HYBRID COMMUNITY PARTICIPATION IN CROWDSOURCED EARLY WARNING SYSTEMS

Nick LaLone
Penn State University
nick.lalone@gmail.com

Andrea H. Tapia
Penn State University
atapia@ist.psu.edu

Nathan A. Case
NASA
nathan.a.case@nasa.gov

Elizabeth MacDonald
NASA
eliz.macdonald@gmail.com

Michelle Hall
Science Education Solutions
hall@scieds.com

Matt Heavner
Los Alamos National Laboratory
matt.heavner@gmail.com

ABSTRACT

In this paper we present Aurorasaurus: a website, a mobile application, and a citizen science initiative that allows a community of users to report and verify sightings of the Aurora Borealis. Through ad-hoc data indirectly offered through social media, a community of citizen scientists verify sightings of the Aurora Borealis. These verified data are tested against currently existing aurora-forecasting models. The insights these data provide are transformed into map and text-based forms. In addition, notifications are sent to interested participants in a timely manner. This is a design test-bed for an early warning system (EWS) that is capable of detecting and communicating the earliest signs of disaster to community members in near real time. Most importantly, this system incorporates community participation in improving the quality of data mined from Twitter and direct community contributions.

Keywords

Aurorasaurus, Early Warning System, Crowdsourcing, Citizen Science, Twitter.

INTRODUCTION

Among the fields that study information and communication during crisis and emergencies, there has been a great rush to make data originating from the crowd useful during a response. The approach has often been: (1) collect data indirectly from the crowd by scraping it from social media sources, (2) Aggregate and process it, and (3) serve it to responders. According to MacEachren, et. al. (2011) research has focused more on the challenges of extracting action items and location information from social media feeds. Ironically, most solutions to this problem do not involve direct contributions from the public itself.

In the larger picture, the goal of this research is to put the public back in public informatics. Automated monitoring of social media data may provide some interesting insights; however, the great majority of scraped data is irrelevant to the needs of the public or responders. The computational power used to identify the relevant and useful elements is awesome though perhaps unnecessary. Unnecessary because community members themselves are in a unique position to view the data of computational processes and with a trained eye localize results to community needs. Data contributed directly by community members toward their own shared community information systems has far greater relevance, accuracy and utility to the public it serves.

Members of an affected public are motivated to work toward their own safety, security, resilience and capacity. Members of the public are more likely to trust data from those they know and offer help to those they know, or at least those identifying as from the same location or group. A motivated and trained public can both contribute directly to a shared information system and help process data scraped from a larger data ecosystem into a hybrid system that has the potential to be more accurate, timely and useful to the same community.

In the remainder of this paper we present a community-computational hybrid system through which community members can both directly contribute and help evaluate scraped data as part of a shared early notification system. Our system, Aurorasaurus, is a hybrid Early Warning System (EWS) for aurora borealis sightings. Aurorasaurus is a website, mobile application, and citizen science initiative that allows a community of users to help scientists better forecast the location and viewing conditions of the Aurora Borealis.

The vision of Aurorasaurus leverages this to gather real-time data about the aurora in two ways, by direct entry into an online (or mobile) form, and by continuous scanning of Twitter for tweets about sightings. Combined with data from Earth-based and satellite observatories, this allows Aurorasaurus to offer near real-time predictions of auroral activity in both text and map form, likely with much greater accuracy and timeliness than the current state of the art. In the middle latitudes sightings of the Aurora Borealis are rare and more difficult to forecast or predict, much like an emergency or crisis. The similarities between a natural disaster occurrence and auroral occurrence offer a chance for researchers to test elements of an Early Warning System.

A COMMUNITY-COMPUTATIONAL HYBRID

Around 68% of all information system (IS) development and deployments fail (Krigsman, 2009). Failure can be outright non functionality or non-use rising to failure of process-taking over 180 percent of target time to deliver; failure of resource control--consuming in excess of 160 percent of estimated budget; or failure of functional match-delivering under 70% of the target required functionality (see: Lyytinen & Hirshiem, 1987; Sauer, 1993). This failure is keenly felt in the public sector where resources and time are even more limited and functionality may be key to improving lives and outcomes of vulnerable populations. IS failure in the public setting is less tolerated and more difficult from which to recover. Therefore, when developing public information systems it is especially important to assure the greatest possible chances for success.

It has been suggested that full engagement with the community for which the IS is being developed and deployed is essential to bringing about this success. In terms of developing and deploying IS for emergency management, several researchers have called for systematic, community integration and guidance of IS projects. (see Berg, 2001; Franco, et. al 2007; Kuziemsky, et. al., 2012; and O'Sullivan et al., 2012). Ahmed, et al. (2012) suggest, "…omitting systematic, process-based community guidance of [IS for emergency management] technical solutions is a prospect that is, at best, doomed to expensive and often predictable failure. At worst, some of the solutions developed have the potential to do more harm than good." There have been several notable collaborative IS for emergency management design and development projects in which local community non-profits, local governments, and local developers have engaged to co-create a system (e.g., Landgren, 2010; Troy, et al., 2008; Franco, et. al 2007; Kuziemsky, et. al., 2012; and O'Sullivan et al., 2012). According to Kuziemsky, et. al. (2012), public engagement is an essential precursor to designing an IS for emergency management and it must be actively pursued before any technology is designed. O'Sullivan, et, al. (2012) state, "the awareness that evolves from a truly collaborative, participative user-engagement process can be empowering, and can stimulate solution-oriented, creative thinking and innovation."

