

Delivering Patients to Sacré Coeur: Collective Intelligence in Digital Volunteer Communities

Kate Starbird

Department of Human Centered Design & Engineering
University of Washington
kstarbi@uw.edu

ABSTRACT

This study examines the information-processing activities of digital volunteers and other connected ICT users in the wake of crisis events. Synthesizing findings from several previous research studies of digital volunteerism, this paper offers a new approach for conceptualizing the activities of digital volunteers, shifting from a focus on organizing to a focus on information movement. Using the lens of distributed cognition, this research describes collective intelligence as transformations of information within a system where cognition is distributed socially across individuals as well as through their tools and resources. This paper demonstrates how digital volunteers, through activities such as relaying, amplifying, verifying, and structuring information, function as a collectively intelligent cognitive system in the wake of disaster events.

Author Keywords

Crisis informatics; collective intelligence; distributed cognition; crowdsourcing

ACM Classification Keywords

H.5.3 [Groups & Organization Interfaces]: Collaborative Computing, Computer-Supported Cooperative Work; K.4.2 [Social Issues].

INTRODUCTION

Bringing Patients to Hôpital Sacré Coeur

On January 12, 2010, a 7.0 earthquake struck Haiti, causing catastrophic damage and leaving in its wake great humanitarian need. In the days after the earthquake, over 300,000 injured people required medical care, a concern made logistically complex by widespread destruction to the country's infrastructure [19,23]. The epicenter of the earthquake was in southern Haiti, near the town Léogâne, and consequently the greatest need resided in the southern regions of the country where many hospitals were either destroyed or seriously damaged [19]. Mobile clinics and field hospitals were quickly established near the most affected areas [19,23], but medical teams there struggled to keep up with the volume of injured patients [24].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2013, April 27–May 2, 2013, Paris, France.

Copyright © 2013 ACM 978-1-4503-1899-0/13/04...\$15.00.

Meanwhile, in the far north of the country in Milot, about 120km from Port-Au-Prince, the Hôpital Sacré Coeur suffered mild damage, but remained fully functional. In the first few days after the quake, while injured people awaited care in other parts of the country, Sacré Coeur had capacity—74 empty beds, a medical staff on hand with more on the way, and two helipads for receiving patients by air—but no patients (personal communication, Tim Traynor). Traynor, Onsite Coordinator at the Center for the Rural Development of Milot (CRUDEM), the organization that runs Sacré Coeur, reports that the hospital tried to communicate with response agencies that it could accept and treat patients, but representatives of the large NGOs who did assess their facility determined that it was too small and too far away to meet their needs.

Believing that their hospital could help people awaiting treatment for in some cases life-threatening injuries, Sacré Coeur staff and volunteers turned to the Internet to get the word out about their facility. On January 13, they posted to their website that the hospital was open, functional, staffed, and ready for patients. Hospital representatives also turned to email, social media, and mainstream media to spread their message. Soon, other social media users, some of them *digital volunteers*, began to help promote their cause. A CRUDEM Facebook post explains what happened next:

Then, through a chain of social networking correspondences, board member Sarah Kane made an online introduction to a Coast Guard commander stationed off Haiti. [They] answered the call for assistance and began the transport of patients from Port au Prince to Milot.

Soon, patients began to stream in, including many airlifted by the US Coast Guard (USCG). On January 16, they received four patients. By the afternoon of January 18, occupancy had increased to nearly 30, and the next day 44 patients were airlifted to their hospital. They were quickly compelled to expand their capacity to accommodate increasing numbers of patients. Eventually, Sacré Coeur performed over 200 amputations for injuries suffered during the quake and treated more than 800 patients (personal communication, Traynor). Traynor claims that this would not have been possible—that his hospital would not have been able to connect with all of these patients—without social media. But what role did social media and social media users play in this incident? This paper addresses this question, exploring the use of social media,

especially the role of digital volunteers using these tools, during this and other crisis events.

Social Media Use during Crisis Events

This research contributes to the expanding literature on the use of social media during crises [e.g. 9,11,20,23,25, 26,31]. Among other uses, social media have begun to play host to *digital convergence* behavior; just as people have been known to converge physically onto the scenes of disaster events [5,15], people are now using social media and other information and communications technology (ICT) to converge “digitally” in the aftermath of disasters [11,20,26], and, like many convergers in the offline context who begin to act as volunteers [5,15], some digital convergers are taking on roles as digital volunteers [26].

The 2010 Haiti earthquake marked a turn for humanitarian responders in recognizing an emerging role for social media and its users in disaster response efforts. The Disaster Relief 2.0 Report explained, “...access to mobile and online communication enabled a kind of *collective intelligence* to emerge—when thousands of citizens around the world collaborated ... to help make sense of a large-scale calamity and give voice to an affected population” [7, p. 11].

Previous work on social media use during the aftermath of the Haiti earthquake has examined the self-organizing of digital volunteers, describing how a group of volunteer Twitterers came together to function as an emergent response organization [26]; and analyzed digital traces left by Twitter users who authored medical-themed tweets, finding the Twitter platform to have been used by smaller organizations for “beaconing” behavior—reaching out to connect to others who could help them do their work [23]. Both of these studies [23,26] focused on how Twitter was used to coordinate and organize. This research takes a slightly different perspective, shifting from a focus on organizing to a focus on information. By looking at the movement and manipulation of information shared in social media spaces after disaster events, this paper describes how “collective intelligence” emerges and offers a conceptual framework for understanding what it is and how it works.

