ANNALS OF GEOPHYSICS, 54, 6, 2011; doi: 10.4401/ag-5324

CITIZEN EMPOWERED SEISMOLOGY/Special Section edited by R. Bossu and P.S. Earle

Crowdsourcing earthquake damage assessment using remote sensing imagery

Luke Barrington[1], Shubharoop Ghosh[2], Marjorie Greene[3], Shay Har-Noy[1], Jay Berger[3,*], Stuart Gill[4], Albert Yu-Min Lin[1], Charles Huyck[2]

[1] *Tomnod Inc., San Diego, CA, USA*

[2] *ImageCat Inc., Long Beach, CA, USA*

[3] *Earthquake Engineering Research Institute, Oakland, CA, USA*

[4] *Global Facility for Disaster Reduction and Recovery, World Bank, Washington, D.C., USA*

Article history
Received July 23, 2011; accepted October 20, 2011.
Subject classification:
Surveys, measurements and monitoring, Instruments and techniques, Seismic risk, Seismology, Data dissemination generation or miscellaneous.

ABSTRACT

This paper describes the evolution of recent work on using crowdsourced analysis of remote sensing imagery, particularly high-resolution aerial imagery, to provide rapid, reliable assessments of damage caused by earthquakes and potentially other disasters. The initial effort examined online imagery taken after the 2008 Wenchuan, China, earthquake. A more recent response to the 2010 Haiti earthquake led to the formation of an international consortium: the Global Earth Observation Catastrophe Assessment Network (GEO-CAN). The success of GEO-CAN in contributing to the official damage assessments made by the Government of Haiti, the United Nations, and the World Bank led to further development of a web-based interface. A current initiative in Christchurch, New Zealand, is underway where remote sensing experts are analyzing satellite imagery, geotechnical engineers are marking liquefaction areas, and structural engineers are identifying building damage. The current site includes online training to improve the accuracy of the assessments and make it possible for even novice users to contribute to the crowdsourced solution. The paper discusses lessons learned from these initiatives and presents a way forward for using crowdsourced remote sensing as a tool for rapid assessment of damage caused by natural disasters around the world.

What is crowdsourcing?

The last decade has seen a proliferation of sophisticated sensors and technology capable of capturing, transferring and storing immense amounts of data, increasing the importance and demand for fast and reliable methods of analysis. Crowdsourcing, as used in our context, aims to quickly and accurately analyze large data sets by creating and leveraging a distributed network of human analysts. For example, after a disaster, large aerial and satellite image datasets can be split into small sections and sent to an online crowd of annotators to identify, classify, and prioritize damaged regions. This crowd can be composed of a small group of experts, the public at large, or a combination of the two.

The three interconnected stages of a successful crowdsourcing campaign are: (1) dividing the tasks into manageable components (*microtask*), (2) motivating a large user base to contribute (*crowdsource*) and (3) combining all responses of possibly varying quality into a complete solution (*consensus*). First, a massive data analysis challenge should be divided into small, manageable microtasks that are sized so that an anonymous, online user can rather quickly generate accurate data yet large enough so that users feel they are making meaningful contributions to the project. Second, the crowdsourcing application should strive to be engaging in order to attract a meaningful number of users and to encourage them to produce accurate contributions. Finally, after the crowd's input has been gathered, it is important to have a back-end infrastructure that distinguishes the quality of the contributions of the various users and combines them to produce a final, reliable solution. This can be as simple as averaging user contributions or as sophisticated as, for example, performing iterative maximum likelihood classification of users [Welinder et al. 2010].

Crowdsourcing methods have been applied to numerous fields where the so-called "wisdom of the crowds" provides insight beyond what individual experts can offer. Participants are motivated through small monetary rewards [Amazon.com Inc. 2011], or, even better, to contribute for free if tasks are disguised as fun games, if they appeal to scientific altruism, or require access to a service of interest. This distributed human computation has been applied, for example, to categorize galaxies [Galaxy Zoo Team 2011], fold proteins [Cooper et al. 2010], transcribe books [von Ahn et al. 2008], search for tombs [National

Geographic Society 2011] and apply descriptive labels to images [von Ahn and Dabbish 2004], web pages [von Ahn 2006] and music [Law and von Ahn 2009]. In the last several years, with the advancement of broadband connectivity and social networks, and the increased use for citizen science and scientific problem-solving, crowdsourcing applications have risen in popularity and in efficacy.