The question that must be asked is, why stop community engagement with the deployment of an information system? If community engagement is essential to the successful design, development and deployment of IS for emergency management, then the public should also play a more integral role in collecting data, inputting data, processing and evaluating data for the system. The public could also enjoy the direct benefits of their labor, as recipients of the outputs of the system. The system could be open so that members of the public can 'look under the hood' and see the system in action and witness how their contributions of time, labor, and data create the value offered to the community via the system. In other words, the public should be more fully integrated in the continuous functioning of the system.

There is evidence that the public could play essential roles as part of such systems. While the strides in the area of big data analytics have met some of the demand, many problems have been identified that show the limitations of computational solutions. There are problems that only human cognition has been able to solve. Perhaps the most famous example of this is CAPTCHA (for an automated public turing test to tell computers and humans apart see von Ahn, et.al, 2004). A CAPTCHA is a program that protects websites against bots by generating and grading tests that humans can pass but current computer programs cannot. For example, humans can read distorted or variable text shapes or melded text, but current computer programs cannot.

Several such human-required problems arise in the analysis of social media data for emergency response. Even with the very best analytical tools social media data defies natural language processing, sentiment analysis, data mining and machine learning interventions. The volume of social media data can be narrowed and focused, but not to the necessary finite point of utility by a community or response organization. A 'community-computational hybrid system is necessary to make the output of analytic approaches useful to communities and responders. For analytic outputs to be useful to community members they must be localized to community needs. For analytic outputs to be used in decision-making in emergency response requires a precise fit in terms of confidence in the analysis, in the data sources and in the format of the information. Human processing of data can meet the needs of communities and responders during a crisis in ways that computers cannot accomplish alone.

In addition, during a crisis, being local to the community affected matters. Starbird, et. al (2012) state, "People who are on the ground are uniquely positioned to share information that may not yet be available elsewhere in the information space. Additionally, locals may have knowledge about geographic or cultural features of the affected area that could be useful to those responding from outside the area." However, in her earlier study she also noted that during a crisis only a small portion of tweets contain information from local Twitter users (2010).

Several researchers have examined the role of community members in filtering or improving data. Collaborative filtering is a technique for using the actions of a large number of people within an interaction space to filter information produced by that same group (Malone et al., 2009). Starbird and Palen (2012) reported that certain characteristics of crowd behavior could act as a collaborative filter for identifying people tweeting from the ground during mass disruption events in Egypt. Starbird, Munzy and Palen (2012) also examined the Occupy Wall Street movement to demonstrate the power of crowd-based action. For example, Mendoza et al. (2010) report evidence that the social media community can collaboratively act to identify bad information. Studying the propagation of rumors through the Twitterverse in the wake of the Chile Earthquake in 2010, they found that tweets containing false information were more likely to be challenged by other Twitterers.

In addition there is growing informal digital response community (Starbird and Palen 2011) that recognizes and fills emergency needs using online tools like social media. These volunteers comprise individual citizens as well as more formalized groups, organizations, and communities. Digital volunteers have contributed to emergency response efforts by monitoring and responding to social media, creating and updating digital maps, and helping to coordinate relief and recovery (St. Denis, Hughes, and Palen 2012; Boehmer 2010; Norheim-Hagtun and Meier 2010; Starbird and Palen 2013; van Gorp 2014). Each of these groups used a combination of crowdsourcing and computational techniques to collect relevant social media data, process and categorize the data, and plot the data on a map for the responding organization.

CROWSOURCING WARNINGS

An early warning system (EWS) is "the set of capacities needed to generate and disseminate timely and meaningful warning information to enable individuals, communities and organizations threatened by a hazard to prepare and to act appropriately and in sufficient time to reduce the possibility of harm or loss" (Othman & Beydoun, 2010)

To increase crisis detection capabilities, several scholars have turned to harvesting data from social media as a means to provide near real-time evidence of impending disaster as part of an early warning system. As the number and frequency of mentions of certain hazard-related terms increase, thresholds are crossed and warnings triggered. While there has been some early success in using social media data to forecast influenza rates, stock market fluctuations, and movie box office sales in near real time, (see Chen et al. 2010; Achrekar et al., 2011; Wolfram 2010; O'Connor et al., 2010; Wakamiya et al., 2011; Lampos et al., 2010; Sakaki et al., 2010) there have been few instances in which social media data has been intentionally used as part of an early warning system for a natural disaster.

To argue for the efficacy of Twitter as part of an early warning system, several scholars have demonstrated its limited power to forecast (see Arias, et al., 2012). Twitter has been used for real-time notifications such as large-scale fire emergencies and downtime on services provided by content providers (Motoyama, et al., 2010).

Most promising is the work by Paul Earle (2012) who developed algorithms that automatically detect large increases in the usage of the term Earthquake, in multiple languages, on Twitter. The algorithm found that TED (Twitter Earthquake Detection) often detected earthquakes in less than 1 minute; in fact, 75 percent of all detections occurred within 2 minutes, which is much faster than the time period of 2 to 20 minutes that traditional sensing methods require (Earle, et al., 2012). While it did detect earthquakes, and quickly, it missed many more during the same period. The US Geological Survey has also started the Citizen Seismology Project (Young, et al., 2013) in which these new detection algorithms are incorporated into citizen science efforts benefitting both scientific knowledge concerning earthquakes and early detection programs.

Ginsberg et al. (2008) showed that the frequency of Google search terms can be used to build a linear model which accurately estimates influenza incidence. Chen, et al. introduced a continuous data collection engine which combines the detection and prediction capability of social networks in discovering real world flu trends (Chen et al., 2010; Achrekar et al., 2011). Wolfram (2010) attempted to predict the price of NASDAQ stock quotes by using Twitter as an additional source of information. Using similar techniques, other researchers have used Twitter data for prediction within presidential polls (O'Connor, et al., 2010), TV ratings (Wakamiya, et al., 2011), and influenza rates (Lampos, et al., 2010). Finally, Sakaki, et al. (2010) used built an autonomous earthquake reporting system in Japan using Twitter users as sensors.