Collective Intelligence

Lévy defines collective intelligence as “a universally distributed intelligence, constantly enhanced, coordinated in real time, and resulting in the effective mobilization of skills [distributed across individuals]” [16, p. 13]. The term has been offered as an explanation for seemingly intelligent behavior in large-scale interaction for both humans [16,17] and non-humans [e.g. 4]. Citing phenomena that occur in army ant colonies, Franks described collective intelligence as problem-solving ability within the group that exceeds the capacity of an individual, noting that, “through communication and mass action, army ants find flexible solutions to complex problems” [4, p. 139]. Though Lévy contests the attribution of collective intelligence to non-humans, suggesting a distinction between instinctive crowd behavior and the conscious and constructive participation

he claims collective intelligence requires [16], this paper resists the intentionality embedded in that claim. Though crowd work in this context is almost always intentional, individual actions often have broader effects than those intended. Here, collective intelligence is defined as a shared intelligence that is distributed across groups of individuals [16], where the problem-solving capacity of the group is greater than the capacity of a single individual [4].

Though collectively intelligent behavior in humans can be observed offline, the connected interaction enabled by ICT has opened up new ways for the phenomenon to both occur and to be observed [e.g. 16,28]. Early hints of machine-connected, collective intelligence appear in Engelbart’s ideas for augmenting the human intellect using machines; e.g. through collaborative online workspaces and hyperlinked documents [2]. Engelbart was influenced by Bush’s proposed Memex machine, which described a knowledge system based on the associative nature of human cognition [1]. Using the “trails” of hypertext, people could tap into existing knowledge in a manner more closely related to their own ways of thinking [1,2]. In recent work, Engelbart extends these views into a vision of “collective IQ,” that increases humanity’s ability to solve complex problems [3].

Real-time coordination of humanitarian response in the wake of crisis events is one of these complex problems to which the collective intelligence of the crowd can be, and is indeed already being, productively applied. This paper aims to extend our understanding of how collective intelligence manifests during crises by analyzing information trails left behind in our hyperlinked information-interaction space.

DATA AND METHODOLOGICAL FRAMEWORK

This research draws from data collected during several studies of remote crowd work during disaster events. These include an extended investigation of the “Tweak the Tweet” project [29], a study of digital volunteers during the aftermath of the Haiti earthquake [26], and long-term participant-observation of Humanity Road, a virtual volunteer organization that responds to disasters [28].

Twitter as a Tool for Digital Volunteers

Twitter has become a popular platform among digital volunteers, for whom it acts as a source of information—including information coming from the ground, a site for working on that information, and a platform for communicating with and connecting to other volunteers, affected people and response agencies [26,28]. Several affordances of Twitter contribute to its utility as a real-time, information-sharing platform during crisis events, including the fact that the vast majority of tweets are public and searchable and that social networking features of the platform allow users to follow anyone they like and permit non-reciprocal ties. Additionally, user-driven linguistic conventions—e.g. mentions, retweets, and hashtags—have served to expand Twitter’s usability and functionality.

Tweak the Tweet

The Tweak the Tweet (TtT) project proposed an extension to the functionality of the hashtag, incorporating it into a crisis-reporting microsyntax [28]. TtT asks users to place certain hashtags in front of specific pieces of information, rendering tweets not just searchable, but machine-readable. From its first deployment, TtT has been adopted by remote Twitterers who work to “translate” information they find elsewhere on Twitter or the web into TtT syntax [26].

Data

This research draws from four distinct datasets.

Dataset	Data Description
<i>Contextual Streams of Haiti Volunteeers</i>	Twitter data. Every tweet sent by 339 Haiti volunteeers from Jan 10-Feb 1, 2010. 292,928 tweets.
<i>Interviews w/ Volunteeers</i>	Interviews with 19 Haiti volunteeers. Completed in July and Aug 2010.
<i>Participant Observation in Virtual Volunteer Org</i>	Varied digital traces. Collected on Skype, Twitter & Google Docs. Jan 2011-Jul 2012.
<i>Keyword Tweets from TtT Deployments</i>	Twitter data. Collected during TtT deployments for 2011 tornados in Alabama (Apr 28-May 12, 102,395 tweets) and Joplin, MO (May 23-Jun 13, 904,710 tweets)

Contextual Streams of Haiti Volunteeers

During the TtT deployment for the Haiti earthquake (January 14-24) [26], we used the Twitter Search API to collect tweets that contained hashtags related to Tweak the Tweet, and from that set identified every Twitterer who sent a tweet or retweet in TtT syntax (339 accounts). Previous research indicates that this dataset includes a subset of an emergent digital volunteer network, many of whom identified as “volunteeers” [26]. Using the Twitter REST API, we collected the full *contextual* stream—every tweet sent between January 10 and February 1—for each of these TtT Twitterers (292,928 total tweets).

Interviews with Haiti Volunteeers

From the *Contextual Streams of Haiti Volunteeers*, researchers identified 74 TtT translators [26]. 19 of these Twitter users were interviewed about their TtT use as well as their broader digital volunteer behavior.

Participant Observation in a Virtual Volunteer Organization

The author of this paper has been a participant-observer within a virtual volunteer organization, Humanity Road (HR), for 19 months (January 2011 to July 2012). The organization is a first-of-its-kind formal non-profit in the domain of digital disaster response [28]. This research relies primarily on the digital traces of HR’s activity, almost all of which is mediated by ICT.

Keyword Tweets from TtT Deployments

During each TtT deployment, as part of the infrastructure for running the project, tweets were collected using the Twitter Streaming API to “track” keywords identified by researchers to be related to the event. This research includes tweet data collected during TtT deployments for two tornado events in the U.S. in 2011.

Using Tweets to Navigate the Larger Information Space

The majority of information on Twitter can be considered “derivative”—retweeted or otherwise repeated information that can be found elsewhere on the web [25]. Much of this derivative information points to a previous source through original author attribution or URL links. Following the methodological strategy employed by Sarcevic et al. [23], this research follows links, retweet attributions and conversations within tweet text to investigate upstream sources. This paper therefore uses tweet data to navigate the larger information space, a strategy that complements an analytical focus on information trajectories.

Note on Tweet Excerpts

Twitter accounts names are anonymized here, except where owners gave permission for direct attribution. Tweet times are in the local time of the event. Approximate times are for tweets where our data collection did not capture the original tweet, but found it through its traces, i.e. retweets.

