Representational\\_State\\_Transfer](http://en.wikipedia.org/wiki/Representational_State_Transfer)). The server is written to enable several people to update the evolving content, and each will immediately see updates made by others.

## 6 Preliminary Usability Testing

After developing a first version of the prototype, the authors conducted a streamlined cognitive walkthrough [18]. We made a number of relatively simple changes to the first version of the interface as a result. (These changes are incorporated in the screens presented earlier in this paper.)

We went on to conduct preliminary, low-ceremony usability testing (see, e.g., [19]). Our testing walks a subject through a generalized version of the scenario, presented in the example section of this paper (e.g., “Access the advice someone has entered as a follow-up. You disagree with the advice. How would you enter a message communicating your disagreement?”). We also ask a number of general questions to elicit the subject’s overall impressions of ease of use, missing

functionality, and the utility of such an application. We have conducted three individual user tests to date with retired and active-duty firefighters. They were not trained in the use of the application before undertaking the test.

All subjects understood the purpose of the application and generally found it to be easy to navigate. As noted earlier, all commented, nearly immediately, that voice input is a must for a fielded version of the application.

To date, we have discovered one fundamental usability issue: All subjects struggled with the idea of categorizing follow-up information entered in response to existing information – at least as the categorization process is afforded by the current interface. For example, in our testing scenario, they are asked to enter an alternative to existing advice with which they disagree; all subjects wanted to simply enter alternative advice as new advice at the same level as the advice with which they disagreed rather than linking their new advice to the existing advice explicitly as an alternative, thus creating the sort of “conversational thread” via which we envision organizing information.

Two possibilities exist for dealing with this issue: The first is to leave the interface as is and to train users to categorize follow-up information appropriately. The second is to make the process of information entry more explicitly procedural, so that a user must go through the steps (via three corresponding screens) of indicating the desire to respond to an existing piece of information, categorizing the response (e.g., as an alternative), and finally entering the follow-up. Since we believe it is unrealistic to insist on training all emergency responders in the process of information categorization and entry for our application, we have decided to implement and usability test the second alternative. We have suspended testing until this change has been made and plan to test three to seven additional users after doing so.

Our testing, to date, has also uncovered a few less fundamental things:

- On-screen instructions, which are white text on a gray background, are not adequately prominent; we will change the background and text colors before further testing.
- Not all categories of follow-up information seem useful in this emergency response context. The system originally had a Big Picture category that was removed after the cognitive walkthrough. Similarly, subjects have not seen the utility of an Examples category; that one will be removed as well before usability testing proceeds.
- A user needs to be able to ask a question as well as to enter information. The user should then be informed when the question is answered, and the answer should be added to the database. This will be added as a future enhancement.
- The map interface should be pre-populated with all units in the area. We have deferred work on this because of the technical issues involved in automatically compiling this information.
- Purpose-directed (or perhaps role-dependent) views of the database would be useful, especially when responders are engaged in a time-critical task. This complex potential enhancement is discussed in the ongoing work section below.

Additional enhancements are also discussed in the ongoing work section.

## 7 Ongoing Work

We are planning four threads of ongoing work: adding new functionality, expanding testing, integrating new sources of information, and envisioning new uses for the technology.

### 7.1 New Work

The most ambitious new functionality, as noted earlier, is a speech-to-text capability that will enable emergency responders to dictate messages instead of the current, rather impractical, alternative of typing them on the small screen of a smart phone. Because of the imperfect reliability of speech-to-text translation (we might realistically expect an accuracy ranging from 70-90% for continuous speech compared to 96-98% accuracy for speech transcription by humans, see, e.g., [20]), we also plan to store the original audio input for reference when the computer-produced text is unsatisfactory.

Another ambitious function might be called a “purpose-directed view.” When a responder is planning a complex, time-critical task, there is large potential benefit to compiling and presenting just the information he or she needs at that moment rather than requiring him or her to browse to acquire the information (e.g., “Here’s what you need to know as you get ready to enter the building.”). Similarly, an incident commander seeking overall situational awareness would benefit from access to a digest of the most important recent messages and/or messages entered by particular responders.

