Crowdsourcing geographic information for disaster response: a research frontier

Michael F. Goodchild[a]; J. Alan Glennon[a]

[a] Center for Spatial Studies, University of California, Santa Barbara, CA, USA

First published on: 20 August 2010

PLEASE SCROLL DOWN FOR ARTICLE

Crowdsourcing geographic information for disaster response: a research frontier

Michael F. Goodchild* and J. Alan Glennon

Center for Spatial Studies, University of California, Santa Barbara, CA 93106-4060, USA

Geographic data and tools are essential in all aspects of emergency management: preparedness, response, recovery, and mitigation. Geographic information created by amateur citizens, often known as volunteered geographic information, has recently provided an interesting alternative to traditional authoritative information from mapping agencies and corporations, and several recent papers have provided the beginnings of a literature on the more fundamental issues raised by this new source. Data quality is a major concern, since volunteered information is asserted and carries none of the assurances that lead to trust in officially created data. During emergencies time is the essence, and the risks associated with volunteered information are often outweighed by the benefits of its use. An example is discussed using the four wildfires that impacted the Santa Barbara area in 2007–2009, and lessons are drawn.

Keywords: emergency management; volunteered geographic information; crowdsourcing; Web 2.0; neogeography; wildfire; Santa Barbara

1. Introduction

Recent disasters have drawn attention to the vulnerability of human populations and infrastructure, and the extremely high cost of recovering from the damage they have caused. Examples include the Wenchuan earthquake of May 2009, Hurricane Katrina in September 2005, and the Indian Ocean Tsunami of December 2004. In all of these cases impacts were severe, in damage, injury, and loss of life, and were spread over large areas. In all of these cases modern technology has brought reports and images to the almost immediate attention of much of the world's population, and in the Katrina case it was possible for millions around the world to watch the events as they unfolded in near-real time. Images captured from satellites have been used to create damage assessments, and digital maps have been used to direct supplies and to guide the recovery effort, in an increasingly important application of Digital Earth.

Nevertheless it has been clear in all of these cases that the potential of such data, and of geospatial data and tools more generally, has not been realized, that the benefits of such technology have fallen far short of expectation, and that research is needed on several key issues if the situation is to improve. In many cases people living far from the impacted area have been better informed through the media than those managing and carrying out the relief effort. The impacted zone often loses power, Internet

connections, and computing capabilities, creating a donut pattern of access to relevant information. A recent report of the US National Research Council (NRC 2007) has documented these problems in detail, based on extensive discussions with responders and emergency managers, and has made a series of recommendations for improving the situation and for needed research.

> The (report)'s central conclusion is that geospatial data and tools should be an essential part of all aspects of emergency management – from planning for future events, through response and recovery, to the mitigation of future events. Yet they are rarely recognized as such, because society consistently fails to invest sufficiently in preparing for future events, however inevitable they may be. Moreover, the overwhelming concern in the immediate aftermath of an event is for food, shelter, and the saving of lives. It is widely acknowledged that maps (and all forms of geospatial data) are essential in the earliest stages of search and rescue, that evacuation planning is important, and that overhead images provide the best early source of information on damage; yet the necessary investments in resources, training, and coordination are rarely given sufficient priority either by the general public or by society's leaders. (NRC 2007, p. 2)

This paper focuses on a specific and rapidly evolving area of geospatial data and tools, a subset of social networking and user-generated web content that has been termed *volunteered geographic information* (VGI; Goodchild 2007) and that is the focus of an emerging body of research. The experience of a series of recent wildfire events in Santa Barbara is used to examine the key issues associated with VGI and its potential role in disaster management. The first major section of the paper provides a brief review of the field of VGI, a survey of its evolving literature, and its relationship to more widely recognized topics. The next section examines the specific issues of data quality in this context, drawing on research on data quality in VGI and examining its relevance to disaster management. This is then followed by a detailed discussion of the wildfire disasters, and the lessons that can be learned from them. The paper closes with a discussion of a vision for the future of VGI, and community activity more broadly, in disaster management.