Application of crowdsourcing in emergency management

Information is crucial for effective response to disasters [Huyck 2005, National Research Council of the National Academies 2007]. The timeliness, accuracy and reliability of information provided to emergency managers and responders directly affects the quality of decisions that steer crisis response [Mehrotra et al. 2003]. Rapidly assessing critical information such as the area of impact, population at risk, damage distribution, economic loss, and potential areas where search and rescue missions are likely to be required has been identified as a priority in catastrophe response [Durham et al. 2008]. Loss estimation tools can give a preliminary indication of potentially problematic regions, but are hindered by data limitations and uncertainty [Eguchi and Seligson 2008]. Remote sensing technologies offer a true picture of events on the ground in the chaos following an event when assessing damage and determining needs can be problematic. However, limited resources and conflicting priorities complicate processing [Huyck and Adams 2002]. Automated damage detection techniques with remote sensing hold promise, but are still in the preliminary stage [Eguchi and Mansouri 2005, Eguchi et al. 2005]. Crowdsourcing can provide a mechanism to develop actionable information and at the same time, alleviate pressures on disaster analysts following an event.

Crowdsourcing currently provides information on the order of weeks and months after considerable review and processing. With institutionalization, information from crowdsourcing could conceivably be available within hours after imagery is acquired, making it a feasible resource for response teams. Search and Rescue teams already use pre-event remote sensing data when planning operations [Thorvaldsdóttir et al. 2011]. Supplementing this information with estimates of collapsed and severely damaged buildings can be instrumental in: 1) deciding whether to deploy international Search and Rescue (SAR) teams, 2) identifying unknown pockets of damage, 3) coordinating response efforts amongst SAR teams, 4) deciding whether and where to send international aid, 5) staging and deploying resources and 6) understanding short-term housing requirements. Additionally, this information can be used in the recovery phase to quantify debris removal and assess reconstruction strategies [European Commission et al. 2010, Ghosh et al. 2011]. Crowdsourced information comes with a degree of uncertainty which can be quantified [Saito et al. 2010]. As with all information contributing to situational awareness, this uncertainty must be factored into the decision making process.

The 2008 pilot project in Wenchuan, China

In 2008, ImageCat Inc. and the Earthquake Engineering Research Institute (EERI) began development of a social networking tool for earthquake impact and damage assessment. EERI was part of an international consortium that supported ImageCat's development of the Virtual Disaster Viewer (VDV). The intent of this viewer was to provide damage and situation assessments by having an international team of expert engineers interpret satellite imagery with GPS-referenced ground photographs and videos recently collected by field teams. Led by ImageCat as VDV developer and data integrator, the consortium included, in addition to EERI, structural engineers, geotechnical experts, and social scientists from the Earthquake Engineering Field Investigation Team (EEFIT, based in the UK), the Multidisciplinary Center for Earthquake Engineering Research (MCEER), the University College London (UCL) Earthquake and People Interaction Centre (EPICentre) and the UK government Engineering and Physical Sciences Research Council (EPSRC).

VDV was considered a first-of-its-kind "social networking tool" for earthquake impact and damage assessment [Bevington et al. 2009]. The initial application of VDV was after the 2008 Wenchuan, China, earthquake. Working within a specially designed online tool developed in MS Virtual Earth, dozens of earthquake experts were assigned specific areas or 'tiles' of the affected area (the city of Ying Xiu) to review and provide damage assessment by comparing before-and-after high-resolution satellite images acquired by DigitalGlobe and Geoeye imagery companies. Initial information gathered by the engineers from the imagery included the number of collapsed, heavily damaged, and intact buildings, the number of collapsed bridges, the area affected by landslides, the length of roads obstructed by landslides, and the location and scale of humanitarian relief operations, determined by the location and number of tent camps.

The 2008 application engaged 85 engineers, geoscientists, and social scientists from the EERI and EEFIT communities. Although feedback from the volunteers about this initiative was positive, one major limitation was the relatively poor quality of the imagery available. Many of the investigators were not familiar with remote sensing data and had trouble identifying seriously damaged buildings in the imagery. However, the potential of the technology and approach was clear to everyone, and many of these experts formed the core of the volunteer group that contributed to the next-applied iteration of the platform after the Haiti earthquake.

The GEO-CAN initiative in Haiti

The Global Earth Observation Catastrophe Assessment Network (GEO-CAN) community was formed to assist the World Bank/Global Facility for Disaster Reduction and Recovery (GFDRR) in quantifying building damage using aerial and satellite imagery collected in the days following the January 12, 2010, Haiti earthquake. GEO-CAN harnessed 'crowds' of experts, managed their contributions online, and allowed critical damage assessment tasks to be completed rapidly by a distributed network. GEO-CAN initially consisted of engineers and scientists who participated in the 2008 Wenchuan, China, earthquake implementation of the VDV. ImageCat and its VDV partners, including EERI, harnessed this loosely organized network to analyze high-resolution post-earthquake satellite and aerial imagery of the Port-au-Prince area of Haiti. This imagery was collected and distributed daily to relief organizations and damage assessment analysts to aid planning of immediate response and recovery activities. The sources included the World Bank–ImageCat–Rochester Institute of Technology (WB-IC-RIT) aerial mission, aerial missions flown by Google and the National Oceanic and Atmospheric Administration (NOAA), and tremendous volumes of high-resolution satellite imagery being transferred to the public domain by DigitalGlobe and GeoEye.

