

During emergency events, information is the most important resource. Updated, verified information is critical to enabling residents, response teams, and government agencies to make evidence-based decisions in a timely manner.

With the spread of COVID-19, we've seen an uptake of social media usage as people are looking for verified information, sharing information and stories with each other, and also planning community based efforts for response.

When faced with an emergency, people turn to the applications that they're familiar with in order to get information and support.

So how can we leverage this increased use of social media?

There's an incredible amount of location-based information being shared, and while the proliferation of social media might first appear as so much noise, with the proper methods and technologies for gathering, sorting, and analyzing data this noise can be transformed into critical information for both understanding and reducing risk.

Starting in Indonesia our team developed an open source software that transforms social media into a collaborative network for participatory data gathering and information coordination during disaster events.

Throughout this module I will be drawing from our experience in developing crowdsourced disaster mapping platforms in Indonesia and South Asia to offer principles and methodologies on harnessing the power of social media to gather and coordinate verified situational reports from people on the ground through the use of AI-assisted chatbots.

To provide some context, I will first begin with a brief overview of our work in Indonesia. Then I will outline key considerations in developing chatbots to crowdsourcing emergency-related information. I will then offer potential applications of these methodologies for disease outbreaks, and finally conclude with key takeaways based on the lessons learned from developing PetaBencana and in looking forward to the various applications of these tools and methodologies.

Our work began in Jakarta, which experiences an annual monsoon season. But flooding has become increasingly unpredictable due to rapid urbanization and climate change.

In this time, residents have developed new ways to self-organize through network communication technologies. As we see with the current pandemic during emergency situations people use the applications that they are already using with an increased frequency to look for and share information.

Social media networks operating over Internet connected smartphones have significantly decreased the barriers for communities to collect share and coordinate hazard reports and self-organize their response accordingly.

In fact during the 2013 monsoon season, we saw that over a period of six months there were over each blue dot on the slide represents a post about flooding in Jakarta and the black line represents the

administrative boundary of the city.

So we see that people from the entire city are sharing information about the emergency event, and we see that even in cities like Jakarta which are often thought to be data scarce, there's actually a lot of data there.

In fact, the tacit knowledge of residents provides a data source of unprecedented resolution for mitigating risk. So we questioned how we could transform the noise of social media into actionable information. How could we filter through this wealth of data, find the messages that are paramount to understanding critical needs, and use that for safety and response?

Current methods to crowdsource situational information during disasters require either manual processes of data sorting, interpretation and management, or rely on computationally expensive rule-based classification.

So our team developed a new methodology to address these issues. By developing an open-source software that automatically gathers, filters, and visualizes verified crowdsource reports, we eliminate the need for expensive and time-consuming data processing.

This way we can collect crowdsource reports and display them on a public map in real time.

So how does it work?

First, we use chatbots to act as a filter. So the software is designed to listen to specific keywords such as "flood" and when it detects this keyword our chatbots send an automatic response asking users if it's really flooding and if they would like to contribute to community flood mapping.

If the user responds, our bot sends them a link to complete a flood report. When the user clicks on this link, they are directed to a simple four step process where they confirm their location, they add the flood height, they add photos, and a description.

Once the report is submitted, it's immediately placed on a public map in real-time. And so this is a method of crowdsourcing that's not passive data mining.

We're not plotting all of the posts that contain #flood. But instead we use that keyword as an initial filter and then employ chatbots to engage in a conversation with residents to verify flooding and gather those confirmed situational reports.

And when you click on any one of those reports, this is what it would look like, where the detail of the report would appear. Here we see the flood height, the time and date of the report, the description that the user has submitted as well as the photo.

And so this is a method of crowdsourcing that allows us to reach millions of users every day and encourage greater public participation in infrastructure monitoring and flood event reporting.

It also leverages capacities for all residents to equally participate in information sharing and decision-making and resource management within their communities because it democratizes decision support tools.

And in this way, it fosters a form of civic core management. By making this critical information freely and transparently available to all residents, humanitarian agencies, and government agencies, everyone has access to the same information and so various sectors of the city can act in much more coordinated ways.

The reason it has worked so well is that it leverages the infrastructures that already exists – the social media channels that people are already spending most of their time on, and through the use of a chatbot, we can direct that attention towards response, relief, and mutual aid.

The chatbot operates on multiple social media platform simultaneously to gather data from users on different networks and it's based on social media activity on in Indonesia.

And so currently we have chatbots operating on Twitter, Telegram and Facebook, but the software is designed to be able to integrate to any variety of social media and instant messaging applications.

The advantage of this is that by co-opting existing channels of communication, the need of additional software on behalf of the user is avoided, which not only increases awareness but makes it far easier and more encouraging for users to participate in the process of crowdsourced information sharing.

They don't have to learn a new application. They don't have to download anything different. They don't have to learn a new user interface.

They're simply using the processes they're already familiar with, and we're intercepting those in order to solicit resident participation in crowdsourced information sharing.

So there are a few important things that we can draw from this method of crowdsourcing.