2. Volunteered geographic information (VGI)

Until recently virtually all geographic information was produced in the form of maps and atlases, by mapping agencies and corporations, and dispersed as paper copies to users — researchers, consultants, and members of the general public — through a system of retail distribution. The geospatial technologies that began to appear in the 1960s did little to change this set of arrangements, since their major impacts were on the acquisition of raw data through new and more efficient instruments, its semi-automated compilation, and its use in such systems as GIS. The transition from paper-based to digital dissemination, from paper map sheets to tapes and eventually internet distribution, left most of the arrangements intact.

In the early 1990s, however, new technologies were emerging that would fundamentally change these arrangements, creating what has been termed a *post-modern* era (Goodchild *et al.* 2007) of geographic information production. First, it became possible for the average citizen to determine position accurately, without the professional expertise that had previously been limited to trained surveyors. This could be done using a simple GPS, or by finding locations using one of a number of services that became available on the Internet — conversion of street address using

a geocoding service, reading a cursor position from an accurately registered map or image provided by a service such as Google Maps, or converting a place name to coordinates using a gazetteer service.

Second, it became possible for anyone to gain the ability to make maps from acquired data, and to employ the kinds of cartographic design skills previously possessed only by trained cartographers. Google's MyMaps service, for example, allows anyone to produce a decent-looking map from custom data, and Open-StreetMap will render raw data provided by the user into a cartographically acceptable street map.

Central production of geographic information had been sustained over the decades by two factors: the need for expertise in map-making and the high capital cost of mapping equipment. By the turn of the new century both of these arguments had essentially disappeared — the cost of entry into map-making had fallen to no more than the cost of a simple PC, and the role of the expert had been replaced by GPS, mapping software, and other technologies (Goodchild 2009).

At the same time individuals with in some cases no expertise in the mapping sciences were suddenly able to perform many of the functions that had previously been the preserve of experts. The term *neogeography* was coined (Turner 2006) to describe this phenomenon, which can be defined as the breaking down of the traditional distinctions between expert and non-expert, in the specific context of the creation of geographic information, since all of the traditional forms of expertise can now be acquired through the use of technology. The default settings of mapping software, for example, now embody the recommendations of professionals, so that software rather than education or instructional manuals becomes the means by which those recommendations are disseminated and employed.

Many web sites have emerged in the past few years to encourage and facilitate the actions of neogeographers. In essence these sites make it possible for the user-generated content that increasingly dominates the Web to include digital material that satisfies the requirements of geographic information — in other words it is formed of facts about specified locations on or near the Earth's surface. Nevertheless the content is *asserted*, in contrast to the *authoritative* output of the traditional mapping agencies and corporations. It is likely not subject to any form of quality control, and the element of trust that accompanies the products of a mapping agency is typically missing. Popular sites include Flickr and its geo-referenced photographs, the OpenStreetMap project described earlier, Wikimapia and its large collection of user-described features, and numerous sites that collect georeferenced observations of plant, animal, and bird sightings. Moreover, it is increasingly common for the content of Twitter, Facebook, and many other social networking sites to be georeferenced.

VGI is closely related to the concept of crowdsourcing (Howe 2008), which has acquired two somewhat distinct meanings. On the one hand it can refer to the proposition that a group can solve a problem more effectively than an expert, despite the group's lack of relevant expertise; advocates of crowdsourcing cite many examples where this proposition appears to be true. On the other hand, and more relevant to VGI, is the notion that information obtained from a crowd of many observers is likely to be closer to the truth than information obtained from one observer. Wikipedia illustrates this meaning, since it relies on the principle that allowing people to edit an entry can produce an accurate encyclopedia, and empirical evidence appears to

support this (Giles 2005). One implication of the crowdsourcing principle is that information in which many people have an interest will be more accurate than information that is of interest only to a few, which in the case of georeferenced Wikipedia entries (a form of VGI) suggests that information about minor features in remote areas will be less accurate than information about major features in heavily populated areas. This pattern of data quality is sharply different from that of traditional mapping agencies, which use quality control procedures to guarantee uniform quality of a given product.

Academic interest in VGI dates from 2007, when a Specialist Meeting in Santa Barbara brought together an international group of experts to discuss the state of knowledge and develop a research agenda (http://www.ncgia.ucsb.edu/projects/vgi/), and was also included in a book by Scharl and Tochtermann (2007). A special issue of *GeoJournal* appeared in 2008 with a collection of research papers, and other publications and research grants have followed. Several key issues form the skeleton of a research agenda:

- What types of geographic information are most suited to acquisition through the efforts of volunteers, and how is this related to the issue of subject-matter expertise?
- What factors determine the quality of VGI, how can quality be measured, and what steps can be taken to improve it?
- What techniques can be developed for synthesizing VGI and conflating it with other data, including authoritative data, and what methods are appropriate for its analysis?
- Who creates VGI, and what are its impacts on society?