A phased approach was adopted to effectively assess the remote-sensing data (World Bank/GFDRR/ImageCat, in preparation). In Phase 1, volunteers identified building damage points using high-resolution satellite imagery. In Phase 2, aerial photography from both the WB-IC-RIT and Google missions were analyzed to delineate building footprints of collapsed or very heavily damaged buildings. The 1998 version of the European Macroseismic Scale[1] (EMS) [Grünthal 1998] was used to classify damage, with damage grade 4 corresponding to buildings appearing to be very heavily damaged and damage grade 5 corresponding to buildings that appeared to have been destroyed. In addition, visual interpretation of aerial photographs was performed to classify land use for Port-au-Prince and to estimate the total square footage of buildings requiring significant repairs or reconstruction. In Phase 3, geotechnical engineers from the Geo-Engineering Extreme Events Reconnaissance (GEER) Association, an NSF-sponsored research organization, focused on earthquake geotechnical issues by identifying and delineating evidence of liquefaction associated with the earthquake. Over 600 individuals from 23 countries participated in this GEO-CAN activation; this included representation from 60 universities, 18 government agencies and non-profit organizations, and 53 private companies [Ghosh et al. 2011].

The GEO-CAN effort in Haiti used a visual-inspection-based methodology to conduct rapid damage assessment and delivery of overall damage estimates. The operational workflow consisted of synchronized viewing of pre- and post-earthquake images, taking into account the characteristics of the remote-sensing sensors (spatial, spectral, and temporal resolution) and the environmental conditions at the time of image acquisition to ensure the images were properly geo-referenced (i.e., showing the same location). A point-and-click function in Google Earth was used to place markers on the approximate centroid of buildings or to delineate the pre-earthquake footprints of damaged buildings, to rapidly assign a damage grade and confidence of assessment, and to save the results to a local file for submission to a central repository. Figure 1 shows the increase in number of buildings analyzed after a general call for volunteers was sent out. EERI sent out a blast email to its general membership and encouraged members to forward it to colleagues, on Friday, January 21, 2010.

Several procedures were undertaken to validate the damage assessment performed by the GEO-CAN community and ensure the best possible classification of building damage from remotely sensed data. These included a thorough internal review to identify errors and omissions, comparison with parallel efforts such as the one conducted by the European Commission (EC) Joint Research Centre (JRC), and a detailed quality assessment evaluation using very high-resolution oblique imagery provided by Pictometry [Saito et al. 2010] and by an engineering team led by Cambridge Architectural Research Ltd. (CAR). Portions of the data were also independently verified using field ground surveys by multilateral bodies, including the United Nations Institute for Training and Research (UNITAR)/Operational Satellite Applications Programme (UNOSAT), JRC, Centre National de l'Information Géo-Spatiale (CNIGS; representing the government of Haiti), and other U.S. reconnaissance teams (EERI, Stanford University, the Pacific Earthquake Engineering Research [PEER] center, and Bertero-Fierro-Perry Inc.).

Figure 1. Number of buildings in damage grades 4 (very heavy) and 5 (destroyed) identified by volunteers, by day.

---

[1] The EMS-98 scale includes five damage grades: 1, no visible damage; 2, minor damage; 3, moderate damage; 4, very heavy damage; and 5, destroyed.

Lessons from the GEO-CAN experience in Haiti

The success of the Haiti GEO-CAN initiative suggests the disaster management community may be on the cusp of using crowdsourced analysis of remote sensing in innovative and rewarding ways. A number of lessons emerged from this experience, many of which have been addressed in follow-on work in Christchurch, New Zealand, and which are briefly reviewed here.

1. To improve the process and make it easier for emergency managers to use the results, it is necessary for the emergency response community to understand the potential application of such crowdsourcing analyses. To share the data more easily with other communities, a single portal should be developed for both the analysts and the potential users of the data/analysis.

2. The response to EERI solicitation e-mail for Haiti was very rewarding. Not only did EERI members willingly give up their weekend to participate in the analysis, many participants forwarded the e-mail to related networks and communities, spawning further enthusiastic participation. As this was the first large-scale effort, some volunteers more clearly understood the purpose of the analyses, and some were better prepared than others to conduct the analyses.

3. GEO-CAN was created very quickly to respond to the Haiti event. To keep volunteers engaged and to solicit participation in future events, it would be helpful to institutionalize GEO-CAN (beginning with maintaining a mailing list and/or a website) so that there can be regular communication with GEO-CAN participants to keep them abreast of remote-sensing developments and applications to new disasters. It will be helpful to highlight the volunteer experiences with GEO-CAN and to work with the volunteer base so they understand the expectations for the next event – e.g., quick turn-around, perhaps weekend work – and also the pay-off, which includes using professional knowledge to contribute in a meaningful way.