Distributed Cognition as a Methodological Framework for Analyzing Collective Intelligence in Action

The theory of distributed cognition [10,12,22] provides a methodological framework for examining connected crowd work from an information-centered perspective. Distributed cognition (dCog) is a theory that claims cognition does not exist solely within individual minds, but is instead distributed socially, across groups of people, as well as physically (and digitally) across tools and artifacts.

A previous study examining “collective intelligence” in the crisis space noted approaching analysis from a dCog view [31], but did not elucidate how this perspective contributed to the understanding of this phenomenon. This paper aims to address that explicitly—to illustrate how dCog can help us understand and describe collective intelligence.

Though presented as a theory of cognition, research studies [e.g. 14,22] have employed distributed cognition as a methodological framework, within which “the unit of analysis is a culturally constituted functional group rather than an individual mind... (permitting researchers) to describe cognitive processes by tracing the movement of information through a system and characterize the mechanisms of the system which carry out the performance, both on the individual and the group level” [14, p. 15]. Cognition, in the dCog view, occurs as transformations of representations, and this model provides an approach for analysis at the level of these transformations. Hutchins has outlined a simple strategy for this, asking, “What information goes where, when and in what form?” [13].

FINDINGS

This paper expands on previous studies, drawing out new findings from existing data sets, and incorporating new data sets for analysis. Importantly, it offers a alternative approach for conceptualizing the activities of digital volunteers, shifting from a focus on organizing to a focus on information movement, and using the lens of distributed cognition to describe collective intelligence in action.

The Information-Processing Work of Digital Volunteers

Earlier work examined the self-organizing of digital volunteers in the wake of the Haiti earthquake, explaining how a group of previously unconnected remote volunteers came together as an emergent organization to help affected people [26]. While that original study described a range of volunteer activities—from performing as “remote operators” to organizing a campaign to fill the cell phones of affected people—this work concentrates primarily on the information-processing work of digital volunteers, activities that we see across individuals and events that have become foundational activities of an emerging practice of “crisis tweeting” and the larger sphere of digital volunteerism [28].




Figure 1. Network of Haiti Volunteeters

Digital Volunteerism within an Emergent Network

The emergent organization of digital volunteers can be visualized as a network, with volunteers as nodes in that network performing information-processing tasks. Figure 1, a diagram showing connections between Haiti voluntweeters determined through analysis of their *Contextual Streams*, provides an initial abstraction for examining digital volunteer work in this way.

In this image, dark nodes represent the 339 voluntweeters identified during the TtT collection. Grey edges connect Twitterers who retweeted or otherwise mentioned another voluntweeter’s account during the Haiti response; this diagram shows a dense network, with all but eight voluntweeters connected to others in the dataset through their tweets. The network can be imagined as one subsection in a much larger information space that contains other individuals, platforms, tools and resources.

Identifying and Amplifying Actionable Information

The most simple and most common form of crowd work during disaster events is amplification. For many Twitter users who later become digital volunteers, amplification acts as an entry-point to volunteerism. Evidence within the *Contextual Streams* and *Interviews with Haiti Voluntweeters* suggests that many began to attempt to help by retweeting

information they thought was important. Others learned the behavior by seeing other Twitterers retweeting actionable information, and then adopted that strategy themselves.

Amplifying occurs as the combination of two separate actions: first, identifying actionable information or information deemed relevant to affected people or responders; and second, rebroadcasting that information to increase its range of exposure. The example below demonstrates how the crowd does this:

@sergegilles (Jan 26, 3:03pm): I Heard that there is a Human Traffic of children in Hopital Espoir, Delmas 75. Can @UNICEF Check this out?

This tweet, a report of human trafficking, was sent by a Haitian Twitterer who was in the country during the response period. The Twitter crowd quickly amplified this tweet, retweeting it in a variety of forms. Within the *Contextual Streams of Haiti Voluntweeters*, this tweet was retweeted 16 times—most in the first hour after the original. 118 other tweets were sent referencing the same report, but with slightly different content from the original. 42 of 339 Twitterers in the set of voluntweeters we identified—or 12% of the network—sent at least one tweet about this report of child trafficking. This collective activity by separate nodes within a network of digital volunteers functioned to amplify relevant, actionable information, in this case information from a trusted source who was on the ground at the time. This example is one among hundreds of mass amplification efforts—for other reports—in the *Contextual Streams of Haiti Voluntweeters*.

In addition to functioning to identify good information, i.e. first-hand information that is actionable or helps to improve situational awareness, amplifying also works by identifying good information sources, i.e. people who are on the ground, have special insight on conditions in the affected area, or represent an official voice.

“I quickly identified the sources of good information ... the people who meant well but got tricked by hoaxers and tricksters ... and the people actually IN Haiti, both locals, journos and aid workers.”

This interview response reveals how some voluntweeters approached this aspect of their work intentionally, but other research suggests Twitterers can also perform this identification function less purposefully [27]. That study shows how seemingly uncoordinated activity of amplification across a large crowd can have an aggregate effect of adding organization to the information space, performing as a sensor network for relevant, actionable information during mass disruption events.

Identification and amplification fits within the first part of Hutchins’ leading question for analyzing dCog, “What information?” These information-processing actions also function to feed other parts of the cognitive system, for instance by providing source material for routing, verifying, and structuring activities described next.

Routing Information

Evidence within the *Contextual Streams* and *Interviews with Haiti Voluntweeters* suggests that many saw routing information as an important task within their digital volunteer work. Routing addresses the next part of Hutchins' question: "What information goes where?" Voluntweeters worked to connect certain information with people who needed or could use it. Though other channels including emails and phone calls were used as well, on the Twitter platform routing was often done using @mentions, public tweets addressed to specific accounts. The following example highlights routing behavior that occurred in the response to the report of human trafficking:

@MarkJones (Jan 26, 3:28pm): @navynews My friend @Sergegilles is getting reports of kids being trafficked @ Hopital Espoir, Delmas 75. Please contact authorities NOW!

