

**Carnegie Mellon University**

---

**From the Selected Works of Adam Hodges**

---

April, 2014

# Rethinking Context: Leveraging Human and Machine Computation in Disaster Response

Sarah Vieweg, *Qatar Computing Research Institute*  
Adam Hodges, *Carnegie Mellon University*



Available at: <https://works.bepress.com/adamhodges/50/>

# Rethinking Context: Leveraging Human and Machine Computation in Disaster Response

**Sarah Vieweg**, *Qatar Computing Research Institute*

**Adam Hodges**, *Carnegie Mellon University*

---

**Human-computer systems that treat context simply as enumerated facts, rules, or axioms about the surrounding physical and social environment will always have trouble handling information requiring human pragmatic interpretation. One way to overcome such limitations is to integrate human computation in the form of microwork with tasks at which machines excel, such as big data extraction.**

---

**R**esearch on context awareness in computing has thus far largely focused on tasks such as identifying user location, user modes of locomotion, and device position.<sup>1</sup> These tasks rely on an understanding of *context* that is synonymous with an objective description of the physical and social environment in which devices and users are embedded.<sup>2</sup> However, context also consists of implied knowledge, assumptions, and cultural scripts that give meaning to a situation.<sup>3</sup> Whereas machines excel at enumerating explicit contextual facts—location, noise level, and so

on—they are no match for humans at processing implicit contextual information.<sup>4</sup>

We argue for an approach to context-aware computing that takes into account the strengths of humans as well as machines. In particular, we discuss how human computation in the form of *microwork* might be more effectively leveraged to draw upon human pragmatic awareness to create human-computer systems capable of both extracting large quantities of data and processing that data in a way that is meaningful to users.

We ground our discussion within the field of crisis informatics with a focus on the use of microblogging in mass emergencies—in particular, those arising from natural hazards such as earthquakes, wildfires, and floods. Disaster response is an ideal lens through which to consider context-aware computing because of the need for time-sensitive information processing. During and after a disaster, microblogging services like Twitter can provide considerable relevant information, but so much data is broadcast that it is impossible for humans to manually locate what is most important to them quickly and easily. By training machines to intelligently sift through tweets and pass these tweets on to users, we have the potential—and the opportunity—to provide government response agencies, NGOs, individual aid workers, and other stakeholders with a set of Twitter communications more likely

to contain personally meaningful information, and to do so more rapidly.

Currently, however, the inability of computers to derive pragmatic meaning presents a challenge for researchers interested in creating tools that can both extract and interpret big data. At issue is how a better understanding of human contextual awareness can lead to the development of more effective human-machine symbiosis.<sup>5</sup>

### MICROBLOGGING IN MASS EMERGENCIES

Internet technologies have changed the way people communicate during crises. This is especially true of microblogging services, online platforms that let users post and exchange messages of limited length (140 characters in the case of Twitter). In addition, users form networks by “following” and being “followed” by other microbloggers. Users can view posts by anyone who maintains a public profile, though some choose to make their profiles private, enabling only approved followers to view their posts or communicate with them via the platform.

When a disaster occurs, affected populations, response and relief organizations, and concerned outsiders often turn to microblogging services and social media sites to help gather and circulate timely, relevant information.<sup>6,7</sup> Among other things, people seek to find out where loved ones are located, whether those at the scene should evacuate, what weather conditions are predicted, and the types and locations of available shelter. Microblogging services such as Twitter also provide a medium to quickly ask various questions and receive answers. Users thus have access to a larger pool of resources than are otherwise provided by conventional methods of communication, such as radio broadcasts, as well as time-critical situational awareness.

### SITUATIONAL AWARENESS

Following Mica Endsley, Betty Bolté, and Debra Jones, we define situational awareness as the state of knowing what is happening in one’s surroundings and understanding what the information means for a given situation, including perception of the elements in the environment and how those elements relate to one another.<sup>8</sup> In their study of air traffic control processes and practices, Endsley and her colleagues point out that air traffic controllers not only must monitor activity in the skies, but also need to grasp the integrated meaning of what they perceive—they must understand the situation as a whole to make optimal decisions.

Tweets contain many types of information that can be used during mass emergencies to inform better decision-making,<sup>6,7,9</sup> but achieving situational awareness involves more than simply reading tweets. Crucially, Twitter users also draw upon their *pragmatic knowledge* to interpret the information communicated in tweets.

Pragmatics is the study of contextual meaning, or the way context contributes to meaning. Pragmatics is

typically juxtaposed with semantics, or the study of denotational meaning as it is encoded in words and sentences. Although there is no sharp line between pragmatics and semantics, pragmatics encompasses aspects of meaning that go beyond what is explicitly communicated to include various forms of implicitly conveyed information—for example, presupposition and suggestion—that requires contextual knowledge to interpret.