4. The size of the imagery files and the capability to work with them quickly will be a major issue in positioning GEO-CAN to participate in future events. To facilitate a wider response, improvements should be made in accessing and working with the imagery and developing the infrastructure to warehouse the data and analyses for future use. Recommendations that emerged from a workshop in May 2010 [EERI 2010] to evaluate the GEO-CAN experience in Haiti documented improvements that would be useful:

- Find a host for the terabytes of data that will likely be available in future disasters is critical to the success of a crowdsourcing initiative. Using tile servers could be one way to help manage the very large data files that are generated as part of this process.
- Develop a web-based application where a user could enter with a password and have immediate access to before-and-after imagery for each grid of analysis would facilitate accurate use by volunteers over a range of remote-sensing experience. Users would just draw their polygons and select criteria for damage state, confidence level, etc., from a drop-down menu.
- Acquire targeted before-event imagery, prior to the next event(s), which could be pre-loaded in the system and perhaps pre-assigned to GEO-CAN volunteers.
- Assign the same block to several people, so that map compilers could compare the results and improve confidence levels and overall accuracy.
- Create a training program that each participant views or listens to before being assigned a grid. After watching a training video, participants take a practice ‘quiz’ where they are asked to analyze a small grid. In order to continue (e.g., be assigned a grid, become a member of the GEO-CAN network), participants would need to ‘pass’ this quiz by successfully completing a certain percentage of the analysis. After analyzing approximately six grids, participants would be invited to watch and listen to a two-minute refresher video, where the most common errors and misunderstandings would be reviewed.
- Provide more detailed guidance on how to determine artifacts in the imagery that are associated with damage, such as shadows, blow-outs, skewed buildings, obvious changes in elevation, amount of debris around building, and color of the debris field.
- Offer more guidance on the EMS-98 scale by including additional photos and cartoons to illustrate damage grades 4 and 5, guidance on what to do if there are multiple damage states within a building, and specific guidance on how to indicate uncertainty with the analysis of each building.
- Include an appendix of common errors.

5. Understanding the cultural context for the remote-sensing analysis and providing guidance on this context would be helpful in future events. For example, in Haiti it would have been helpful to inform the analyst on what is unique about Haitian construction practices.

6. Given the extensive field reconnaissance conducted in Haiti, some of which included ground-truthing of the aerial imagery, examples could be developed in a future training program to present together the remote-sensing imagery and the photograph of the building from the ground. In Haiti there are at least 40,000 buildings that have been evaluated by the United Nations Office for Project Services (UNOPS) – each building has a photo and its global positioning system (GPS) coordinates [Lallemant 2010]. These could be used for the examples.

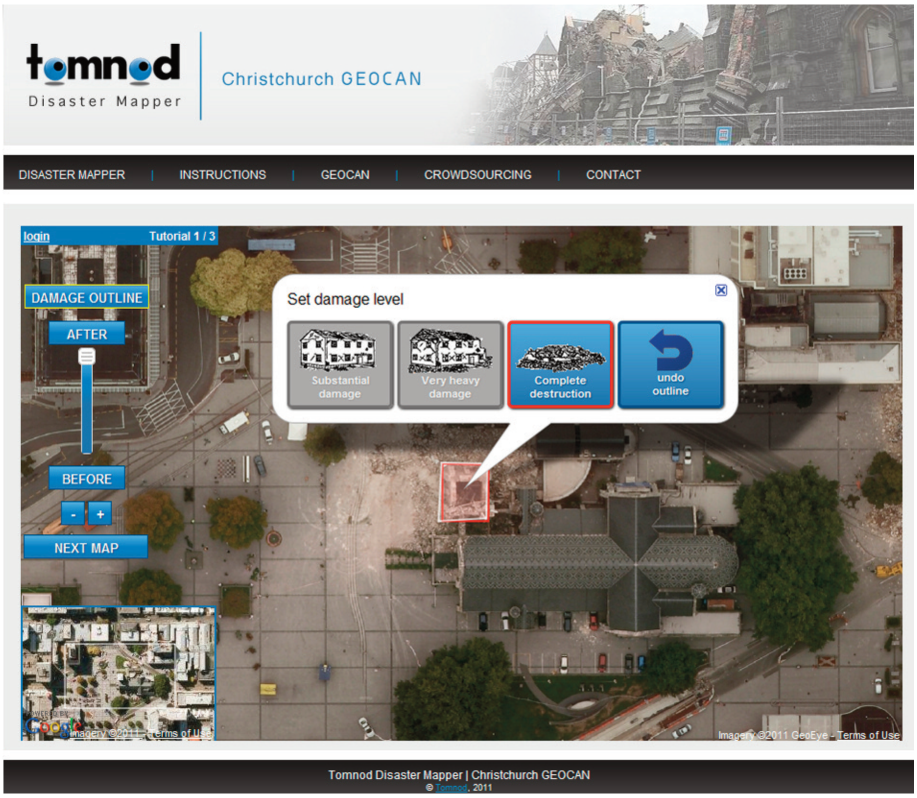

Figure 2. The Disaster Mapper developed by Tomnod, Inc for GEO-CAN effort with the February 2011 earthquake in Christchurch, New Zealand.