One project now run by the European–Mediterranean Seismological Centre (EMSC) seeks to both detect earthquakes using new detection methods and display these on their website in real time (see: Young, et al., 2013). However, this website gains more traffic after an earthquake has already occurred and is based on visits (pull) rather than alerts (push) interactions with the user.

A project that is closest to our model comes out of the USGS (United States Geological Survey) initiated in 2006 called the Earthquake Notification Service (ENS https://sslearthquake.usgs.gov/ens/). ENS is a subscription-based service, meaning that users must go online and create profiles that specify their alert needs. Users can customize the types of alerts that they want to receive, including magnitude, location, and time of an earthquake. Users can also choose to receive alerts via email or phone. Unfortunately this service, despite its age, is not well used.

EARLY WARNING SYSTEMS COMPONENTS AND DESIGN

A comprehensive EWS usually consists of four key elements: (1) prior knowledge of the risks, (2) presence of a monitoring and warning service, (3) multi-layer information dissemination system, and (4) the capacity to take timely actions (UN, 2005). In order to be a functioning EWS, these integrated systems must be designed a specific disturbance in mind. An Integrated EWS is comprised of components designed to work together in order detect a disturbance that is pre-defined as environmental-stability threatening (United Nations, 2005). Integrated EWS' are, "chain[s] of information communication systems comprising sensor, detection, decisions, and broker subsystems" (United Nations, 2005). These communication systems work together to observe the known behaviors of a particular phenomenon and signal brokers who then decide if action or response is needed.

The best practice for EWS design was most recently put forth by Waidyanatha (2010) who spoke of the link between EWS components. The initial components are those that send the signal of detection – sensors. Sensors are instruments that can detect a change in the environment. These instruments could be biological, technological, or social. Sensors are the origin system of any EWS. Detectors are the second stage of sensing and these are instruments that allow the gathered information to be engaged and ultimately acted upon. Detectors are pre-formatted systems or protocols that allow either an automated or person-based decision making system to be engaged.

Most important to this work is the concept of the broker. Brokers are mediators between the sensors, who notice, and the responders, who act. A broker can be an individual or group, a system or human. Typically, these mediators are under strict protocols to gather information in order to send commands to dispatch for response. A broker is a multiple input multiple output system. We can call those who input messages, publishers, and those who receive messages, subscribers. The broker serves as a means of communication between these two types of participants. A broker serves as a central or origin source of communication (Waidyanatha, 2010). A broker may transform incoming messages from the language of sensors to the language of responders.

If we examine the four key elements of an EWS mentioned previously, current approaches to integrate social media fall into the first two categories (detector and sensor) and ignore the second two (broker and response). Many have examined social media data for its ability to serve as a sensor, or input, often in a passive way. These current models may be improving their ability to detect known crisis events through algorithmic manipulation of data gathered and scanned from the crowd, but the results of these analyses are not integrated into the third and fourth key elements of an EWS. The results must be disseminated to the population at risk and these results must be able to produce timely action.

There are three problems with designing EWS, though only the first is crippling. The core issue first encountered with building any early warning system is testing. Natural disasters are rare and unpredictable. As such, artificial data, scenarios and laboratory experiments, or the analysis of past data, can never deliver the essential quality needed for an early warning system—the ability to detect and deviate from planned procedures given the unexpected. Second, early warning systems must be an integrated chain of links that begin with detection of known disasters and end with notification and action (Waidyantha, 2010; Sakaki, et al., 2010; Stankiewicz, et al., 2013).

While social media data analysis is improving at the detection phase, the results of these social media sensors have not yet been integrated into a whole chain where they can deliver real-time actionable alerts. Third, high functioning early warning systems require multiple sensors as inputs, multiple modes of alert notification as outputs, and a community of semi-passive observer-participants able to receive and act of an alert. While there have been some projects that have attempted to integrate both physical and social media data into an alert system, none have integrated these algorithmic attempts with a community of observer-participants capable of accurate analysis in real time.

AURORASAURUS

Predicting when and from where the aurora can be seen is a non-trivial task. Space weather scientists use a suite of different methods to try to predict when the aurora might be seen (e.g. Newell, 2009; Sigernes, 2012). The strength of an aurora, and thus its location, is driven by the Sun's activity and the conditions in interplanetary space. In addition to driving the aurora, it is well known that space weather can have a strong, damaging, effect on both space- and ground-based technologies (e.g. NERC, 2012). While predicting when these space weather dangers might occur is an important task; we focus only on the less dangerous, but more beautiful, northern lights.

Figure 1 - Aurorasaurus Homepage.

An alpha version of Aurorasaurus (Aurorasaurus 1.0) has been live on the web for two years and has shown promising results. At its core, Aurorasaurus offers an estimate of the aurora's location and character while simultaneously seeking participatory data for improving forecasts of where the Aurora will be and is currently. It does this through the data it aggregates. First, we gather ad-hoc data indirectly offered by the crowd through social media. These data are filtered via keyword and with the help of citizen scientists, are further refined and verified. These verified data are then tested against predictive data from space weather observatories from around the world.

