

Crowdsourced translation for emergency response in Haiti: the global collaboration of local knowledge

Robert Munro

Department of Linguistics
Stanford University
Stanford, CA, 94305
rmunro@stanford.edu

Abstract

In the wake of the January 12 earthquake in Haiti it quickly became clear that the existing emergency response services had failed but text messages were still getting through. A number of people quickly came together to establish a text-message based emergency reporting system. There was one hurdle: the majority of the messages were in Haitian Kreyol, which for the most part was not understood by the primary emergency responders, the US Military. We therefore crowdsourced the translation of messages, allowing volunteers from within the Haitian Kreyol and French-speaking communities to translate, categorize and geolocate the messages in real-time. Collaborating online, they employed their local knowledge of locations, regional slang, abbreviations and spelling variants to process more than 40,000 messages in the first six weeks alone. According to the responders this saved hundreds of lives and helped direct the first food and aid to tens of thousands. The average turn-around from a message arriving in Kreyol to it being translated, categorized, geolocated and streamed back to the responders was 10 minutes. Collaboration among translators was crucial for data-quality, motivation and community contacts, enabling richer value-adding in the translation than would have been possible from any one person.

1 Background

Collaborative translation is not new to linguistics. Researchers have always engaged multiple languages consultants in transcription, analysis and translation, especially for less widely-spoken languages. However, the recent proliferation in crowdsourcing for language processing has its roots more in industry than academia. To pick one catalyst, it is probably when the natural language search startup *Powerset* started using the micro-tasking/crowdsourcing platform *Amazon Mechanical Turk* to make semantic annotations of web-based text. From 2007-2008 at least three independent directions in crowdsourcing were inspired by this: Snow et al (2008) sparked a new direction in natural language processing by showing that annotation via crowdsourcing could be as accurate as expert annotation at a fraction of the cost, inspiring innovative work in crowdsourced translation (Callison-Burch 2009); the company *CrowdFlower* was founded, specializing in quality-control for crowdsourced work; and linguists began using crowdsourcing for experimental paradigms (Munro et al, 2010) with the majority of empirical research in some departments now employing crowdsourcing technologies.

Social development organizations had also begun leveraging crowdsourcing technologies. *SamaSource* were setting up computer centers in a number of countries so that people could be employed through microtasking. Platforms like Amazon Mechanical Turk pay, at best, \$2-\$3 per hour, but this is a competitive income for many people in regions of the world with low costs of living.

Within social development ‘crowdsourcing’ also means sourcing data from the public, as with the *Ushahidi* platform which was created to monitor post-election violence in Kenya, and has subsequently been deployed for crowd-based reporting on many projects globally.

While text-messaging has quickly become the dominant form of remote communication in many parts of the world there are/were few controlled studies focused on processing large volumes of text messages in less-resourced languages. We had recently completed a study using machine-learning in partnership with *FrontLineSMS:Medic* (Munro and Manning 2010), but did not explicitly investigate crowdsourcing for privacy reasons relating to the use of medical data.

2 Mission 4636

In the wake of the January 12 earthquake in Haiti it quickly became clear that the existing emergency response services had failed and while communication channels were damaged, most text messages were still getting through. A number of people quickly came together, all connected by Josh Nesbit of *FrontLineSMS:Medic*, to establish a text-message based emergency reporting system through a free ‘4636’ phone number, giving a voice to the crisis-affected population. From launch the messages were tied into an ad-hoc system reporting trapped people to search and rescue teams within Haiti. Within days the US Military had recognized the quantity and reliability of actionable information in these messages, especially in remote areas, and set up task forces to respond.