In the heat of an emergency, a responder might not always remember to check for new or updated information. Pushing messages (or notifications) of certain types or from certain people would ensure that critical messages are seen when they are sent. Such a capability might be implemented as a subscription mechanism to allow different responders to select the messages and message senders most relevant to their responsibilities.

Several organizations have standard formats for various types of information, e.g., a firefighter’s size-up or report of units on the scene. We plan to add standardized, but customizable, forms to facilitate the entry of such information.

While we believe a mobile application, such as the one described herein, is a useful tool for emergency response, it tends to provide a personal and local view of a disaster, so it is not optimal for providing incident commanders with a big-picture view of a large-scale disaster. We plan to integrate our application with the Golden Gate Safety Network Common Operating Picture [1] to enable an incident commander to see such a big picture and also to interact with responders using the same structured hypermedia capability as the mobile application provides.

The current Semantic Geotagging system produces a “forest” of trees of conversational threads. There are, however, potential connections across threads and even across trees, which if made available would provide access to potentially useful information that is otherwise likely to be missed. Previous research (e.g., [21, 22]) suggests that fully automated crosslinking is impractical but that it is possible to

propose reasonable links that a human user can examine for suitability. We plan to revisit these techniques in light of advances in artificial intelligence research and to prototype one or more approaches for crosslinking.

## 7.2 Expanded Testing

While our cognitive walkthroughs and usability testing have been useful in identifying and addressing basic interaction design issues, we have not yet tested the application in an authentic usage context. Such a context would likely involve multiple, geographically separated responders addressing different aspects of a complex incident. We plan to increase the realism of our testing, first by engaging multiple users in scenario-based “tabletop” exercises and eventually by deploying the system for experimental use in “real-world” disaster preparedness exercises (see, e.g., [23]).

## 7.3 New Information Sources

Finally, we recognize that professional responders typically do not have all of the information relevant to emergency response. Members of the public can often supply much useful information, especially in the early stages of an incident or when professional resources are stretched thin. That said, professional responders cannot unconditionally accept public-provided information because such information might be inaccurate or even intentionally misleading. We are planning an interview-based approach to formulate a set of heuristic rules for determining which public-provided information is likely to be trustworthy. Rather than giving members of the public access to this application, however, we plan to add a new interface layer, which will display filtered, geotagged messages from social media tools, such as Twitter, when it is activated.

## 7.4 New Uses

We believe that the applicability of an application of this sort goes far beyond emergency response, and we are considering additional uses. One intriguing possibility is to use Semantic Geotagging as a platform for citizen participatory journalism, providing a simple tool for both access and updates to a large corpus of evolving news.

# 8 Conclusion

A system of this type will be most useful in a major disaster, but it is also likely to be most difficult to get responders to use it in such a situation. In particular, rank-and-file responders, such as front-line firefighters, are, to quote one of our test subjects, “a bunch of adrenalin junkies who are going 100 miles per hour. It will be hard to get them to slow down to enter information,” although they might ultimately be consumers of such information. Command personnel might be more receptive to

using such a system; in a fire, these include the incident commander, battalion chiefs, and safety officers. These responders are likely to be older and less tech-savvy than younger ones, so ease of use and absolute reliability are paramount. It is thus imperative that the system be thoroughly tested before even experimental deployment and thoroughly proven in exercises before attempting use in an actual emergency.

Of course, focusing on senior personnel as users is counter to our initial goal of “democratizing” the development of situational awareness as an alternative to the currently prevalent command-and-control model of information collection, curation, and dissemination. Perhaps change in the direction we envision will begin with providing all responders the opportunity to contribute to situational awareness even if most do not take advantage of it; perhaps a simplified information entry interface will be required. We will continue to explore this issue as the Semantic Geotagging project progresses.

**Acknowledgments.** This work was funded, in part, by a grant from Carnegie Mellon University’s Cylab under grant DAAD19-02-1-0389 from the Army Research Office. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ARO, CMU, the U.S. Government or any of its agencies. The authors would like to thank Jeannie Stamberger, Sean Lanthier, and a number of firefighters for their contributions to this work. We would also like to thank the anonymous reviewers for their careful reviews and advice.