Risks are identified and sensors calibrated for those risks. In our case the "risk" is the event of a visible aurora over large population centers. Sensors are of three types, (1) physical, those scientific instruments that measure the atmospheric charged particles and solar emissions, (2) social media data in-directly gathered form Twitter in real-time that speak of aurora sightings, and (3) human. To address the third part of an EWS, notifications of aurora events are sent to community members who have registered to receive alerts when an aurora is seen within a certain radius of their location. These notifications are provided on multiple platforms and encourage the receiver to get a camera, go outside and look up.

PASSIVE DATA COLLECTION AND IMPROVEMENT

In general, the tweets rise when auroral activity rises above background levels and when the aurora starts to be visible over more populated areas. The Twitter data stream requires a significant level of filtering to select aurora-related tweets from the daily 300-500 million tweets shared in total. Our current classification algorithm, trained on a set of tweets manually classified by us, divides a tweet stream initially filtered by keyword into the following classes: positive sighting (i.e., someone saw the aurora), negative sighting (someone looked for the aurora but failed to see it), desire to see the aurora, not aurora related, and unclassifiable. Geo-location is accomplished either by using the geotag provided by Twitter directly (typically derived from GPS or other automated location services) or a gazetteer-based approach applied to toponyms (place names) in the tweet text or user profile. This yields approximately 400 tweets on a typical day and 1,000 or more during a non-routine event.

We bring users into the tweet filtering loop by allowing them to give feedback on the classification algorithm's judgments, using up- and down- votes similar to those found on the Reddit social news site (Potts, et al); this will both improve the algorithm's accuracy and make it more nimble in the face of changing language. Second, we integrate improved location inference techniques (for example, Priedhorsky et al. 2014).

The crux of how Aurorasaurus functions as an EWS is how it gathers, verifies, and distributes its data. Any time there is an aurora, viewable from an area with a human settlement that has widespread access to the internet; there is a constant stream of tweets associated with it. As Aurorasaurus gathers tweets, some server-side filtering is undertaken (for example to remove tweets from Twitter users with "aurora" in their username). Tweet location is extracted via the open-source CLAVIN geo-parsing software and stored on our server for users to verify manually.

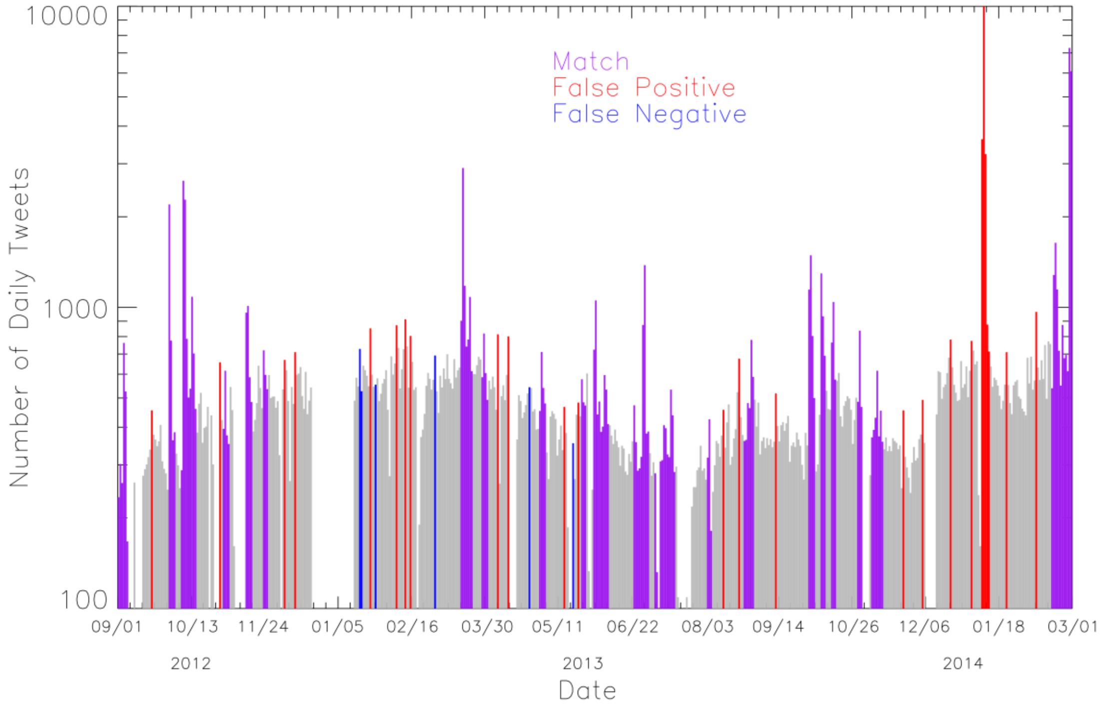

Figure 2 – A histogram of the daily number of aurora-related tweets spanning from September 2012 to March 2014.

Verification involves allowing our users to vote on whether a tweet is a real observation of the aurora. These observations are then combined with the direct entry of auroral observation (via a mobile or website form) in order to create alerts. If a certain score is achieved, the tweet is classified as a real sighting and is treated like a reported observation. In this way, it functions both as a sensor and broker.

In Figure 2, we present a histogram of the daily number of unverified aurora-related tweets, recorded between September 2012 and March 2014. The histogram bar is purple if a peak in the number of daily tweets is detected and this peak coincides with strong geomagnetic activity (i.e. a Match). The bar is colored red if a peak in the number of tweets is detected but this does not correspond to strong geomagnetic activity (i.e. a False Positive) and is colored blue if there is strong geomagnetic activity but no corresponding peak in the number of tweets (i.e. a

False Negative). The grey bars indicate that there was no strong geomagnetic activity on that day and that the number of tweets did not constitute a peak value.

We find that, during this 18 month period, 81% of geomagnetic storms are detectable by a peak in the number of aurora-related tweets. This detection rate is improved by modifying how we define the background level, however, this results in an increased number in false positives.