Firstly, chatbots are a great way to capture attention.

Especially in this kind of a situation where we are saturated with information, a bot interface will be natural for users since they use messengers every day to communicate and connect with friends and family, and we've found that people are very receptive of the chatbot because it engages in direct conversations through them, through the channels they are used to, requiring very minimal additional effort on their part to participate in this process.

With careful design, chatbots can be incredibly effective in encouraging participation in crowdsourced information sharing. What we've learned through our work is the importance of designing both the tone and visual communication strategies of the chatbot to be effective in soliciting participation.

Of course, this is something that will vary across different cultural contexts. For our particular case, context and hazard the chatbot has a very friendly tone with the simplest language. And so we're using rhetorics that are familiar and casual so that the emphasis is placed on mutual aid instead of the trauma of a disaster.

And these types of nuances are especially important to consider in

crisis events because we don't want to amplify anxiety, but rather channel that for a collaborative response and safety.

Thirdly, the report cards allow us to structure this crowd source information. So whereas the chat box acts as an initial filter and to engage conversation, the report cards begin to structure crowdsourced information.

Because when we started the project we noticed that residents have a wealth of information to share, but of course, there's no consistency to what they're sharing. Some people would share very detailed reports on social media that included location, that included information about road accessibility, information about the flood height, but others would simply include a picture that was often very vague and not much can be gained from this.

However they're still interested in sharing information, and so by developing these report cards we've created a mechanism by which we can begin to structure crowdsourced information in a standardized manner so that all reports contain some baseline information and so that data can be presented in the useful manner for people trying to make decisions.

At the same time there's always this free text option available so that people can add as much additional information as required. Of course, this is a very simple design, but it arose out of the process of iterative co-research and testing with the widest variety of stakeholders in order to determine the critical pieces of information that would be necessary to help people make decisions.

For flooding, of course, it was the location and flood height that were deemed the most critical pieces of information.

But designing the interface to report the flood height is also culturally specific. In Indonesia, people describe the flood based on bodily measurements, saying that the water is ankle deep or hip deep, and so the graphic is designed to be as intuitive to the way people are already speaking about the disaster so that there's a useful way to begin to parameterize that information.

On the one hand, the chatbots create a seamless integration into the tools that people are already using for communication, and on the other hand, the graphic of the report card also helps to create this very uncumbersome method of sharing information and submitting a report.

Of course, this cultural specificity of that design will vary across geographies and hazards and different contexts. When the software was adapted to power a flood mapping platform for Vietnam, the bot no longer asks people to describe the flood height based on a human, but rather on the height of a motorcycle wheel, as you can see here. Does it cover a quarter of the wheel, a half of the wheel, or the full wheel? As we expand to other hazards, we continue to learn the different ways people speak and understand risk. For example, when speaking about haze, people often describe it based on respiratory symptoms.

So we've designed the reporting mechanism based on the way people

speak about the risk and by speaking with a variety of stakeholders and experts we found a way to calibrate those colloquialisms to assess the severity of the hazard.

This helps everyday people understand each other's reports, but when the information is structured and calibrated it also helps formal agencies assess risk and make appropriate decisions for response. Finally, the report cards also allow for an automatic process of verifying reports, eliminating the need for expensive or time-consuming data processing. Through this mechanism of reporting, we've designed a threshold of participation that is low enough to incentivize reporting because it's a quick and simple four-step process.

But, that threshold of participation is still designed to be high enough to disincentivize false reports because there's still a level of effort needed to submit a report beyond simply submitting a hashtag on social media, and so this was designed through a process of research and testing where we finally ended up with this configuration for flood reporting.

And it's been an incredibly useful tool for us and the platform has never seen a case of a false report since the platform launched in 2014.

Of course, the verification of information not only has to do with the report cards but critically also in building trust and transparency in the platform.

So in building trust and transparency it's important for people to know where the information is going and that when it is received it's acted upon. So to give the example in Indonesia, we work with government emergency management agencies who are monitoring the map in real-time, using this to understand resident needs, and respond to what is going on.

In any context, we've learned that there tends to be a level of friction between formal and informal sectors of the city and while there may have been an initial hesitation even when we started, through processes of collaborative building and socialization, and as the platform has been used over the last several years, there's a greater sense of trust where residents really see the difference in response when they're submitting reports.

At the same time, government agencies are recognizing the viability of crowdsourced information in providing the source of granular data that they otherwise would not have.

So this trust is something that is built over time. But even in the very initial stages, it's important for any contributor to the platform - any contributor willing to share crowdsourced information - it's important for them to understand that when they're making a report, that report is being recognized.

There has to be some sort of feedback mechanism that allows them to understand this, whether it's by allowing them to see the report appear on a map in real time or by socializing enough and understanding that the report is being used by their fellow neighbors

or response teams.

And it's critical that contributors understand how that information is being used, because besides incentivizing engagement it also disincentivizes false reports. In Indonesia, there's also an additional tool that we've developed to build this continued trust and transparency and to facilitate additional information sharing.