The way forward: GEO-CAN in Christchurch

The past efforts at crowdsourcing disaster assessment have done an exceptional job at engaging a motivated team of researchers, engineers, and scientists to label large datasets in a short amount of time. Looking towards the future, the GEO-CAN initiative is being expanded with the goal of making crowdsourced damage assessments of disaster areas faster and more accurate. To accomplish this, researchers with Tomnod, Inc. have joined the GEO-CAN initiative and have been working on the development of an improved interface, with other GEO-CAN partners. The improved disaster mapper (www.tomnod.com/geocan) runs in a desktop web browser and invites an online crowd to view and annotate pre- and post- aerial and satellite imagery. The current effort is using imagery provided after the Christchurch, New Zealand, earthquake, where users are asked to delineate building footprints of collapsed or very heavily damaged buildings (Figure 2). The interface has been improved to make it simpler and more intuitive to use, so that more of the user's time is spent analyzing data versus learning how to use the interface. Using a simple interface that runs in a web browser, rather than an "experts-only" geographic information system (GIS) platform, opens the initiative to a larger group of untrained analysts drawn from the general Internet public. For example, an integrated training module walks the user through the interface without requiring prior experience or the reading of any instruction manuals.

To accommodate an anticipated high volume of users and large amount of data transfer after the next major event, the system has been built on a scalable cloud-computing infrastructure. The intent is that immediately following the next major earthquake, new imagery can be quickly transferred to the cloud where it is processed (tiled, enhanced, pan sharpened, etc.) and made available to a crowd of users. Individual image tiles would be prioritized (e.g., focusing on areas of high building density, important infrastructure or damage centers) and distributed for viewing and analysis. As new data are collected, priority areas can be changed, both automatically and with administrator supervision, to ensure strategic and comprehensive coverage of the area under investigation. In addition, every crowd click, tag, label, annotation, outline or comment is referenced to the corresponding image tile and stored in a GIS-enabled back-end database. Furthermore, the Disaster Mapper tracks available user information (e.g., profession, organization, Facebook profile, etc.), making it possible to distinguish between analyses produced by various groups, such as earthquake engineers, academics, industry professionals, or possibly members of the public at large. Many of these features are being tested with the current initiative in Christchurch, New Zealand.

The final step in a crowdsourcing workflow involves merging all the inputs from a crowd of hundreds or even thousands of image analysts with varying skill levels, professional backgrounds, and motivations, into a final 'product' that constitutes a complete and reliable solution to the remote-sensing analysis challenge (e.g., mapping all the damage caused by the earthquake). Crowdsourced data, summarizing damaged and destroyed buildings in Port-au-

Prince fed into the Preliminary Damage and Needs Assessment prepared by the Government of Haiti and several international organizations. Data coming from the current Christchurch effort will be given to the Government of New Zealand, with preliminary results highlighted in Figure 3. The hope is that over time the crowd-sourced product or map of damage will become a reliable source of preliminary or additional data for governments and researchers. To get to that point the crowd performing such analyses needs to grow; although an individual user may be untrained, inexperienced, or even malicious, when crowds of users all agree on an analysis, it is likely to be reliable.

To date, over 200 users have analyzed more than 77 km^2 of the Christchurch central business district (CBD), contributing over 1,400 building outlines and damage assessments.

Several tools that have been prototyped for Christchurch can help GEO-CAN sift through the crowdsourced data and identify consensus more easily in future efforts. One technique in identifying consensus is visual clustering: by presenting all the collected data as a geo-referenced image overlay, areas of high consensus (i.e., lots of tags, overlapping outlines, etc.) pop out to the administrator, who can quickly zoom in on areas of interest. Figure 3 illustrates this by overlaying a subset of the more than 1,400 crowd- contributed building outlines on a small area of the Christchurch CBD where the areas of most destruction are readily apparent. A more quantitative approach is to perform statistical clustering on the crowd's data. For example, nearest-neighbor clustering of crowdsourced points identifies the most-tagged regions, or geometrical intersection of building footprints prioritizes the most-likely-to-be-damaged structures. These statistical methods can also take account of the 'unseen' data. For example, a building that has been examined by ten users but only outlined by two users is less likely to be of interest than another that was viewed by two users and outlined by both. Using this approach, Figure 3 also highlights three of the areas of maximum crowd consensus where it is clear that there has been significant destruction.