We define a geomagnetic storm by a period in which the Dst index is less than -40nT over a sustained period. The Dst index is a measure of the disturbance of the Earth's magnetic field (Sugiura, 1964). Increasingly large negative values of Dst (especially less than -40nT) indicate a geomagnetic storm is underway. As stated before, strong auroral displays, visible from lower magnetic latitudes, are associated with geomagnetic storms.

It is clear that the peaks in the number of daily Tweets match well with times of geomagnetic storms. This indicates that the number of aurora-related Tweets increases during times of strong geomagnetic activity (i.e. when an aurora is most likely to be visible to larger numbers of people).

DIRECT DATA COLLECTION VIA WEB FORMS

Each user of the Aurorasurus website may click on a link that states that they are currently viewing the aurora and want to report a sighting. The site asks if it can use the current location of the user as the location of the aurora sighting but also gives the user the option of adding a custom location. It then asks the user to report the date and time, the color of the aurora, the shape and movement of the aurora, the height in the sky and activity level of the aurora. Lastly, the user is encouraged to make comments about the aurora. During our initial data collection phase, of September 2012 through May 2014, we collected 2109 web observations via this form.

Not all observations were viable observations, but of those that were we were able to compare the number of observations against the Kp index of the period. Of the 2109 we received, 595 were real reports (i.e. not spam). Of those, 59 were observations of an aurora (the other 536 were saying they could *not* see an aurora What we found was interesting and is shown in Figure 3. The rise in the number of web observations preceded the Kp index by a few hours to a day and remained high during the event. We attribute this to the users of the site being strongly aware and observant of news reports concerning solar flares and coronal mass ejections and anticipating aurora.

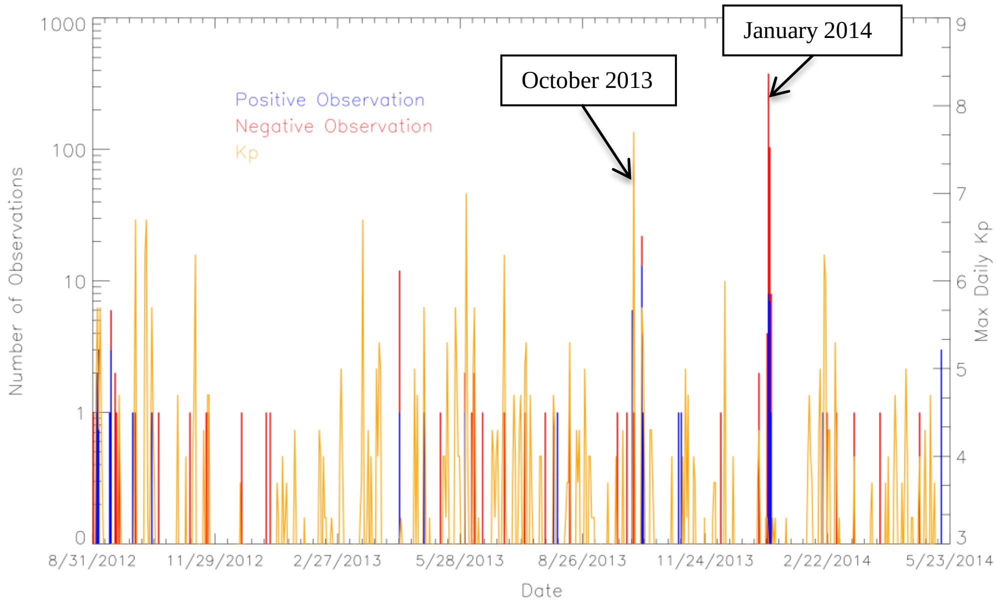

Figure 3– A histogram of the daily observation levels and the maximum daily Kp index.

Shown in Figure 3 is a histogram of the daily number of observations reported to the Aurorasaurus website during the initial data collection phase. Positive observations (where the user saw the aurora) are colored blue; negative observations (where the user did not see the aurora) are colored red. Plotted over the histogram in orange is the daily maximum Kp value.

Owing to the small number of observations, direct correlations between the number of observations and the Kp index are not meaningful. However, as expected and as with tweets, the observations tend to be reported when Kp levels are elevated.

There are two specific periods of interest which are worthy of note. In October of 2013, the number of observations increased by a factor of ten as the Kp index rose significantly and the aurora moved southward into more densely populated areas. In January of 2014 news outlets were reporting a potential visible aurora so activity on the site increased, but the Kp index did not rise and the auroras were in fact not visible at lower latitudes. This could be seen as a false positive. More valuable are the small peaks, as represented at the end of February and beginning of March 2014 where we see a small rise in both the Kp index and the number of observations.

In addition, we make Aurorasaurus more fully available in the field where its services are most useful. If the user makes a positive aurora sighting the user may then use the mobile application (Aurorasaurus mobile) to submit a positive sighting of the aurora directly. Users can take a photo or video of the event and tag it with location, color, and height in the sky and aurora type. The user enters this data on a simple mobile form. This entry contributes to improving the prediction algorithms of aurora activity in general and contributes to the predicted aurora oval made by Aurorasuarus in the specific at that time and location, allowing other users to be more confident in their seeking a sighting. We have IRB approval from two institutions (Los Alamos National Labs and Pennsylvania State University) to collect data from human subjects for this research.

The third and fourth elements of an early warning system are notification and action. To address the third part of an EWS, notifications of aurora events are sent to community members who have registered to receive alerts when an aurora is seen within a certain radius of their location. Notifications of an aurora event can be provided via text or email in near real time based on location. Push notifications to alert users can be of ongoing or even expected aurora activity in their locale. These notifications are provided on multiple platforms and encourage the receiver to get a camera, go outside and look up, in other words, take action.