So for this particular case, we developed an interface for the government management emergency agencies where they can input formal information onto the map.

And so, the platform becomes a way for transparent two-way information sharing which facilitates further coordination and verification of the flood reports, but it also has been critical in building trust in the platform and among contributors.

Of course, designing engagement and building trust are also closely linked to questions about privacy and security. Even though the data is made publicly available, we anonymize all of the reports and completely remove the users' metadata, and it's important for our users to know that we only keep the parts of the data we need for the map.

By ensuring privacy and anonymity, users also feel encouraged to contribute to crowdsource information sharing. They're not hesitant to submit a report, and this becomes a more pertinent issue as a disaster or emergency is more politicized contains more sensitive information. Not everyone is willing to share a report if they're worried that this will compromise their safety, and so for our particular case, we decided that it wasn't necessary for us to store any user data and it's important that the user is aware of this.

CogniCity is a free and open source software designed to be translatable to other geographies and challenges. It's already been adapted to power disaster mapping platforms in India, Vietnam, and the Philippines.

We are also currently adapting it for other hazards including earthquakes, forest fires, volcanoes, and haze. However, the software is not limited to collect reporting about hazards, but it's meant to be a tool to support various forms of urban coordination and there are already various experiments underway including forms of food distribution for agricultural uses etc.

And so there's several ways that the software can be adapted for a disease outbreaks as well.

One of the applications that we see it could be useful for in this context is developing community level coordination features, where people could submit reports about things or services that they want, need, or have.

We have preliminary studies of what this might needs to be and although it's not operationally standing by, we think it would be useful to develop for the coordination of future emergency events such as the one we are seeing today.

Because we've seen that with the spread of COVID-19, there are already numerous examples from around the world where people are using social

media, or have developed simple forms or are spreading messages via instant messaging applications where volunteers can offer help and vulnerable populations or caretakers or people who need it can ask for help.

And right now, we see that there's a manual process of pairing these two together: the want and the need.

But CogniCity open source software could be used to enhance the capacities of initiatives like these at a much larger scale, eliminating much of the manual processes that are currently hindering rapid coordination.

Of course, there various other applications of the types of things that can be reported, of the types of coordinations that can happen, and various combinations as well. But especially in a pandemic like this, it's important that such tool is designed with careful consideration so as to respect privacy, be secure, and be a useful information source that enables people to take action rather than spreading further confusion or panic.

And so the example we're offering here is only one potential application. Of course, there are many potential uses and we're excited to see what might come out of it, but we caution that these are developed with careful consideration, again to be secure and useful.

And so to conclude, just a few key takeaways. First, it's critical to identify the information that would be useful to crowdsource.

Just because something can be collected doesn't mean it should. Some tips for this process that have helped us is first it's useful to map the information flow, and this includes mapping the existing information flow, because this will help identify gaps and how crowdsourced information might help fill the gap.

In this process it's important to define who will be using the information that is to be crowdsourced and how will it support them. What will this information help them do?

Throughout this process it's always important to consider the implications of collecting such information and the different ways that it might be used or misused.

Especially in light of the current pandemic, it's important to be careful that the crowdsourced information doesn't lead to increased panic or disorder but instead that it enables people to make accountable decisions.

Another key importance is to build accountability and trust. Once submitted, how is the information acted upon?

How will residents know that their information that they're being shared is acted upon? How is this conveyed to them so that they're incentivized to report? And how do we make sure that there is a process of an accountability in that follow-up?

It's also incredibly important to safeguard the privacy and sensitive information. Especially in the current pandemic, surveillance is becoming an issue of increasing concern.

It is paramount to ensure that crowdsourced information will not

infringe on the privacy of residents or put people in dangerous situations. What information is useful to share with everyone and what information is sensitive?

Are there aspects of crowdsourced information that are not made publicly available to everyone, and if so what are the implications of doing that?

Also, it's important to design simple information gathering prompts, and this must be done through carefully considering the language and interface, adapting to an understanding of culturally specific or context specific or time specific behaviors.

Finally, throughout the whole process it's important to ensure inclusivity. This has been always an important consideration for us, but especially in the current context the virus is making inequalities increasingly visible and some of these pertain to the accessibility to social media networks.

So how do we ensure that the methods of crowdsourcing are inclusive and not only representative of the voice of a small demographic who have access to these networks, but that this process of crowdsourcing also supports the concerns of a much larger group of people?

In addressing all of these points above, what has really helped our platform become a democratic platform for mutual aid is designing the systems completely out of a process of collaborative co-research, and this is perhaps the most important takeaway in development being any tool.

In doing so we've seen that crowdsourcing can really have a positive impact. It can have an incredible impact in minimizing risk and loss and supporting practices of mutual aid.

And so to end, just to say that in consideration of the heightened use of social media that we've seen, that we continue to see during emergency events, there are ways to intercept these communication channels to transform these networks into more collaborative networks in order to support aid, response, and care if done through processes of careful and considered design.