### CONTEXT

Given the importance of context in pragmatics, it is useful to rethink the definition of context. Doing so is more than a mere theoretical exercise: because conceptions of context shape the development of context-aware systems, understanding what context means from a pragmatic perspective holds important implications for future systems.

Although the complexity of context is well recognized both outside and within the human-computer interaction community,<sup>3,10-12</sup> context-aware computing researchers have typically concentrated on the more tangible and tractable features of context that machines can readily sense,

---

**Twitter users draw upon their pragmatic knowledge to interpret the information contained in tweets.**

---

such as location and noise level. This focus is understandable given the nature of technological development: we often break down a technical challenge into smaller, more manageable ones. However, this approach becomes problematic when the goal is to develop more intelligent systems capable of dealing with context in all its human complexity.

Traditionally, context has been viewed as synonymous with the “environment” or “situation.” Context, then, includes everything that surrounds some sort of entity engaged in an interaction. For language scholars, this entity is a person involved in linguistic interaction; for technologists, it can be a device user as well as the device itself. With this understanding of context, the analytic focus turns to describing “the context” by enumerating various features of the setting and its surroundings. In this way, context becomes “any information that can be used to characterize the situation”<sup>11</sup> or “that attempts to describe something about the conditions in which an application executes.”<sup>12</sup> Although such information involves facts about the physical world in which interaction takes place, it also extends to “pertinent supplementary facts, rules, or axioms”<sup>12</sup> about the social world.

Attempts to define context by human-computer interaction, user-centered computing, and context-aware computing researchers resemble the ways that

anthropologists and sociocultural linguists have theorized about context. Different formulations pinpoint different aspects of the context the analyst wishes to highlight, with each definition attempting to break down the enormity of context into objectively describable chunks.

Given that context is “whatever might be considered relevant to the current act from sleeping to printing a document,”<sup>10</sup> reducing context to an objective description of the surroundings prioritizes the analyst’s ideas of what is important to the process of deriving meaning from that situation. In addition, context comes to be seen as something external to the interaction rather than a dynamic part of an encounter where meaning is produced.<sup>14</sup> In sum, traditional treatments reify context and treat it as an object describable in positivistic terms, even while recognizing the impossibility of that task. In light of these problems, how else might context be defined?

Our view builds on Lucy Suchman’s “situated action” perspective,<sup>4</sup> which “suggests a deeper role for context in interaction,”<sup>5</sup> and on the shift among sociocultural linguists away from the reification of context as an object

---

**Because much information is implicitly conveyed in tweets, microworkers can help machines to interpret the significance of content in Twitter communications.**

---

separate from interaction toward an understanding of *contextualization* as an active process of meaning negotiation. Crucially, language and technology users arrive at meaning interactionally—that is, meaning is a dynamic sensemaking process that not only relies upon an interaction’s semantic content and a catalogue of facts about the surrounding physical and social world, but also upon a rich array of resources that we use to give meaning to the world.

Whereas the dominant view within the human-computer interaction community is that context-free propositional content ultimately carries meaning, we hold that the meaning of any utterance—which we define as any human communication mode, including tweets—is contingent and interactionally derived. From this perspective, attempts by machines to deterministically describe context will eventually encounter difficulty because the descriptive task takes place amidst an ever-shifting interpretive landscape. Humans, however, excel at navigating this landscape.

### LEVERAGING HUMAN COMPUTATION THROUGH MICROWORK

In seeking to develop context-aware systems capable of collecting and identifying information relevant to any number

of situations, researchers and practitioners are increasingly turning to microwork: small, computer-based tasks performed by humans. We view microwork as a form of human computation, which is “a paradigm for utilizing human processing power to solve problems that computers cannot yet solve.”<sup>15</sup> For us, these problems involve interpreting tweets. More specifically, we are interested in harnessing human pragmatic intelligence to categorize tweets that concern disasters into situational awareness “information types.”<sup>6,7,9</sup>

A previous analysis outlined 35 information types that can benefit people in a mass emergency including reports of damage; preparation and evacuation plans; details about dead, missing, and injured persons, incidents of crime; requests for help; shelter options and locations; public property status; and weather updates.<sup>7</sup> As more research takes place, this list will inevitably grow. Notably, a single tweet can contain up to five information types.<sup>7</sup>

Because much information is implicitly conveyed in tweets, microworkers can help machines to interpret the significance of content in Twitter communications. Tagged tweets not only can be disseminated more effectively to end users, but they can also be used to form datasets for supervised machine learning.


### ARTIFICIAL INTELLIGENCE FOR DISASTER RESPONSE

While services such as Amazon Mechanical Turk, Crowdcrafting, and Samasource provide a means for people to complete any number of microtasks, successful identification of tweets containing actionable information in disaster situations requires a system that can ingest tweets and then let people easily tag those tweets with relevant category information very quickly—ideally, in near real time as the situation unfolds. Toward this end, the Qatar Computing Research Institute is developing the Artificial Intelligence for Disaster Response (AIDR) platform.