@skinterfy (Jan 26, 4:51pm): @redcross RT @MarkJones 100% sure of source. It may be too late by the time we get there. Kids being sold @Hopital Espoir, Delmas 75

These tweets are selections from 64 different tweets sent in the wake of the trafficking report and during the resulting amplification efforts, directing that information to specific accounts. The recipient accounts were typically ones that voluntweeters thought could act on the information. In the cases of @navynews and @redcross, those accounts belonged to organizations with people in Haiti, and voluntweeters might have believed account monitors could pass the information to operatives on the ground. Below, a high-volume and highly-connected voluntweeter describes how she attempted to send actionable information to people on the ground by combining strategies for finding on-the-ground Twitterers with directing behavior:

"By searching Twitter and finding these ppl, we could send them details on where to go and who needed what... If we saw that they were headed to a particular area..."

Directed-to accounts also included media and celebrities who voluntweeters felt could further amplify the report—e.g. several directed tweets went out to @andersoncooper, a CNN reporter who covered the Haiti response. Twitterers also passed tweets on to accounts of other voluntweeters, hoping to continue to move the information if not directly to, then at least closer to someone who could use it.

Voluntweeters intentionally moved information across nodes in the network, attempting to push it directly to nodes that could use it, or move it closer to such nodes. While amplification activity is a mass broadcast out to every node in the network that an account can reach, routing includes identifying *what information* as well as *where* to send it. Routed tweets may carry a stronger, more focused signal than mass broadcasts, with recipients being more likely to tune in—while a typical Twitterer may see thousands of tweets a day while monitoring a popular hashtag stream, they will receive a much smaller number directed to them.

Verifying Information

Another way voluntweeters help shape the information flowing through Twitter and other channels is through verifying information [26]. Many of the interviewed voluntweeters reported that verifying information became an important aspect of their work. Like many other incidents, the report of child trafficking eventually became a focus of this work. Several hours after the initial tweet and after most of the amplification efforts faded, a few voluntweeters began to question the validity of the report and attempted to have someone confirm it.

@JaneSM (Jan 26, 8:31pm): @sergegilles my friend can you please elaborate on the source that came to you with the trafficking report. We thank you

@radioto (Jan 27 12:15am): Is this story is true or not ? "...Heard that there is a Human Traffic of children in Hopital Espoir, Delmas 75..."
PLEASE OFFICIAL SOURCE

In this case, verification work took place across a section of the digital volunteer network, with multiple volunteers joining in on a collaborative effort. At times, interactions like these could be contentious, because reputations were on the line. Digital volunteers may lose the trust of others if they are thought to be passing along bad information. When one openly questions information, others sometimes step in publicly to support the questioner, as did the two in this example. This verifying work had direct impact on the spread of this report. Some time later, after attempting to confirm the information by contacting his source, @sergegilles determined it to be a rumor and shared that with the network, which then worked to spread the retraction—though significantly less broadly than the original report.

Within a dCog perspective, the work of verifying speaks to the issue of "what information?" Through their public, collective work to confirm information, the crowd works as a collaborative sensor network for good and bad information as well as relevant and actionable information. The descriptive account corroborates previous quantitative research suggesting that large-scale patterns of activity can be used to identify rumors on Twitter after a disaster [18].

Though the examples presented here all revolve around the Twitter platform, verifying work moves to (and originates in) other parts of the information system as well, as volunteers look to triangulate reports and find outside sources to confirm. In interviews, some voluntweeters reported calling numbers listed as contact information in tweets to verify that information was true and still valid.

Structuring Information

Structuring, or adding structure to information, is another component of crowd work, one that addresses both "What information?" and the last component of Hutchins' question, "In what form?" Adding metadata to information and entering it into a form or a spreadsheet are types of structuring, as is TtT translation.

Focusing on individual transformations, each act of translation into TtT syntax takes place as a combination of first, the identification of actionable or otherwise relevant information; and second, the performance of one or more transformations that reconfigure the information into a state of increased structure. The following example from the *Joplin TtT Collection* demonstrates:

@TVJoe (May 23, 2011 ~12:20am): Animal shelter established in #Joplin, Missouri. Call (620) 674-3634 if you need animal shelter services.

@TheFireTracker2 (May 23, 2011 12:21:54am): MT @TVJoe: #Offer Animal shelter #Loc Joplin, Missouri. #contact Call (620) 674-3634 if you need animal shelter services. #Joplin

Shortly after the Joplin tornado in May 2011, a handful of TtT translators, many with previous experience using the syntax, worked to identify tweets with actionable information and translate them into TtT syntax. In the above example, a veteran digital volunteer, identifies information that might be useful to affected people and creates a TtT tweet by adding three TtT tags (*#offer*, *#loc*, and *#contact*) and reorganizing the tweet text to match the prescribed format. The consistent word order between the original and the translation suggests a copy and paste technique. *@TheFireTracker2* uses “MT” to indicate a “modified tweet,” a strategy adopted by TtT translators to give attribution and note that text has been changed. For practiced TtT translators, using a series of micro-structuring actions like adding a few tags and copying-pasting pieces of text, it can take several seconds to a few minutes to generate a tweet like this. These transformations create new representations that are later circulated through the system. During TtT deployments, software processes TtT-formatted tweets into public spreadsheets and Google Maps.

Few cases of TtT translation are as simple as the base translation case: TtT translators recognizing information entering the system for the first time and rapidly translating it into the format. Information often takes a less direct route from original source to structured state. The following example demonstrates this effect, revealing an important role that the larger network plays in relaying the signal through the information space until it reaches a node with the potential to process it.