All of these questions have a broader context. For example, the question of who creates VGI is related to the broader question of volunteerism in society, and why certain people are willing to devote time and effort to tasks for which they receive no monetary reward. Quality questions are related to broader questions of crowdsourcing and collective intelligence, but also must be studied in the context of the special nature of geographic information and what is already known about measuring and modeling its quality (Guptill and Morrison 1995). Societal impacts should be addressed within the broader context of participatory GIS and the role of information in empowering individuals and communities. Budhathoki (2009, personal communication) has developed a conceptual framework for VGI research that places it within many of these broader contexts, and points to key references.

3. The quality of volunteered geographic information (VGI)

As noted above, geographic information can be defined as information linking a property to a location on or near the Earth's surface and perhaps a time. Because many of these components must be measured, and because the potential amount of such information is infinite, it is inevitable that all geographic information be subject to uncertainty. While early literature on the topic (Goodchild and Gopal 1989) emphasized accuracy, implying the existence of a truth to which a given item of information could be compared, more recently the emphasis has been on uncertainty, reflecting the impossibility of knowing the truth about many aspects of the geographic world. Research over the past two decades has focused on the sources, measurement,

and modeling of uncertainty; on its propagation into the products of analysis and modeling; and on the relative merits of the theoretical frameworks of probability and fuzzy sets (Zhang and Goodchild 2002). Standards of data quality exist for many of the products of the mapping agencies and corporations (Guptill and Morrison 1995), and data quality is an important component of metadata.

Quality is perhaps the first topic that suggests itself to anyone encountering VGI for the first time. If the providers of VGI are not experts, and if they operate under no institutional or legal frameworks, then how can one expect the results of VGI creation and publication to be accurate? Similar concerns are often expressed regarding many other types of information provided by amateurs, reflecting the deep association society makes between qualifications, institutions, and trust.

Nevertheless there are several grounds for believing that the quality of VGI can approach and even exceed that of authoritative sources. First, there is evidence that the crowdsourcing mechanism works, at least in some circumstances. In the case of Wikipedia, for example, research has shown that accuracy is as high as that of more traditional encyclopedias, according to some metrics (Giles 2005). Mention has already been made of geographic effects on the crowdsourcing mechanism, an argument that leads one to suppose that less important features, and features in little-known areas of the planet, would be less accurately described, and preliminary results from research on Wikimapia appear to bear this out. This topic is revisited below.

Second, geographic information is remarkably rich in context. Information about a location x can always be compared to other information that is likely to be available about the same location from authoritative sources, and to information about the surroundings of x. Tobler's First Law (Sui 2004) tells us that any location is likely to be similar to its surroundings, so information that appears to be inconsistent with the known properties of the location itself or of its surroundings can be subject to increased scrutiny. Wikipedia, for example, uses elaborate mechanisms for flagging and checking contributions that appear dubious, and these mechanisms are likely to be more effective for geographic information than for many other kinds. Companies that create and market street centerline databases for vehicle navigation and rely increasingly on volunteered corrections and updates, have developed elaborate, fully automated mechanisms for detecting doubtful contributions. Formalization of such methods would be a suitable topic for future VGI research, since it could lead to the ready availability of error-checking tools.

Third, discussions of geographic information quality (Guptill and Morrison 1995) emphasize the importance of completeness or *currency* as a dimension of quality — the degree to which the data are true and report all existing features *at the time of use*. Unfortunately traditional methods of map-making by government agencies, which required expert teams to travel to every part of the area and were constantly subject to funding constraints, led to lengthy delays in the updating of maps, and as a result the average map may have been years or even decades out of date by the time it was used. By contrast VGI may be produced much more quickly, and may capture changes in the landscape almost as fast as they occur. In comparing VGI with authoritative sources, therefore, one is often comparing current data with much older data. Moreover the technologies of measurement, especially measurement of position, have improved greatly over the past decade, and the expectations of users have risen accordingly. Thus a map made in 1980 at a scale of 1:24,000 may have a published positional accuracy of

12 m, but may pale in comparison with VGI acquired in 2009 using a differential GPS with a positional accuracy of 1 m.