Future studies and conclusions

In future work, it may be possible to further improve the combination of crowdsourced data by including user-specific information in the consensus calculation. This can be explicit, for example, by assigning a higher weight to contributions by

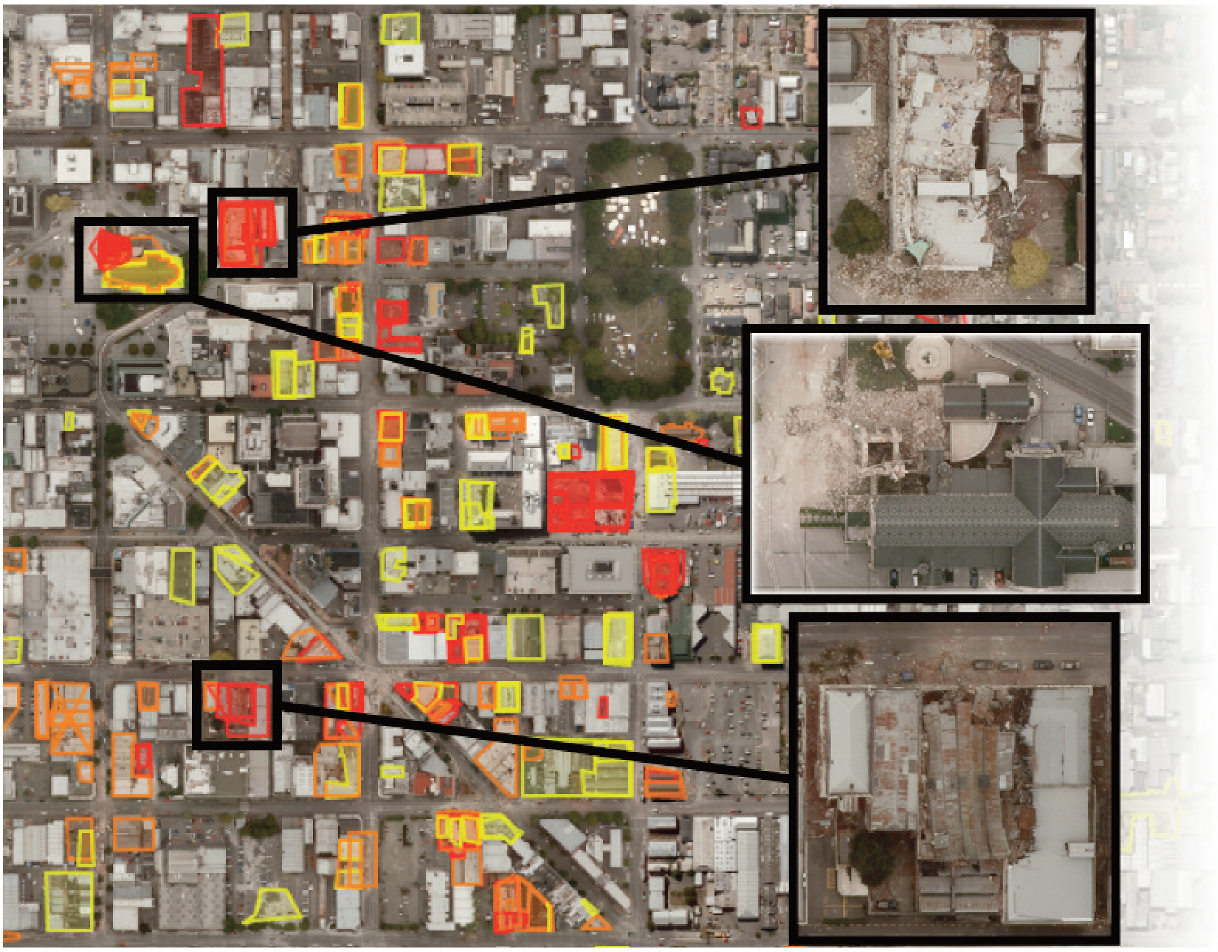

Figure 3. Sample of crowdsourced damage map of Christchurch, New Zealand. Each polygon was drawn by a crowd contributor and marked with a damage level (**red**: complete destruction, **orange**: very heavy damage, **yellow**: substantial damage). Three locations of crowd consensus are highlighted on the right, revealing reliable damage assessments.

users from professional earthquake or imagery analysis communities. Alternatively, user accuracy can be implicitly derived from the data by examining how an individual user has agreed with the crowd's consensus on previous microtasks (including tasks where the 'correct' solution is known in advance). It would then be possible to gauge their reliability in future analyses [e.g., Whitehill et al. 2009, Welinder et al. 2010]. The improved GEO-CAN interface that Tomnod has developed deploys visual, statistical and user-driven crowd analytics algorithms to appropriately weight contributions from multiple users and increase the reliability of the final, combined analysis. As an extra step, in the future it may also be possible to use crowd-derived consensus as reliable examples to train machine-learning algorithms to automate future analysis [e.g., Snow et al. 2008, Vijayanarasimhan and Grauman 2011]. A hybrid of human perceptual expertise that guides the development of automatic analysis methods potentially presents a scalable, reliable approach to extracting information from massive amounts of remotely sensed imagery.