@medmrsmith (May 23 01:22): #joplin Emergency Animal Hosp on 7th and Illinois behind Sonic is OPEN and taking found animals

@klboard (May 23 01:25): RT @medmrsmith: #joplin Emergency Animal Hosp on 7th and Illinois behind Sonic is OPEN and taking found animals

@Jeannie_Hartley (May 23 01:37): RT @klboard: #joplin Emergency Animal Hosp on 7th and Illinois behind Sonic is OPEN and taking found animals @franyafranya

@maryslate (May 23 01:42): RT @Jeannie_Hartley: RT @klboard: #joplin Emergency Animal Hosp on 7th and Illinois behind Sonic is OPEN and taking found animals ...

@SoutholdSMEM (May 23 01:47): RT @maryslate: RT @klboard: #joplin Emergency Animal Hosp on 7th

and Illinois behind Sonic is OPEN and taking found animals

@CrisisMappers (May 23 02:57): #tornado #offer Emergency Animal Hosp behind Sonic is OPEN & taking found animals #loc 7th & Illinois #Joplin #src @medmrsmith @SoutholdSMEM+

In this example, *@medmrsmith* initially tweets this report of an emergency animal hospital announcing that it will accept animals. Within minutes, the information bounces off three other Twitterers who retweet the text and give attribution to an upstream author, though the original author is soon excluded. Eventually the tweet is translated into TtT syntax by *@CrisisMappers*, who tracks down the original author and also cites the author immediately upstream.

Here the cognition is two-fold, measurable in both the form of the information and the movement of information across different nodes in the network. *@CrisisMappers*' TtT translation at the end of this excerpt demonstrates dCog as transformation of representation. Cognition also occurs through the actions above—the many retweets whereby nodes in the network acted to pass on the information. This relaying activity, a product of amplification, is a feature of a collaborative information-processing system where different nodes have different information-processing capabilities. Tweets are essentially ephemeral, but retweets serve to keep information alive by repeatedly posting it to the public stream [25]. Though the intermediary accounts might not have used the information themselves, their work to rebroadcast information can be seen to increase the odds that it will reach the account of someone who can use it.

Though TtT is one (relatively complex) way of structuring tweets, it is not the only way. The most common form of structuring on Twitter is simply adding hashtags. Often, volunteers comb the Twitter stream and add event hashtags to relevant information before retweeting. *Tag-added RTs* allow the original content to show up in hashtag searches. Digital volunteers can also be seen, across events, developing their own, more complex hashtagging systems. An interview with one voluntweeter revealed a subgroup of those volunteers to have created a system where they would tag tweets with *#rescuemehaiti* to call them to the attention of response agencies and other digital volunteers. Members of that sub-community tuned into this stream of *#rescuemehaiti* for source information for their other coordination activities. This example of user-driven structuring shows activity similar to TtT translation work occurring “naturally” within digital volunteer communities.

During response efforts after the Alabama tornados in the Spring of 2011, Twitterers introduced special hashtags for categorizing certain types of response-related information.

@jsandford (May 3 7:27pm): Please use #ALHaves or #ALNeeds if you have or need supplies in Alabama storm recovery to help better org tweets. #WeAreAlabama @spann

@two_slice (May 6 8:55pm): Theres a need for doctors to help National Guard soldiers w/ meds and such, in Greensboro. #WeAreAlabama #ALneeds

More than 1000 tweets were sent containing one of the two suggested terms (#ALHaves or #ALNeeds). Weeks later, when Joplin, MO was struck by a devastating tornado, digital volunteers from Alabama encouraged those responding to the Joplin event to use variations of those earlier tags (#MOneeds and #MOhaves), which many new volunteers did.

User-driven tweet structuring can be seen as a self-organizing human computation system providing real-time information processing. Structuring also taps into previous work, like amplification and verification, of other volunteers in the network:

There were some ppl that were REALLY good at [TiT translation] and they would catch many of our forgotten tweets and redo them. Sometimes we were going so fast that it was hard to remember to use it.

This suggests the emergence of a multi-tiered system where some volunteers monitor incoming streams to identify and amplify actionable information and other volunteers tune in to those amplification efforts and then add structure to the information, sometimes through multiple transformations. The cognitive system works to process information in real-time during crisis events through the collective action of individual nodes that transform information by moving it from one place to another, rebroadcasting it to keep it alive in a temporal context, or altering its form by adding micro-structure to its textual content.

Synthesizing Information into Resources

Synthesizing information is a primary activity of the Humanity Road organization. HR devotes considerable effort to creating “situation reports” during active disaster events, where they bring together multiple types of information, including damage reports, nearest airports, location and capacity information for hospitals and shelters, lists of local government and response media accounts, etc.

Synthesizing is the work of pulling information together from disparate sources to create more complete representations of an unfolding event. At the micro level, synthesis work is again enacted as transformations of representations, and like many of the activities described above—e.g. amplifying, verifying, and routing—it relies first on the work of identifying actionable or otherwise useful information during an event. The HR team works together to filter information, identify important pieces, synthesize them into a single resource, and publicize that product of their work to affected people, responders, and other digital volunteers. They complete this work together within conversations in Skype windows as well as through tweets, emails, and other channels.

Significantly, the synthesizing work of Humanity Road volunteers is also embedded within and connected to other digital volunteer communities and the massive connected crowd surrounding them. Volunteers pull information from other Twitter accounts, media reports, official websites, and blogs. Many are connected through Twitter to other digital

volunteers and in many cases, to other communities to which those volunteers belong—e.g. CrisisMappers, who work by structuring and synthesizing information into maps [7]. Information filtered by one network or community becomes input for another. HR volunteers often notice new events by seeing tweets from others in their personal networks. They also incorporate information gleaned from these connections into their resources. And other groups may get information for their resources from HR. The cognitive system of HR is plugged into a much larger system that works to process information during crises.

Emergent Organization as Cognitive Architecture

Figure 1 shows the network of connections between the 339 Haiti voluntweeters. Previous research suggests this network was as an emergent one—only a small portion of links between voluntweeters for Haiti had been established before that event [26]. Within the dCog view, these connections and this emergent organization can be seen as (part of) a “cognitive architecture” [10,12], a collection of features of the physical-social-digital environment that shapes the cognitive processes of the group. The emergence of this organization shows this architecture to also be *shaped* by the activities of the group.