Studies of Wikimapia conducted by my group are detecting what may be the first case of a VGI project life-cycle. Wikimapia's mantra is 'Let's describe the whole world,' which it does by enabling volunteers to identify significant features of interest on the Earth's surface, and to provide descriptions and links to other information. In effect, Wikimapia is a crowd-sourced *gazetteer*, the traditional authoritative form of place-name index (Goodchild and Hill 2008). But in contrast to gazetteers, Wikimapia has no limits to the richness of description that can be associated with a feature, allows representation of the feature's full extent instead of a single point, and accommodates both officially recognized (*gazetted*) features and unofficial ones. At the time of writing the number of entries in Wikimapia exceeded 11 million, much larger than authoritative gazetteers (see, for example, the products of the US Board on Geographic Names, http://www.geonames.usgs.gov).

Despite these benefits, once Wikimapia had reached a sufficient size and visibility, it began to attract erroneous and sometimes malicious content. The complexity of the project and the obscurity of many features made it difficult for crowdsourcing mechanisms to work to correct errors. For example, it is tempting to look for an unnamed feature in a remote part of the world and to give it a name, perhaps naming it after oneself. In the Santa Barbara area, an entry was made in mid 2009 outlining the nine-hole Ocean Meadows Golf Course, incorrectly identifying it as the 18-hole Glen Annie Golf Course, and giving the feature a detailed description that matches the latter and not the former. The distance between the two features is approximately 2 km. Although the entry was edited once in late 2009, the position had not been corrected at the time of writing.

As the volume of errors increases and crowdsourcing mechanisms fail to assure quality, the reputation of the site begins to deteriorate. Eventually the motivation to maintain the site is impacted and the site fails. It seems that Wikimapia is now entering this phase of decline and it will be interesting to see how it fares in the next few years.

Wikipedia, on the other hand, appears to have sufficient mechanisms in place to avoid this problem. Wikipedia entries are reviewed by a hierarchy of volunteers who employ well-defined criteria that are appropriate to crowdsourcing. Each entry is assessed in relation to the size of the interested crowd, as represented by the number of contributors and editors, and if that number is too small and interest fails to materialize the entry is deleted as unimportant. Moreover each entry is assessed in terms of general, permanent interest; entries describing newly coined terms (neologisms), for example, are actively discouraged. By contrast the size of the crowd interested in a Wikimapia entry describing a small feature on the Earth's surface will often be very small and Wikimapia makes no effort to use geographic context to assess the validity of entries. Thus, there may be no one sufficiently interested in the Glen Annie Golf Course, or similar examples worldwide, to correct the kind of error identified earlier, and no basis for doubting that such a golf course could exist at the identified location. The difficulties experienced by Wikimapia seem to be due at least in part to its focus on geographic features.

It is important to recognize that the quality of geographic information may have different effects, depending on the use to which the information is to be put. For any given level of uncertainty, there will be some applications for which uncertainty is not an issue and some for which it is. A 15-m error in the location of an electrical

transformer, for example, will have little effect on the operations of the utility that owns it, except when the error results in it being assigned to the wrong property and thus to land owned by someone else. Similarly a 15-m error in the position of a street may have little effect on an in-vehicle navigation system, but will be glaring if the street is superimposed on a correctly registered satellite image.

Suppose the existence of an event at a location, in other words a time-dependent item of geographic information, is critical in determining a response. For example, the event might be a chemical spill that requires evacuation of the surrounding neighborhood. Two types of errors may exist in this situation: a *false positive*, in other words a false rumor of a spill, or a *false negative*, in other words absence of information about the existence of the spill. The information is also time-critical, and a delay in its availability amounts in effect to a false negative. To reduce the chance of errors the information can be checked, by requiring independent verification or by waiting for more accurate and more authoritative information to become available. But this takes time.

In such situations decision-makers, including those residents who may need to evacuate, are faced with a choice between acting on less reliable information and waiting for more reliable information. Each has its costs, but in general acting unnecessarily, in response to a false positive, is likely to have smaller costs than not acting if the emergency turns out to be true — in other words false negatives are likely more costly and less acceptable than false positives. The next section explores these arguments in the context of a series of wildfire emergencies that impacted the Santa Barbara area in 2007–2009.