To recruit a larger crowd of volunteer damage assessors in future events, GEO-CAN also leverage users' Facebook and Twitter communities to encourage greater participation. Users can share the project as a whole with their social network or share a particular image region and request assistance with it. The goal is to engage a large audience to analyze data, and then to combine the analyses with segmentation of user groups and back-end annotation filtering algorithms, to allow for fast, scalable, and accurate analysis of massive amounts of remote sensing imagery.

Harnessing the power of crowdsourcing as a mechanism for damage assessment to improve disaster response was clearly demonstrated by the GEO-CAN initiative. A full validation of the reliability of the obtained crowdsourced data against ground truth information is beyond the scope of this paper, however, this type of study could lend additional insight into the strengths and weaknesses of the approach. Furthermore, several operational and methodological challenges exist in the current GEO-CAN structure and workflow. In its current form, the GEO-CAN community is an informal network of individuals, agencies, and organizations using imagery and expertise towards a united common goal of disaster response. The disaster management profession, which includes the engineering and scientific communities, would benefit from a more formal structure for GEO-CAN. The benefits of GEO-CAN would extend to society as a whole by improving the ability to quickly understand the scale of disasters, leading to faster and more effective emergency response. Developmental organizations and professional engineering and scientific societies would also benefit from the GEO-CAN pre-event baseline analysis and post-event preliminary damage assessments. GEO-CAN has the potential to improve the understanding of vulnerability (pre-event) and post-event data analysis. It can also advance the Learning from Earthquakes Program, an EERI-initiated program involving multi-disciplinary teams of researchers investigating the damaging effects of earthquakes from the field, the results of which are viewed by many worldwide as a primary source of earthquake reconnaissance information [EERI 2011]. Finally, GEO-CAN would benefit scientists in the global research community by providing a unique opportunity to channel their expertise towards humanitarian applications and share analyses online.

In conclusion, we have presented previous and current state-of-the-art on the emerging use of crowdsourced analysis of remote sensing imagery as a tool to aid in disaster assessment and management. We showed the evolution of crowdsourced platforms from a prototype for Wenchuan, China in 2008, through Google-Earth-based mapping and the foundation of GEO-CAN for Haiti in 2010, to the current, web-based platform deployed for Christchurch, New Zealand, in 2011. As these platforms and technologies evolve, it becomes easier to engage large online audiences to assist in rapid analysis of disaster imagery and more efficient to collect and combine reliable results on the backend. Moving forward, widespread online dissemination of crowdsourced tasks and intelligent aggregation of crowd contributions offer a new source of insight in response to earthquakes and other natural disasters and present a platform to engage the world to assist with global challenges.

References

Amazon.com Inc. (2011). Amazon Mechanical Turk; URL: https://www.mturk.com.

Bevington, J., B. Adams, E. Verrucci, P. Amyx, C. Huyck and R. Eguchi (2009). Distributed Information Sharing for Disaster Response and Recovery: www.VirtualDisasterViewer.com. DHS Workshop on Emergency Management: Incident, Resource, and Supply Chain Management (EMWS09).

Cooper, S., F. Khatib, A. Treuille, J. Barbero, J. Lee, M. Beenen, A. Leaver-Fay, D. Baker, Z. Popovic and Foldit Players (2010). Predicting protein structures with a multiplayer online game, Nature, 466, 756-760.

European Commission - EC, Joint Research Centre - JRC, United Nations Institute for Training and Research - UNITAR, Operational Satellite Applications Programme - UNOSAT, World Bank Global Facility for Disaster Reduction and Recovery - GFDRR, and Centre National d'Information Géo-Spatial - CNIGS (2010). Building Damage Assessment Report Haiti earthquake 12 January 2010 Post Disaster Needs Assessment and Recovery Framework (PDNA), Report to the Haitian Government.

Durham, T.S., P. Johari and D. Bausch (2008). Strategic Directions in Seismic Modeling: HAZUS® Development and Current Applications for Catastrophe Planning, In: A. Bostrom, S. French and S. Gottlieb (eds.), Risk Assessment, Modeling and Decision Support: Strategic Directions, Springer-Verlag, Berlin Heidelberg, 101-116.

EERI, Earthquake Engineering Research Institute (2010). Workshop Report: Remote Sensing and the GEO-CAN Community: Lessons from Haiti and Recommendations for the Future; URL: http://www.eqclearinghouse.org/20100112-haiti/wp-content/uploads/2010/02/EERI-Workshop-Report-6-30-10_FINAL_APPENDIX.pdf.

EERI (2011). Introduction to LFE; URL: http://www.eeri.org/site/lfe-introduction.

Eguchi, R. and B. Mansouri (2005). Use of Remote Sensing Technologies for Building Damage Assessment After the 2033 Bam, Iran Earthquake –Preface to Remote Sensing Papers, Earthq. Spectra, 21 (S1), S207-S212.

Eguchi, R., C. Huyck and B. Adams (2005). An Urban Damage Scale Based on Satellite and Airborne Imagery, In: Proceedings of the 3rd International Workshop on Remote Sensing for Post-Disaster response (Chiba University, Japan, September 12-13).