Due to the public nature of tweets, the affordance of non-reciprocal ties, and the ability to send public tweets to the attention of any account, connections can and do develop between people who were previously strangers. Across the entire network, many nodes are already connected, but there remain an immense number of potential connections. For voluntweeters, their digital volunteer work—i.e. amplifying information through retweeting, directing information through mentions, marking up tweets with the “correct” hashtags—are all catalysts for creating new connections between nodes. In another view, these activities *are* these connections. And once connections are established between nodes, information may move more freely between them. Thus, the practice of crisis tweeting determined how this structure (see Figure 1) emerged *and what it looked like*. The structure, in turn, shaped how information moved during the event, demonstrating a reciprocal relationship between action and structure, one described by Giddens [6].

The Case of Hôpital Sacré Coeur in Milot, Haiti

The cognitive architecture of digital disaster response extends across networks and platforms. Returning to the example in the Introduction, the account of how Hopital Sacré Coeur used social media to connect with patients, demonstrates this.

Beaconing: Broadcasting an Urgent Message

Sarcevic et al. [23] describe how medical organizations used Twitter for *beaconing* behavior, to broadcast their message as widely as possible. Though CRUDEM did not use Twitter directly, one volunteer there, Sarah Kane, turned to ICT to beacon, sending mass email messages and posting to forums in an attempt to bring patients to the hospital. Kane was “responding” virtually, from outside

Haiti, in the days after the earthquake. This text excerpted from one of her email messages was posted January 15 to a website set up to assist in the response (www.rockmasters.com/haiti-earthquake/communications.htm):

URGENT: *We have a full-service hospital in Milot, Haiti that is *not* being utilized!! Hopital Sacre Coeur... We can take up to *200 injured patients* immediately! We have operating rooms and beds. ... We have cleared a soccer field for helicopter landing. ... We cannot understand why our facility is not being utilized! As of yesterday, we have received only two patients that came by car. ... Please get the word out to news organizations to get rescue teams... It makes no sense to be 25 minutes away by chopper with beds and doctors waiting to help! ...*

*Sarah Kane - Crudem Support
Hopital Sacre Coeur in Milot, Haiti, 555 555-5555*

CRUDEM, perhaps through Kane, also posted information to their website and made periodic updates to their Facebook page, including this one on January 17:

People are dying in the streets AND in the hospitals in Port au Prince. We have a full-service hospital with two ORs, trauma team and ortho team waiting for patients!

Kane continued to send email messages, until at least January 18, pleading for organizations to bring patients to Sacré Coeur. Many of these were directed to media outlets, urging them to “please broadcast this anywhere you can in media & send to everyone you can” (www.thetakeaway.org/blogs/takeaway/2010/jan/17/aid-agencies-struggle-bring-help) and suggesting that she needed help getting her message to the “highest levels of U.S. government” (Caribbean Free Radio website).

Relaying, Amplifying and Structuring across Platforms

Echoes of these posts show up in several different interactive spaces on the web. As with the excerpts above, several of Kane’s email messages were posted within the content of blogs and websites of other small organizations and media outlets. In at least three instances, individuals who were not connected to the organization posted information about Sacré Coeur as comments to other blogs and articles in mainstream media. For example, the following was added as a comment to a New York Times blog on January 15 (cityroom.blogs.nytimes.com/2010/01/15/the-spirit-of-port-au-prinmce-now-broken/):

A University of Iowa physician in Haiti now has found a 200-bed hospital, fully equipped, in Milot with no patients in it. He’s asking us to try to get the word to people in Haiti. Contact there is Sarah Kane, Crudem Support, Hopital Sacre-Coeur in Milot. 555-555-5555, Sarah@email.net.

This activity shows digital volunteerism in the form of relaying and amplifying, similar to what this paper has demonstrated on Twitter, occurring elsewhere online.

Not surprisingly, traces of CRUDEM’s beaconing behavior also show up on Twitter. The first tweets with links pointing to the CRUDEM website appear on January 13, and more than 600 tweets were sent about the hospital between that day and February 1. Most involved trying to

get patients to the facility, though some tweets in the latter part of the collection window dealt with a new issue of obtaining resources (food, cots, etc.) for the hospital. Volunteers worked to move information from other sites onto Twitter and also to amplify messages posted there. Over time, they added more information to the tweets, e.g. listing the contact info for Kane, noting the availability of a landing pad, and including GPS coordinates:

@rqskye (Jan 15): @carelpedre-SacredHeart Hosp, Milot, also has landing pad cleared to accept 200 patients NOW. Call Sarah Kane Crudem 555-555-5555

@redcrossmom (Jan 18): @firesideint Milot (19° 37' 0" N, 72° 13' 0" W) Sacre Coeur is open #Haiti Milot,555-555-5555:FullStaff, 2ORs,TeamsWAITING.

As these tweets show, volunteers often directed information about Sacré Coeur to those they thought could help. They also used TtT syntax to structure tweets, sending 66 TtT tweets and retweets about the hospital.

Other structuring activity moved to Ushahidi, a “crisis mapping” platform that was accepting reports about Haiti via SMS and a web portal [7]. Someone, perhaps a digital volunteer, added information about Sacré Coeur to the map on January 16. Later tweets linked to this Ushahidi report.

In addition, two CNN iReports were created about Sacré Coeur, both by individuals unaffiliated with CRUDEM. The first was posted by a woman in Haiti on the morning of January 15. Later that morning, information in that post began to spread through the tweet streams of digital volunteers. CNN picked up the story of Sacré Coeur for the first time that evening, mentioning the hospital on air during the AC360 program. From their transcript:

There's apparently a hospital north of here that actually has room for patients. It's a hospital called Sacre-Coeur, it's located in Milot...80 miles north of Port-au-Prince...We are told, I've not seen this with my own eyes but we're told that it's fully operational ...but again, it's far away. A lot of people can't get there.