## References

1. Golden Gate Safety Network Common Operating Picture, <http://comopview.org/ggsn>
2. Luqman, F.B., Griss, M.L.: Leveraging Mobile Context for Effective Collaboration and Task Management in Disaster Response. In: Proc. ISCRAM, Portugal (2011)
3. Luqman, F.B., Griss, M.L.: Overseer: A Mobile Context-Aware Collaboration and Task Management System for Disaster Response. In: The 8th International Conference on Creating, UC San Diego, CA, USA (2010)
4. Cheng, H.T., Buthpitiya, S., Sun, F.T., Griss, M.L.: OmniSense: A Collaborative Sensing Framework for User Context Recognition Using Mobile Phones. ACM HotMobile (2010)
5. GeoCAM Share Project at NASA Ames, <http://geocamshare.org/>
6. GeoTwitter, <http://geotwitter.org>
7. GeoChat, <http://instedd.org/technologies/geochat/>
8. Rinner, C., Kebler, C., Andrulis, S.: The use of Web 2.0 concepts to support deliberation in spatial decision making. Computers, Environment and Urban System 32, 386–395 (2008)
9. Niko, D.L., Hwang, H., Lee, Y., Kim, C.: Integrating User-Generated Content and Spatial Data into Web GIS for Disaster History. In: Lee, R. (ed.) Computers, Networks, Systems, and Industrial Engineering 2011. SCI, vol. 365, pp. 245–255. Springer, Heidelberg (2011)
10. Roche, R., Propeck-Zimmerman, E., Mericskay, B.: GeoWeb and crisis management: issues and perspectives of volunteered geographic information. GeoJournal (June 2011)

11. Zook, M., Graham, M., Shelton, T., Gorman, S.: Volunteered Geographic Information and Crowdsourcing Disaster Relief: A Case Study of the Haitian Earthquake. *World Medical and Health Policy* V2, 7–27 (2010)
12. Maiyo, L.: Collaborative Post-Disaster Damage Mapping via Geo Web Services. MSc Thesis, International Institute for Geo-Information Science and Earth Observation (ITC), The Netherlands (2009)
13. Bareiss, R., Osgood, R.: Applying AI Models to the Design of Exploratory Hypermedia Systems. In: *Proceedings of Hypertext 1993*, pp. 94–105 (1993)
14. Ferguson, W., Bareiss, R., Birnbaum, L., Osgood, R.: ASK Systems: An Approach to the Realization of Story-Based Teachers. *Journal of the Learning Sciences* 2(1), 95–134 (1992)
15. Johnson, C., Birnbaum, L., Bareiss, R., Hinrichs, T.: War Stories: Harnessing Organizational Memories to Support Task Performance. *ACM Intelligence* 11(1), 16–31 (2000)
16. Schank, R.: Rules and Topics in Conversation. *Cognitive Science* 1(4), 421–441 (1977)
17. Anonymous, Fire Point Size-Up,  
<http://www.fireengineering.com/index/articles/display/96802/articles/fire-engineering/volume-154/issue-2/features/five-point-size-up.html>
18. Spencer, R.: The Streamlined Cognitive Walkthrough Method: Working Around Social Constraints Encountered in a Software Development Company. *CHI Letters* 2(1), 353–359 (2000)
19. Krug, S.: *Rocket Surgery Made Easy: The Do-It-Yourself Guide to Finding and Fixing Usability Problems*. New Riders Press, Berkeley (2010)
20. NIST, The History of Automatic Speech Recognition Evaluations. National Institute of Standards and Technology, NIST,  
<http://www.itl.nist.gov/iad/mig/publications/ASRhistory/index.html>
21. Cleary, C., Bareiss, R.: Practical Methods for Automating Linking in Structured Hypermedia Systems. In: *Proceedings of Hypertext 1996*, pp. 31–41 (1996)
22. Osgood, R., Bareiss, R.: Automated Index Generation for Constructing Large-Scale Conversational Hypermedia Systems. In: *Proceedings of AAAI 1993*, pp. 309–314 (1993)
23. *Golden Guardian* (2011),  
<http://www.oes.ca.gov/WebPage/oeswebsite.nsf/Content/AEC4591D507E40F3882576480064882D?OpenDocument>