4. Volunteered geographic information (VGI) and the Santa Barbara wildfires of 2007–2009

Wildfire has always been a part of life in Southern California, but from July 2007 to May 2009 a series of four large and damaging fires occurred in rapid succession. The Zaca Fire was ignited in July 2007 and burned for 2 months, consuming 120,000 hectares largely in the Los Padres National Forest, before finally being brought under control. Although the fire threatened the city of Santa Barbara, especially should a strong wind have developed from the north, in the end no inhabited structures were destroyed and the fire never threatened the city, though the costs of fighting the fire ran into the tens of millions. The long duration of the fire created ample opportunity for the setting up of information kiosks, news releases, and other ways by which the agencies dealing with the fire communicated with the general public.

Santa Barbara lies to the south of a large wilderness area, and many homes are built in close proximity to areas covered by the local scrub, known as chaparral. The lack of wildfires in the immediate vicinity in the past 40 years had allowed large amounts of combustible fuel to accumulate. Moreover, many chaparral species contain oils that make them highly flammable, especially after a series of comparatively dry winters. In July 2008 the Gap Fire ignited in the hills immediately north of the city, threatening many homes at the western end of the urbanized area. This fire was brought under control in 7 days, but led to evacuation orders for a number of homes. The short duration of the fire and the severity of the threat meant that approaches used to inform the public during the Zaca Fire were no longer adequate

and instead numerous postings of VGI, using services such as Flickr, provided an alternative to official sources.

Traditionally the community at large has remained comparatively mute in such emergencies. Citizens rely on official agencies to manage the response, to issue and enforce evacuation orders, and to issue authoritative information. But the ease with which volunteers can create and publish geographic information, coupled with the need for rapid dissemination, has created a very different context. Agencies that are responsible for managing information are often under-funded and under-resourced, and compelled to wait while information can be verified, whereas volunteers are today equipped with digital cameras, GPS, digital maps, and numerous other resources. Multiply the resources of the average empowered citizen by the population of the city and the result is an astounding ability to create and share information.

In November 2008 the Tea Fire ignited in the hills behind Santa Barbara, and spread extremely rapidly, driven by a strong, hot Santa Ana wind from the northeast. VGI immediately began appearing on the web in the form of text reports, photographs, and video. Although search services such as Google take several days to find and catalog information from all but the most popular sites, the community had by this point learned that certain local sites and repositories were effective ways of disseminating information. These included sites run by the local newspapers and by community groups, as well as services that immediate update their own catalogs, making it possible for users to find new content quickly rather than wait for Google's robots to find and catalog it. Moreover, citizens were able to access and interpret streams of data from satellites with comparatively fine temporal and spatial resolution such as Moderate Resolution Imaging Spectroradiometer (MODIS). Several volunteers realized that by searching and compiling this flow of information and synthesizing it in map form, using services such as Google Maps, they could provide easily accessed and readily understood situation reports that were in many cases more current than maps from official sources. The Tea Fire burned for 2 days and destroyed 230 houses.

In May 2009 the Jesusita Fire ignited, again in the chaparral immediately adjacent to the city, burning for 2 days and consuming 75 houses. Several individuals and groups immediately established volunteer map sites, synthesizing the VGI and official information that was appearing constantly. For example, the officially reported perimeter of the fire was constantly updated based on reports by citizens. By the end of the emergency there were 27 of these volunteer maps online, the most popular of which had accumulated over 600,000 hits and had provided essential information about the location of the fire, evacuation orders, the locations of emergency shelters, and much other useful information (Figure 1).

In all of these activities it is clear that users were conscious of the need to balance the rapid availability of VGI with the unverified nature of much of its content. A homeowner who evacuated based on information from an online map created by a volunteer might be responding to a false positive, and by waiting for official, verified information might have avoided the need to evacuate altogether. But on the other hand the delay in acquiring information from official sources might have made the difference literally between life and death.