It is not clear how CNN reporters first heard about Sacré Coeur, but they and other mainstream media were often the focus of CRUDEM’s beaconing as well as digital volunteers’ directing activities. Both groups appear to have perceived those channels as having the ability to further amplify their messages.

Delivering Patients to Milot

Eventually, patients began to arrive in Milot. On January 16, the day after the CNN report, the U.S. Navy delivered one group of injured people to Sacré Coeur, an event remarked upon in a Facebook post by CRUDEM that also included another plea for more patients—suggesting that the coordination problem had not yet been fully resolved.

The USCG airlifted two groups of patients, nine total, to Sacré Coeur on the 17th, a day that witnessed the most Twitter activity about the hospital (155 tweets). Though there is no evidence of a direct or causal relationship in the

tweet streams, it is possible that Twitter activity and the delivery of patients are themselves related to a common factor: the information, moving in and across networks of social media users and digital volunteers, was beginning to reach nodes that could use it in forms that were usable.

On January 18, while Kane's pleas for patients continued to spread throughout the web, the USCG began to deliver a steady stream of patients. Others arrived via automobile. On the 23rd, tweets began to contain requests for supplies for the facility, which had by then begun to exceed capacity.

Previous studies were careful to avoid claims of impact by digital volunteers during Haiti [23,26], as it is difficult to connect social media activity to actions on the ground. However, evidence in the digital traces analyzed here suggests a role for social media and digital volunteers in delivering patients to Sacré Coeur. According to CRUDEM's Facebook page, Kane eventually made contact with a USCG Commander, and after that the USCG began to coordinate a steady flow of patients to the facility. It is unclear in the digital traces exactly when or how the connection was made, but CRUDEM credits social media. That the USCG first contacted CRUDEM through Kane, a volunteer who was neither located in Haiti nor acting as an official leader of the organization, suggests that Kane's social media posts—and the efforts by others to help distribute the information in those posts—played a role.

It is likely that CRUDEM's message reached people who could act on it through multiple paths. With repeated movement across platforms and media, there was considerable redundancy in the network. It may be hard to assign credit for generating this connection to a single node of the digital volunteer network or to a single post or digital artifact. However, if digital volunteers had not been acting on this information, the connection may not have happened, or it may not have happened as quickly—a critical concern. Traynor recalls that a member of one USCG helicopter team, during one of their first patient deliveries, remarked, "God, had we only known you were here earlier!"

To answer the question posed earlier, social media did not deliver patients to Sacré Coeur, but volunteers did use these tools to coordinate action that saved lives. The perspective advocated here views this not as any individual act by any one person, but collective action by—and collective intelligence within—a network of people, connected through ICT, who helped organize the information space through many small acts of transformation.

CONCLUSION

Collective intelligence, along with the related concept of "crowdsourcing," has become a popular term for characterizing crowd activity online. But what is collective intelligence and how does it work?

Describing his view, Franks wrote, "all intelligence involves rational manipulation of symbolic information" [4, p. 139, citing 8]. Aligned with this view, this empirical

analysis of digital traces within the domain of digital disaster response demonstrates how collective intelligence manifests as transformations of information across an ICT-connected crowd. Using the theory of distributed cognition as a methodological framework, this paper shows how seemingly simple actions from individuals collectively organize the information space.

Though the theory of distributed cognition is intended to apply to all systems, the methodological framework has most often been employed to examine interactions between groups of individuals—and their tools—collocated in physical space [e.g. 10,12,14,22]. This research extends the dCog perspective to examine collective action within the massively connected information space of social media and the surrounding Internet. Though a previous study on collective intelligence in social media spaces used this theory as part of a broader analytical frame [31], this paper explicitly employs dCog as a methodological framework for examining the work of digital volunteers.

The dCog framework offers an effective tool for thinking about how this work is coordinated and organized—how these individuals are able to come together as remote response organizations in some instances and massive sensor networks in others. It allows us to conceptualize how this organization emerges without preplanning and without any single individual or node having a complete grasp of the entire system. Heverin and Zach describe how, in the wake of violent crises, individuals come together on social media sites to make sense of the event, a collective sense that can only develop in interaction with others [9]. Digital volunteers go beyond collective sensemaking to collective action, so their goals may be different, but the processes whereby they achieve those goals are similar. Using the dCog methodological framework, we see how local actions create global effects, and how the cognition of the system extends far beyond the cognition of any individual.

dCog also gives us a lens through which to view this emergent activity, by focusing on transformations of information. This approach fits well with the affordances of social media platforms that leave the digital traces of many of these transformations in public view and allow us to easily "step inside the cognitive system" [12].

Collective intelligence can be seen to emerge from the individual actions of digital volunteers to relay, amplify, verify, structure and synthesize information. These actions also form a cognitive architecture [12] for the volunteer network, one that is itself plugged into the cognitive architecture of the larger information space, and one that structures future actions, e.g. by establishing information pathways between trusted volunteers.

Digital volunteerism is becoming an established feature of the online ecosystem after disaster events. It has begun to stabilize, somewhat, as volunteers establish commonly understood ways of doing things and organize into formal

virtual volunteer organizations. But it also continues, in event after event, to take on new forms and emerge within new networks of individuals. This research offers a detailed description of the core activities of digital volunteerism and explores the utility of using dCog to both analyze and conceptualize this collective activity online.

ACKNOWLEDGMENTS

This work would not have been possible without the support of Project EPIC at the University of Colorado. I thank Leysia Palen for her guidance throughout this multi-year research effort, and Ken Anderson and Aaron Schram for helping collect the Twitter data used in this analysis. I also thank Ed Hutchins for his early guidance, the reviewers of this paper for their thoughtful feedback, and Taylor Scott for his feedback on a late draft. Early research for this paper was supported by an NSF Graduate Research Fellowship and NSF grants IIS-0546315 and IIS-0910586.