There was one large hurdle: the majority of the messages were in Haitian Kreyol which was understood by very few of the responders. No professional translators of Kreyol were sitting idle and resources among all aid organizations were stretched. We therefore decided to crowdsource the translation of messages, allowing willing volunteers from within the Haitian Kreyol and French-speaking communities to translate and categorize the messages. By embedding a map into the crowdsourcing platform, the translators could click on the location to generate the exact coordinates. Therefore, the emergency responders did not just receive plain-text messages in Kreyol, they also received translated, categorized messages with latitude and longitude. The median turn-around from

Fanmi mwen nan Kafou, 24 Cote Plage, 41A bezwen manje ak dlo
My family in Carrefour, 24 Cote Plage, 41A needs food and water

Moun kwense nan Sakre Kè nan Pòtoprens
People trapped in Sacred Heart Church, PauP

Ti ekipman Lopital General genyen yo paka minm fè 24 è
General Hospital has less than 24 hrs. supplies

Fanm gen tranche pou fè yon pitit nan Delmas 31
Undergoing children delivery Delmas 31

Figure 1: Examples of messages

receiving a message to having it translated, categorized and geolocated and streamed back to responders in Haiti was less than 10 minutes.

Following the launch, we immediately began looking for a robust translation platform to ensure data-integrity and explicitly track workers. We engaged CrowdFlower who immediately agreed to host this service. At the first meeting we discovered that their long-term social development partners, Samasource, had signed an agreement with *1,000 Jobs/Haiti* to establish a microtasking center in Mirebalais outside of Port-au-Prince, just 20 minutes before the earthquake. The earthquake had initially delayed the plans to establish the center, but the local organizers were pleading that people needed work now more than ever. By bringing them into the partnership we were able to expedite the delivery of equipment to Mirebalais and immediately provide work, with the paid workers in Haiti gradually taking over the translation task from volunteers through March 2010.

The mapping component of the system was hosted by *Ushahidi*, launched for the first couple of days as a parallel crowdsourcing effort out of Tufts University. They took the structured messages from the translation platform, refined the coordinates, clustered reports geographically and worked with the responders to identify actionable incidents. FEMA director, Craig Fugate, identified this as the most up-to-date crisis mapping service and many organizations began using the maps as their primary crisis report data, including Charity Water, USAID, World Food Program, SOUTHCOM and UNDP. It was also transferred to local ownership over the following months, with *Ushahidi* helping a local tech company *Solutions*



Figure 2: Locations of interactions between people collaborating to translate messages in the first week of Mission 4636. Each dot represents a person contributing to the online chat-room used for collaboration between translators.

develop a similar platform, *Noula*, with a new number '177' that combined both text-messages and phone calls.

3 Emergency response outcomes

This was the only emergency reporting and response service available to people within Haiti following the earthquake. According to the responders it saved hundreds of lives and directed the first aid to tens of thousands.

The translators and Ushahidi mappers were both aided by a parallel crowdsourcing effort. Volunteers combined satellite imagery, offline maps and reports from people in Haiti using GPS devices to add thousands of data points to Open Street Map, taking the number of labeled roads and landmarks from dozens to thousands in just a few days.

4 Collaborative translation

In the first week alone more than 1,000 people came online to help translate the messages as they arrived (See Figure 2). We interacted in an online chat-room that served as a Q&A for the people newly joining the volunteer efforts, an avenue for volunteers to interact directly with the people coordinating the process, and a collaborative space for translators to work with each other.

This paper focuses on the last of these – collaborations between translators – drawing on observation from conversations within the chat-room. Collectively, they shared their knowledge of locations, regional slang, abbreviations and spelling variants to process more than 40,000 messages in the first six weeks, with most people working

S: dont know the name in english for bidonville

J: it means real poor neighborhood

M: shanty town

M: Hi, Wondering what is "akwatab"? Is this short for potable water? Thanks, M

K: "akwatab " is some kind of pill that you put in water so that it can sanitize it

R: @M - sounds like aquatab - can we more a bit more context to be sure...?

D: I need Thomassin, A please

A: Kenscoff Route: Lat: 18.495746829274168, Long:-72.31849193572998

A: This Area after Petion-ville and Pelerin 5 is not on Google Map. We have no streets name

A : I know this place like my pocket

Dalila: thank God u was here

Figure 3: Chat-room interactions showing collaboration in translation (actual names reduced to initials for confidentiality)

(physically) alone from within at least 49 countries.