The hybrid-forecasting engine combines data from physical, social and community sources to produce and estimate of the aurora oval. This is displayed on the map, which is visible on all platforms, including mobile platforms, again providing information to the aurora seeker in the field.

DISCUSSION

The central take-away from this paper should be that twitter alone can be seen as a good predictor of visible auroral activity, however, when this data is combined with data that has been improved by community members and data that has been directly contributed by community members, the ability of the system to predict auroral activity significantly rises. When community members help collect the data and improve the data, the utility of the system for the community also rises. Lastly, when community members participate in the data collection and improvement they become interested, and vested, in the overall project. It becomes, in part, *their* project and motivation ceases to become insurmountable. At the time of writing this paper data was available from one significant solar storm in which community members made direct contributions (59 posted actual sightings to the website), this, in combination with the up-and-down voted tweets, helped to improve the predictor engine.

This strong engagement between the designers and researchers involved in the project and the community of users through the continued, active engagement in data collection and improvement, is the model that is carried through to our extension to an early warning system. As mentioned above, in order to perform at their highest potential, early warning systems must present four key elements: (1) prior knowledge of the risks, (2) presence of a monitoring and warning service, (3) multi-layer information dissemination system, and (4) the capacity to take timely actions.

A comprehensive and effective early warning system would fully engage the public in all four elements of such a system. In the case of aurorasaurus, the greatest impact can be seen in improvements of parts 2 and 3, a monitoring and warning system and an information dissemination system. Community contributors act as part of the system, scanning the skies, taking pictures, recording data and uploading this material in a timely fashion. Community contributors monitor the changing map displayed on the website homepage seeking knowledge concerning the position of the auroral oval and then taking action based on its movement. Community contributors read collected semi-filtered tweets concerning auroral activity and improve their quality by selecting actual sightings, which then appear on and improve the auroral map. All of these activities involve and engage the community in monitoring both the physical and online worlds and offering warning when the aurora moves closer to a given geographic location. Aurorsaurus also engages the community via notifications triggered by both data contributed by scientists and community members, addressing the information dissemination element of an EWS.

Aurorasaurus has the potential to be a community-driven EWS broker. The essential link between the first two elements and the second two elements of an EWS is the broker, the link between monitoring and notification. Current attempts to integrate data from the crowd or social media have not yet moved beyond the stage of contributing the results of social media monitoring to existing monitoring systems. Aurorasaurus functions as sensor, detector, and broker. There are multiple inputs and multiple outputs to the same system. We believe that the map and the visualization of the auroral oval, and positive sighting pins, in real time is the key to this broker function. It translates data from three sensors types, (1) physical, those scientific instruments that measure the atmospheric charged particles and solar emissions, (2) social media data indirectly gathered from twitter in real time that speak of aurora sightings, and (3) data that is directly contributed by community members on confirmed sightings—to a single visual pane. This allows the viewer to interpret the current location of the visible aurora in relation to themselves in that moment and decide to take action. As the map updates with new data it allows the user to interpret changes in the location of the visible aurora and make predictions and change their behavior accordingly. Ultimately, Aurorasaurus serves as a prototype early warning system that integrates the community of users at each phase.

CONCLUSIONS

Seeing the aurora is often an emotionally meaningful moment in someone's life, one which merits membership in "bucket lists" and/or major travel at financial and time cost. In this context, the significance of a solar maximum, auroral activity moves southward, become viewable by larger populations; that is, the aurora comes to the people rather than the people coming to the aurora. However, while visible aurora becomes more likely at middle latitudes, it is still a rare event — one cannot intuitively know when and where aurora will appears. Therefore, a mechanism, which alerts interested parties to this event, is valuable in itself. Communities are already interested in the aurora and would be further motivated to participate if they could be alerted to potential sightings in real time. Aurora sightings have the potential to yield insight for both applications, serving as the motivation for participation in an early warning system as well as citizen science projects.

We argue that this broadly shared deep interest provides two valuable opportunities. First, it can contribute to the body of scientific theory and practice around the aurora. Finally, it can provide a testing environment for a crowdsourced and community-based early warning system.

We have demonstrated that the peaks in the number of Tweets concerning aurora match well with times of geomagnetic storms. This indicates that the number of aurora-related Tweets increases during times of strong geomagnetic activity (i.e. when an aurora is most likely to be visible to larger numbers of people). Despite this strong correlation, we have also demonstrated that human intervention is essential to increasing the predictive power of the system. At best we were able to predict 81% of storm activity. With community-based filtering by members the predictive ability of the system increases.

Most importantly we have demonstrated that the combination between all three sensors, the Tweets themselves, the filtering by community members and the sightings directly contributed by community members leads to a stronger prediction engine, better nowcasting of the aurora event, better visualization through mapping and more timely and geographically precise notifications.

While this paper does not reflect the citizen science elements of the project due to space constraints, the community of users for aurorasaurus is both interested in and contributes to space weather science. Participants are offered educational modules about space weather and training on how to filter images and tweets as they arrive. This encourages learning about space weather beyond that of common knowledge. Free text observations in existing social media, structured observations entered via questionnaire, both text and visual, comprise scientific contributions to improve understanding auroral activity and prediction. This research will have a positive impact on the concept of nowcasting and forecasting as they relate to space weather.

We expect that our tools will serve as a proof-of-concept for crowdsourcing other forms of early warning system. Such a system could be employed during other kinds of disasters such as earthquakes, tornados, tsunamis, and flooding. That is, real-time alerts informed by crowd participation could save lives.