Several lessons can be learned from the experience of these four fires. First, due to lack of resources, the need to verify, and imperfect communication, authoritative information is much slower to appear than VGI. Citizens constitute a dense network of
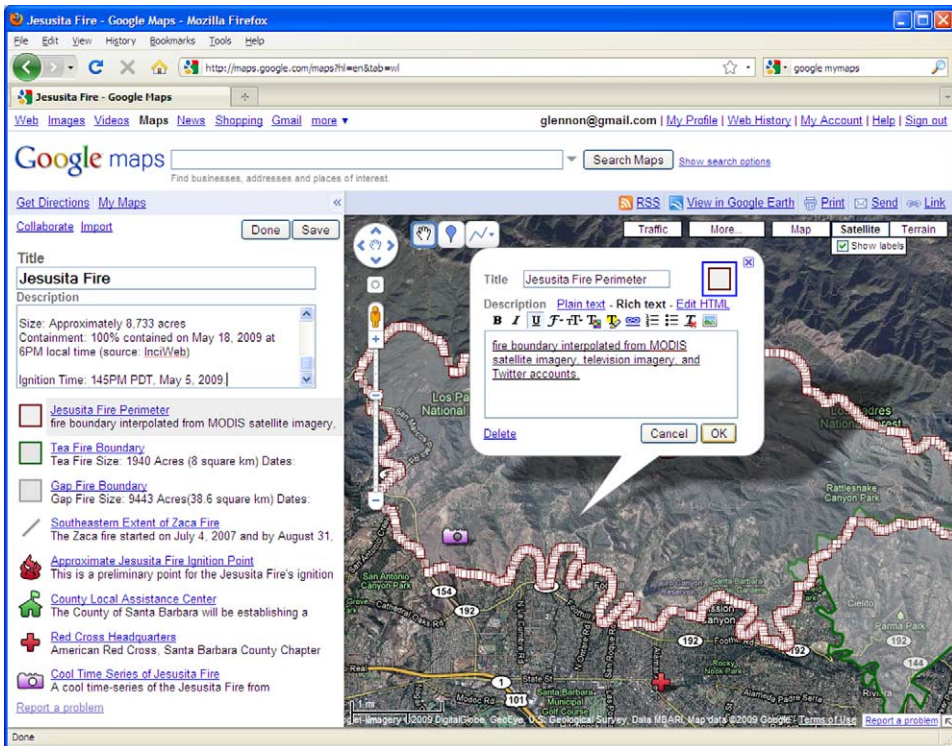
Figure 1.   Screen shot of one of the Web map sites created by amateurs during the Jesusita Fire of May 2009, showing information synthesized from a wide range of online sources, including Tweets, MODIS imagery, and news reports.

observers that is increasingly enabled with the devices and software needed to acquire, synthesize, and publish information. Second, asserted information is more prone to error, and there were many instances during the fires of false rumors being spread through Web sites. These probably led to errors on synthesized maps and to many unnecessary evacuations. Crowdsourcing mechanisms may have led to correction in some of these cases, but not all.

Third, it is clear from these experiences that the costs of acting in response to false positives are generally less than the costs of not acting in response to false negatives. Lack of information is often the origin of false negatives, as when a resident fails to receive an evacuation order, or an agency fails to verify and report a new area of fire. In essence, emergencies such as this create a time-critical need for geographic information that is quite unlike the normal, sedate pace with which geographic information was traditionally acquired, compiled, and disseminated. VGI, with its dense network of observers, is ideally suited to fill the need for near-real time information.

Fourth, in practice it is difficult during emergencies for citizens to distinguish between authoritative and asserted information. During the fire, VGI data collection and dissemination had no consistent or expected form. For example, local authorities, news media outlets, and community members all used Google MyMaps and Twitter to offer fire information, whatever its source. Each map or Twitter post was bounded by the user-interface choices of the original Google and Twitter software designers, and

thus possessed very similar visual characteristics. An elaborate system of feature-level metadata would be needed in order to distinguish information based on its source and provenance.

Finally, because of the maps' similar appearances, map popularity appears to have been bolstered by perceptions about the size and energy of a map's underlying community. During the Jesusita Fire, the most popular online VGI channels were those that provided both information and a parallel interactive discussion forum. The characteristic chaotic nature of the crowdsourced maps and their associated discussion boards may have yielded the appearance of a more active, uninterrupted channel, thus bolstering their popularity.