4.1 The knowledge of the crowd

The combined knowledge of the translators exceeded that of any one individual. Figure 3 gives examples of where volunteers were struggling with particular translations and reached out their colleagues. Individual lexical items were the most frequent items being queried, especially as the slang varied greatly from region to region. As part of this effort, the volunteers created spreadsheets of common slang words found across the regions using online collaborative software, while others made translation lists and even first-pass machine-translation tools. Researchers at Microsoft also began importing the messages to improve their machine translation services. Collaboration was crucial for identifying locations, too. Emergencies were coming in from areas without online, labeled maps. If someone knew the location they could simply click on a map to generate the coordinates.

For the emergency responders this local knowledge was vital.

4.2 The commitment of the crowd

The most active translators still working after several weeks were, overwhelmingly, also those who had been the most active in the chat-room shortly after launch and it was the sense of community that kept many going. Not all translators collaborated in the chat room – a few were professional translators volunteering who were more used to working alone and others were simply not comfortable with that particular interaction environment. For some that did not feel comfortable participating in an online chat-room they still reported that they found comfort in having the chat-room open while they worked, knowing that they were part of a larger effort if not actually messaging themselves.

For many, the ability to contribute in real-time was in itself a motivation, combating the feeling of helplessness that accompanies being physically removed from a crisis affecting loved ones:

“When I found out about the project I thought it was a really good thing because I could help people, instead of just donating to charity”

Or simply:

“You kept us going”

4.3 The contacts of the crowd

Telecommunications between people in Haiti and their overseas friends and relatives became fairly reliably within about a week of the earthquake. Members of the Haitian diaspora who were working on translating emergency messages were also adding reports directly to the Ushahidi mapping platform. This information flowed both ways. For example, following the arrival of this message:

“Rue Casseus no 9 gen yon sant kap bay swen ak moun ki blese e moun ki brile”

Street Casseus no 9, there is a center that helps people that are wounded or burnt

one volunteer immediately stated:

“I will pass it on that is my cousins hospital”

This is the hardest part to quantify – the interactions between people working on the translations and their contacts within Haiti. However, it is clear that by the second or third week these individual communications were out-numbering the actual emergencies. Beyond one-way processing, the technology had become a collaborative space.

5 Conclusions

We learned very quickly the potential power of collaborative translation. While the emergency response service was largely successful we did not have the resources at the time to carefully investigate of the quality of the data beyond quick quality control measures. For example, what are the precise trade-offs between crowdsourced, professional and machine translation for sudden onset crises? A more thorough analysis is currently underway and we hope to share the results soon. This should help contribute to our knowledge of crowdsourced translation and how it might contribute to future crisis response strategies, especially in the context of less-resourced languages.

Acknowledgements

With far many people to thank I will limit the expression of gratitude here to the Haitian population sending reports. Their selflessness in the face of a crisis of this scale is humbling. They shared information not just about emergencies but also any information that they thought would benefit the relief efforts. Despite the tragedies surrounding them, *mèsi* (‘please’) was among the most frequent words in the reports.

References

- Chris Callison-Burch. 2009. Fast, cheap, and creative: evaluating translation quality using Amazon’s Mechanical Turk. In *EMNLP ’09: Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*
- Robert Munro, Steven Bethard, Victor Kuperman, Vicky Lai, Robin Melnick, Christopher C Potts, Tyler Schnoebelen, and Harry Tily. 2010. Crowdsourcing and language studies: the new generation of linguistic data. In *Proceedings of the NAACL-2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*.
- Robert Munro and Chris Manning. 2010. Subword modeling for text-message classification. In *Proceedings of the NAACL 2010*.
- Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew T. Ng. 2008. Cheap and fast—but is it good?: evaluating non-expert annotations for natural language tasks. In *EMNLP ’08: Proceedings of the Conference on Empirical Methods in Natural Language Processing*