Eguchi, R. and H. Seligson (2008). Loss Estimation Models and Metrics, In: A. Bostrom, S. French and S. Gottlieb (eds.), Risk Assessment, Modeling and Decision Support: Strategic Directions, Springer-Verlag, Berlin Heidelberg, 135-170.

Galaxy Zoo Team (2011). Galaxy Zoo; URL: http://www.galaxyzoo.org.

Ghosh, S., C.K. Huyck, M. Greene, S.P. Gill, J. Bevington, W. Svekla, R. DesRoches and R.T. Eguchi (2011). Crowdsourcing for rapid damage assessment: The global earth observation catastrophe assessment network (GEO-CAN), Earthq. Spectra, 27 (S1), S179-S198.

Grünthal, G., ed. (1998). European Macroseismic Scale 1998 (EMS-98), Cahiers du Centre Européen de Géodynamique et de Séismologie 15, Centre Européen de Géodynamique et de Séismologie, Luxembourg, 99 pp.

Huyck, C.K. and B.J. Adams (2002). Emergency Response in the Wake of the World Trade Center Attack: The Remote Sensing Perspective, MCEER Special Report Series on Engineering and Organizational Issues Related to the World Trade Center Terrorist Attack, Vol. 3, Multidisciplinary Center for Earthquake Engineering Research.

Huyck, C. (2005). Suggestions for the effective use of remote sensing data in emergency management, NRC Planning for Catastrophe Study Workshop on Geospatial Information for Disaster Management, National Academy of Sciences, Washington, D.C.

Lallemant, D. (2010). Structural Building Evaluation Project, presentation at ImageCat/EERI Workshop on Remote Sensing and the GEO-CAN Community: Lessons from Haiti and Recommendations for the Future (May 2010).

Law, E. and L. von Ahn (2009). Input-agreement: A new mechanism for collecting data using human computation games, In: 27th International Conference on Human Factors in Computing Systems (ACM CHI).

Mehrotra, S., C. Butts, D. Kalashnikov, N. Venkatasubramanian, R. Rao, G. Chockalingam, R. Eguchi, B. Adams and C. Huyck (2003). Project rescue: challenges in responding to the unexpected, In: S. Santini and R. Schettini (eds.), Internet Imaging V: Proceedings of 16th SPIE Annual Symposium, Vol. 5304, 179-192.

National Geographic Society (2011), Field Expedition Mongolia; URL: http://exploration.nationalgeographic.com.

National Research Council of the National Academies (2007). Successful Response Starts With a Map: Improving Geospatial Support for Disaster Management, The National Academies Press.

Saito, K., R. Spence, E. Booth, M. Madabhushi, R. Eguchi and S. Gill (2010). Damage Assessment of Port-au-Prince using Pictometry, In: Proceedings of the 8th International Workshop on Remote Sensing and Disaster Management (Tokyo, Japan, September 30 - October 1).

Snow, R., B. O'Connor, D. Jurafsky and A.Y. Ng (2008). Cheap and fast--but is it good? Evaluating non-expert annotations for natural language tasks, In: 13th Conference on Empirical Methods in Natural Language Processing (EMNLP).

Thorvaldsdóttir, S., E. Birgisson and R. Sigbjornsson (2011). Interactive on-site and remote damage assessment for urban search and rescue, Earthq. Spectra, 27 (S1), S239-S250.

Vijayanarasimhan, S. and K. Grauman (2011). Cost-Sensitive Active Visual Category Learning, Int. J. Comput. Vision, 91 (1), 24-44.

von Ahn, L. and L. Dabbish (2004). Labeling images with a computer game, In: 22nd International Conference on Human Factors in Computing Systems (ACM CHI).

von Ahn, L. (2006). Games with a purpose, IEEE Comput. Mag., 39 (6), 92-94.

von Ahn, L., B. Maurer, C. McMillen, D. Abraham and M. Blum (2008). re-CAPTCHA: Human-Based Character Recognition via Web Security Measures, Science, 321 (5895), 1465-1468.

Welinder, P., S. Branson, S. Belongie and P. Perona (2010). The Multidimensional Wisdom of Crowds, In: 24th Conference on Neural Information Processing Systems (NIPS).

Whitehill, J., P. Ruvolo, J. Bergsma, T. Wu and J. Movellan (2009). Whose Vote Should Count More: Optimal Integration of Labels from Labelers of Unknown Expertise, In: 23rd Conference on Neural Information Processing Systems (NIPS).

World Bank/GFDRR/ImageCat (in preparation). Post-Disaster Building Damage Assessment using Satellite and Aerial Imagery Interpretation, Field Verification and Modeling: 2010 Haiti Earthquake Final Report.

---

*Corresponding author: Jay Berger,
Earthquake Engineering Research Institute, Oakland, CA, USA;
email: jberger@eeri.org.