5. Conclusion

Agencies are inevitably stretched thin during an emergency, especially one that threatens a large community with loss of life and property. Agencies have limited staff, and limited ability to acquire and synthesize the geographic information that is vital to effective response. On the other hand, the average citizen is equipped with powers of observation, and is now empowered with the ability to georegister those observations, to transmit them through the Internet, and to synthesize them into readily understood maps and status reports. Thus the fundamental question raised by this paper is: How can society employ the eyes and ears of the general public, their eagerness to help, and their recent digital empowerment, to provide effective assistance to responders and emergency managers?

Many aspects of the data quality problem need further research. As noted earlier, an important item for further research is the formalization of rules that permit contributed geographic information to be assessed against its geographic context, and the prototyping of software tools that would implement these rules. Research is also needed to interpret what is known about trust and volunteerism in the specific context of crowdsourced geographic information, to devise appropriate mechanisms and institutions for building trust in volunteer sources.

The recent experience of the Santa Barbara fires suggests that a community can indeed contribute effectively. There are risks, of course, and more research is urgently needed to understand and minimize them. Not discussed in this paper, but also of critical importance, is the role of the citizen in those parts of the world that lie beyond the 'digital divide,' where the Internet and its services are largely unavailable. It is clear, however, that society has now entered a new era where geographic information will not only be used by all, but created by all, or at least by a dense and distributed network of observers. This leads to an entirely new vision for Digital Earth, one of a dense network of distributed, intelligent observers who are empowered to create geographic information, and particularly the types of geographic information that remote sensing and other large-scale acquisition systems are unable to produce. Protocols and institutions will be needed to ensure that the result is as reliable and useful as possible.

Notes on contributors

Michael F. Goodchild is Professor of Geography, University of California, Santa Barbara, CA, USA and Director of Center for Spatial Studies. He received his BA degree from Cambridge University in Physics in 1965 and his PhD in Geography from McMaster University in 1969. He was elected as a member of the National Academy of Sciences and Foreign Fellow of the Royal Society of Canada in 2002, and member of the American Academy of Arts and Sciences in 2006, and in 2007 he received the Prix Vautrin Lud. His current research interests include geographic information science, spatial analysis, and uncertainty in geographic data.

J. Alan Glennon is a doctoral candidate in the Department of Geography at the University of California, Santa Barbara, CA, USA and a Research Associate in its Center for Spatial Studies. He received a master's degree from Western Kentucky University. His research concerns the representation and analysis of networks in GIS, and in addition he is a leading authority on geysers.

References

Giles, J., 2005. Special report: internet encyclopaedias go head to head. *Nature*, 438, 900–901.

Goodchild, M.F., 2007. Citizens as sensors: the world of volunteered geography. *GeoJournal*, 69 (4), 211–221.

Goodchild, M.F., 2009. Neogeography and the nature of geographic expertise. *Journal of Location Based Services*, 3 (2), 82–96.

Goodchild, M.F., Fu, P., and Rich, P., 2007. Sharing geographic information: an assessment of the geospatial one-stop. *Annals of the Association of American Geographers*, 97 (2), 249–265.

Goodchild, M.F. and Gopal, S., 1989. *Accuracy of spatial databases*. London: Taylor and Francis.

Goodchild, M.F. and Hill, L.L., 2008. Introduction to digital gazetteer research. *International Journal of Geographical Information Science*, 22 (10), 1039–1044.

Guptill, S.C. and Morrison, J.L., 1995. *Elements of spatial data quality*. New York: Elsevier.

Howe, J., 2008. *Crowdsourcing: why the power of the crowd is driving the future of business*. New York: McGraw-Hill.

National Research Council, 2007. *Successful response starts with a map: improving geospatial support for disaster management*. Washington, DC: National Academies Press.

Scharl, A. and Tochtermann, K., 2007. *The geospatial web: how geobrowsers, social software, and the web 2.0 are shaping the network society*. London: Springer.

Sui, D.Z., 2004. Tobler's first law of geography: a big idea for a small world? *Annals of the Association of American Geographers*, 94 (2), 269–277.

Turner, A., 2006. *Introduction to neogeography*. Sebastopol, CA: O'Reilly.

Zhang, J-X. and Goodchild, M.F., 2002. *Uncertainty in geographical information*. New York: Taylor and Francis.