ACKNOWLEDGMENTS

This material is based upon work supported, in part, by the National Science Foundation (NSF) under Grant #1344296. This work was supported in part by the U.S. Department of Energy through the Los Alamos National Laboratory LDRD Program. LANL is operated by Los Alamos National Security, LLC for the Department of Energy under contract DE-AC52-06NA25396. Approved by LANL for public release, LA-UR 14-20644.

REFERENCES

1. Achrekar, H., Gandhe, A., Lazarus, R., Yu, S. H., & Liu, B. (2011). Predicting flu trends using twitter data. *Proceedings of the 2011 IEEE Conference*.
2. Ahmed, S., Nelson, D., Biedrzycki, P., Sandy, M., Opel, S., Franco, Z. (2012). Community Engagement for Disaster Research Proceedings of the 9th International ISCRAM Conference.
3. Arias, M., Arratia, A., & Xuriguera, R. (2013). Forecasting with twitter data. *ACM Transactions on Intelligent Systems and Technology*, 5, 1, 8.
4. Asur, S., & Huberman, B. A. (2010). Predicting the future with social media. In Web Intelligence and Intelligent Agent Technology (WI-IAT). *IEEE/WIC/ACM International Conference*, 1, 492-499.
5. Berg, M. (2001). Implementing information systems in health care organizations: myths and challenges. *International Journal of Medical Informatics*, 64, 2, 143-156.
6. Boehmer, Erin. (2010). Coordinating Efforts by Volunteer and Technical Communities for Disaster Preparedness, Response, and Relief. Science and Technology Innovation Program - Woodrow Wilson International Center for Scholars.
7. Chen, H., Achrekar, B., Liu, & R. Lazarus. (2010). Vision: towards real-time epidemic vigilance through online social networks: introducing sneft – social network enabled flu trends. *ACM Mobile Cloud Computing and Services*.
8. Earle, P., Guy, M., Buckmaster, R., Ostrum, C., Horvath, S., & Vaughan, A. (2010). OMG earthquake! Can Twitter improve earthquake response? *Seismological Research Letters*, 81, 2, 246-251.
9. Earle, P. S., Bowden, D. C., & Guy, M. (2012). Twitter earthquake detection: earthquake monitoring in a social world. *Annals of Geophysics*, 54, 6.
10. Franco, Z. E., Zumel, N., & Beutler, L. E. (2007). A ghost in the system: Integrating conceptual and methodological considerations from the behavioral sciences into disaster technology research. *Proceedings of the 4th ISCRAM Conference*.
11. Ginsberg, J., Mohebbi, M. H., Patel, R. S., Brammer, L., Smolinski, M. S., & Brilliant, L. (2008). Detecting influenza epidemics using search engine query data. *Nature*, 457, 7232, 1012-1014.
12. Glantz, M. (2004). Usable Science 8: Early Warning Systems: Do's and Don'ts. Report of Workshop held 20-23 October 2003, *ISSE.*
13. Gillmor, Dan. (2006). We the Media: Grassroots Journalism by the People, for the People. Sebastopol, CA: O'Reilly Media.
14. Goodchild, Michael F. (2007). Citizens as Sensors: The World of Volunteered Geography. *GeoJournal* 69 (4): 211–21.
15. Hall, D. L., Llinas, J., McNeese, M., and Mullen, T. (2008). A Framework for Hard/Soft Fusion. *Proceedings of The 11th International Conference on Information Fusion*.
16. Ichinose, F., Maeda, Y. Kosaka, N., Higashida, M., Sugiyama, M., Takeda, H., Yamamoto, T., Hayashi, H. (2014). A Fundamental Study of Efficiency of Information Processing in Emergency Operations Center. *Journal of Disaster Research* 9, 2, 206-215.
17. Kringsman, M. (2009). Study: 68 percent of IT projects fail, taken from Kieth Ellis, *The Impact of Business Requirements on the Success of Technology Projects.* IAG Consulting.
18. Kuziemsky, C., O'Sullivan, T., Corneil, W. (2012). Upstream-Downstream Information Systems Design In *Proceedings of the 9th ISCRAM Conference*.
19. Lampos, V., De Bie, T., & Cristianini, N. (2010). Flu detector-tracking epidemics on Twitter. In *Machine Learning and Knowledge Discovery in Databases*, 599-602. Springer, Berlin.
20. Landgren, J. (2010). Principles of radical research in the area of information systems for crisis response. In *Proceedings of the 7th ISCRAM Conference*.
21. Lyytinen K, Hirschheim R. (1987). Information failures—a survey and classification of the empirical literature. *Oxford Surveys in Information Technology*; 4, 257–309.
22. MacEachren, A. M., Robinson, A. C., Jaiswal, A., Pezanowski, S., Savelyev, A., Blanford, J., & Mitra, P. (2011, July). Geo-twitter analytics: Applications in crisis management. *25th International Cartographic Conference*, 3-8.
23. Malone, T. W., Laubacher, R. & Dellarocas, C. N., (2009). Harnessing Crowds: Mapping the Genome of Collective Intelligence. *MIT Sloan Research Paper* No. 4732-09.
24. Mendoza, M., Poblete, B., & Castillo, C. (2010, July). Twitter Under Crisis: Can we trust what we RT?. *In Proceedings of the 1st Workshop on Social Media Analytics*, 71-79.
25. Motoyama, M., Meeder, B., Levchenko, K., Voelker, G. M., & Savage, S. (2010). Measuring online service availability using twitter. *Proceedings of the 3rd Workshop on Online Social Networks*, Boston, MA.
26. Newell, P. T., Sotirelis, T., & Wing, S. (2009). Diffuse, monoenergetic, and broadband aurora: The global precipitation budget. *Journal of Geophysical Research: Space Physics* (1978–2012), 114, A9.