REFERENCES

- Bush, V. As We May Think, *Atlantic Monthly*, (July 1945), 101-108.
- Engelbart, D. *Augmenting Human Intellect: A Conceptual Framework*. AFOSR-3233, Summary Report. Stanford Research Institute, 1962.
- Engelbart, D. Augmenting society's collective IQs. *Proc. of Hypertext and Hypermedia* (2004). New York: ACM, 1-1.
- Franks, N. Army Ants: A Collective Intelligence: A neural network seems an apt analogy as a colony of army ants navigates the tropical rain forest. *American Scientist*, 77, 2 (March-April 1989), pp. 138-145.
- Fritz, C. E. & Mathewson, J. H. *Convergence Behavior in Disasters: A Problem in Social Control*, Committee on Disaster Studies, National Academy of Sciences, National Research Council, Washington DC, 1957.
- Giddens, A. *Central problems in social theory: action, structure and contradiction in social analysis*, University of California Press, Berkeley, 1979.
- Harvard Humanitarian Initiative. Disaster Relief 2.0: The Future of Information Sharing in Humanitarian Emergencies. Washington, D.C. & Berkshire, UK: UN Foundation & Vodafone Foundation Technology Partnership, 2011.
- Haugeland, J. *Artificial Intelligence: The Very Idea*, MIT Press, 1986.
- Heverin, T. & Zach, L. Use of Microblogging for Collective Sense-making during Violent Crises: A Study of Three Campus Shootings. *J. Am. Soc. Inf. Sci. Technol.* 63, 1 (January 2012), 34-47.
- Hollan, J., Hutchins, E. & Kirsh, D. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Computer-Human Interaction*. 7, 2 (June 2000), 174-196.
- Hughes, A., L. Palen, J. Sutton, S. Liu, & S. Vieweg. "Site-Seeing" in Disaster: An Examination of On-Line Social Convergence. *Proc. of ISCRAM 2008*.
- Hutchins, E. *Cognition in the Wild*. Cambridge, Mass: MIT Press, 1995.
- Hutchins, E. The Social Organization of Distributed Cognition. Course Lecture, COGS 102a, UCSD, 2011. <http://hci.ucsd.edu/102a/10Lectures/SODC-basics.pdf>
- Hutchins E. & Klausen, T. Distributed Cognition in an Airline Cockpit. In Middleton, D. & Engeström, Y. (Eds.) *Communication and Cognition at Work*. (pp. 15-34) Cambridge University Press, 1996.
- Kendra, J. M. & Wachtendorf, T. Reconsidering Convergence and Converger: Legitimacy in Response to the World Trade Center Disaster, *Terrorism and Disaster: New Threats, New Ideas: Research in Social Problems and Public Policy*, 11, (2003), 97-122.
- Lévy, P. as translated by Bononno, R. *Collective Intelligence*. Perseus Book, USA, 1997.
- Malone, T. W., Laubacher, R. & Dellarocas, C. N., Harnessing Crowds: Mapping the Genome of Collective Intelligence (February 3, 2009). MIT Sloan Research Paper No. 4732-09.
- Mendoza, M. Poblete, B. & Castillo, C. (2010). Twitter under crisis: can we trust what we RT? In *Proceedings of the First Workshop on Social Media Analytics (SOMA '10)*, 71-79.
- PAHO/WHO. Earthquake in Haiti—One year later. *Report on the Health Situation*, January 2011.
- Palen, L & Liu, S. B. Citizen Communications in Crisis: Anticipating a Future of ICT-Supported Participation, *Proc. of CHI 2007*, 727-736.
- Pape, J.W., Johnson, W.D. and Fitzgerald, D.W. The earthquake in Haiti—Dispatch from Port-au-Prince. *New England Journal of Medicine* 362 (2010), 575-577.
- Rogers, Y. & Ellis, J. Distributed Cognition: An Alternative Framework for Analyzing and Explaining Collaborative Working. *J. of Information Technology*, 9, 2 (1994), 119-128.
- Sarcevic, A. Palen, L., White, J., Starbird, K., Bagdouri, M. & Anderson, K. "Beacons of Hope" in Decentralized Coordination: Learning from On-the-Ground Medical Twitterers during the 2010 Haiti Earthquake. *Proc. of CSCW 2012*, 47-56.
- Sheridan, M.B., Wilgoren, D., & Leonnig, C.D. Patients Overwhelm Medical Teams at Haiti Clinics. *Washington Post*. (Jan 17, 2010). <http://www.washingtonpost.com/wpdyn/content/article/2010/01/17/AR2010011701144.html>
- Starbird, K, Palen, L, Hughes, A & Vieweg, S. Chatter on The Red: What Hazards Threat Reveals about the Social Life of Microblogged Information. *Proc. of CSCW 2010*, 241-250.
- Starbird, K. & Palen, L. "Voluntweeters:" Self-Organizing by Digital Volunteers in Times of Crisis. *Proc. of CHI 2011*, 1071-1080.
- Starbird, K. & Palen, L. (How) Will the Revolution be Retweeted? Information Propagation in the 2011 Egyptian Uprising. In *Proc. of CSCW 2012*, 6-17.
- Starbird, K. & Palen, L. Working and Sustaining the Virtual Disaster Desk. *Proc. of CSCW 2013*.
- Starbird, K. & Stamberger, J. Tweak the Tweet: Leveraging Microblogging Proliferation with a Prescriptive Syntax to Support Citizen Reporting. *Proc. of ISCRAM, 2010*.
- Surowiecki, J. *The Wisdom of Crowds*. First Anchor Books Edition. New York: Anchor Books, 2005.
- Vieweg, S., Palen, L., Liu, S., Hughes, A., & Sutton, J. Collective Intelligence in Disaster: Examination of the Phenomenon in the Aftermath of the 2007 Virginia Tech Shooting. *Proc. of ISCRAM, 2008*.