27. Norheim-Hagtun, I., and P. Meier. 2010. Crowdsourcing for Crisis Mapping in Haiti. *Innovations: Technology, Governance, Globalization* 5, 81–89.
28. North American Electric Reliability Corporation (NERC) (2012), Special reliability assessment interim report: Effects of geomagnetic disturbances on the bulk power system.
29. O'Connor, B., Balasubramanyan, R., Routledge, B. R., & Smith, N. A. (2010). From tweets to Polls: Linking Text Sentiment to Public Opinion Time Series. *Proceedings of the 5th International AAAI Conference on Weblogs and Social Media*, 11, 122-129.
30. O'Sullivan, T. L., Kuziemsky, C. E., Toal-Sullivan, D., & Corneil, W. (2013). Unraveling the complexities of disaster management: A framework for critical social infrastructure to promote population health and resilience. *Social Science & Medicine*, 93, 238-246.
31. O'Sullivan, T., Corneil, W., Kuziemsky, C., Lane, D. (2013). Citizen Participation in IS Design for Asset-Mapping. In *Proceedings of the 10th ISCRAM Conference*.
32. Othman, S. H., & Beydoun, G. (2010). A disaster management metamodel (DMM) validated. *In Knowledge Management and Acquisition for Smart Systems and Services*, pp. 111-125. Springer Berlin Heidelberg.
33. Priedhorsky, R. Aron Culotta, and Sara Y. Del Valle. (2014). Inferring the Origin Locations of tweets with Quantitative Confidence. *Proceedings of the the 17th Annual Conference on Computer Supported Cooperative Work and Social Computing*, Baltimore, MD.
34. Sakaki, T., Okazaki, M., & Matsuo, Y. (2010). Earthquake shakes Twitter users: real-time event detection by social sensors. *Proceedings of the 19th international conference on World wide Web,* 851-860.
35. Sauer C. (1993). Why information systems fail: a case study approach, information systems series. Alfred Waler, Henley-on-Thames, UK.
36. Sigernes, F., Dyrland, M., Brekke, P., Gjengedal, E. K., Chernouss, S., Lorentzen, D. A., Oksavik, K., & Center, S. W. (2011). Real time aurora oval forecasting-SvalTrackII. In *other words*, 3, 3.
37. Sugiura, M. (1964), Hourly values of equatorial Dst for the IGY. *Annals of the International Geophysical Year*, 35, 9.
38. St. Denis, Lise Ann, Hughes, A., and Palen, L. (2012). "Trial by Fire: The Deployment of Trusted Digital Volunteers in the 2011 Shadow Lake Fire." In *Proceedings of the 9th ISCRAM Conference.*
39. Stankiewicz, J., Bindi, D., Oth, A., & Parolai, S. (2013). Designing efficient earthquake early warning systems: case study of Almaty, Kazakhstan. *Journal of seismology*, 17, 4, 1125-1137.
40. Starbird, K., & Palen, L. (2010). Pass it on? Retweeting in mass emergency. *Proceedings of the 7th ISCRAM Conference*.
41. Starbird, K., & Palen, L. (2011). Voluntweeters: Self-organizing by digital volunteers in times of crisis. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1071-1080.
42. Starbird, K., Muzny, G., & Palen, L. (2012). Learning from the crowd: Collaborative filtering techniques for identifying on-the-ground Twitterers during mass disruptions. *Proceedings of the 9th ISCRAM Conference.*
43. Starbird, K. & Palen, L. (2012). (How) Will the Revolution be Retweeted? Information Propagation in the 2011 Egyptian Uprising *Proceedings of the ACM 2012 CSCW Conference.*
44. Tapia, A.; Lalone, N.; MacDonald, E.; Hall, M. (2014). Crowdsourcing Rare Events: Using Curiosity to Draw Participants into Science and Early Warning Systems. *Proceedings of the 11th ISCRAM Conference*
45. Troy, D. A., Carson, A., Vanderbeek, J., and Hutton, A. (2008). Enhancing community-based disaster preparedness with information technology. *Disasters:* 32, 1, 149-65.
46. United Nations International Strategy for Disaster Reduction (UN/ISDR). (2005). Four Elements of People Centered Early Warning Systems. *Early Warning Systems—Interdisciplinary Observations and Policies from a Local Government Perspective.*
47. Van Gorp, Annemijn F. (2014). "Integration of Volunteer and Technical Communities into the Humanitarian Aid Sector: Barriers to Collaboration." *Proceedings of the 11th ISCRAM Conference*.
48. Von Ahn, L., Blum, M., & Langford, J. (2004). Telling humans and computers apart automatically. *Communications of the ACM.* 47, 2, 56-60.
49. Waidyanatha, N. (2010). Towards a typology of integrated functional early warning systems. *International Journal of Critical Infrastructures*, 6, 1, 31–51.
50. Wakamiya, S., Lee, R., & Sumiya, K. (2011). Towards better TV viewing rates: exploiting crowd's media life logs over Twitter for TV rating. *Proceedings of the 5th International Conference on Ubiquitous Information Management and Communication*.
51. Wolfram, M. S. A. (2011). Modelling the Stock Market using Twitter. Masters Thesis for The University of Edinburgh.
52. Young, J.C., Wald, D.J., Earle, P.S., Shanley, L.A. (2013). Transforming Earthquake Detection through Seismology. *Wilson Centers. The Commons Lab*. Washington. D.C.