Figure 1 outlines AIDR’s multistage process of tweet collection, human annotation, and classification and dissemination.

For a given disaster, AIDR automatically collects tweets containing keywords chosen by a system administrator based on an initial interpretive assessment of Twitter content. For example, an administrator searches for “Pakistan earthquake” or “Florida hurricane” and quickly determines what keywords and hashtags are present in relevant tweets. The administrator can also instruct the system to collect tweets geotagged in a certain region.

As it collects tweets, the system shepherds them to microworkers, who read each tweet and tag them with an information category previously designated by an AIDR administrator. Categories may be high-level, such as “related to situational awareness” or “not related to situational awareness,” or fine-grained, such as “evacuation” or “offer of shelter.” For quality assurance, each tweet is tagged by



**Figure 1.** The Artificial Intelligence for Disaster Response (AIDR) platform leverages both machine and human computation to collect, process, and disseminate relevant information from tweets broadcast during a natural disaster to end users.

three microworkers, at least two of whom must agree on the category.

AIDR uses the annotated tweets to train an automatic classifier, which classifies tweets containing actionable information.

Due to the intricacies and nuances of human communication, and our reliance on pragmatics in interpreting meaning, it is unlikely that any automatic classifier—no matter how well-trained—will be able to locate all Twitter content relevant to situational awareness. Humans thus must continuously verify that the system is correctly classifying tweets.

### HUMAN PRAGMATIC INTERPRETATION

To illustrate the uniquely human capability of pragmatic processing, we provide several examples from previously collected Twitter datasets.<sup>7</sup>

Consider the following tweet from the 12 January 2010 earthquake in Haiti:

Watching 60 minutes on Haiti doctors cutting off a Childs leg with a saw ... Jesus help these people.

An estimated 200,000 Haitians lost a limb due to the earthquake, and the distressing phrase “doctors cutting off a Childs leg with a saw” indicates that the child mentioned likely had what is termed a “crushing injury” requiring amputation to avoid infection or further injury. However, due to the implicit nature of the communication, a machine likely could not accurately categorize this tweet as containing information about injuries. Although the machine could identify the word “doctors” and thereby deduce that the content relates to a medical situation, it could not infer that the tweet is specifically conveying information about an injury. In contrast, humans can easily recognize the tweet’s reference to amputation because they understand that this type of extreme treatment is often necessary in critical situations such as in the aftermath of an earthquake.

The following tweet is from the same event:

Heroes’ actor looking for parents – Haitian native Jimmy Jean-Louis learned that a house he had grown up in in Haiti ... <http://ow.ly/16kJ6G>

In this example, a particular person is “looking for” his parents, which indicates that their location or whereabouts are unknown. The semantic meaning of this phrase is likely insufficient for a machine to correctly categorize this tweet as being about missing persons, but humans can draw upon contextually relevant pragmatic information—we only look for people or things when we do not know where they are—to make this inference.

Another example comes from an outbreak of wildfires that ravaged central and southern Oklahoma and parts of north Texas in April 2009:

It is amazing how people are coming together and helping to water neighbors yards ... <http://www.koco.com/news/19140975/detail.html>

One way people prepare for the threat of fire to their homes is to water their surroundings—especially highly flammable areas covered by grass, brush, and trees—to deter flames from starting or spreading. The phrase “people are coming together and helping to water neighbors yards” in the tweet alludes to such preparatory measures, but it is doubtful that a machine would make this inference based solely on the mention of watering yards.

As these examples show, there is much more to understanding context than cataloguing a potentially infinite array of facts, rules, and axioms. It is difficult for machines to model behavior, because humans are unpredictable and wont to act capriciously in any given situation,<sup>10</sup> and we constantly improvise in response to our environment and circumstances.<sup>4</sup> Yet where machines fail in such situations, people are uniquely qualified to process pragmatic inferences to richly interpret human language.

### TOWARD MORE EFFECTIVE HUMAN-COMPUTER SYMBIOSIS

Our research yields important insights into context-aware computing. First, it underscores the inadequacy of relegating complicated forms of pragmatic processing to machines alone; using both human and machine computation is the key to developing more effective information processing systems, particularly in time-sensitive situations such as mass emergencies. In the case of AIDR, machines are tasked with extracting tweets that likely

contain situational awareness information regarding a specific disaster, while humans are tasked with using pragmatic knowledge to label those tweets using a predefined rubric of information types.

Moreover, given that rich interpretation of context relies on often unspoken, shared understanding among a community or group, criteria for selecting microworkers to perform manual tagging must take into account their familiarity with or willingness to learn about a particular mass emergency. In other words, having a general knowledge of natural disasters—although necessary—is not sufficient to accurately label tweets if the microworker lacks a common-ground starting point with those tweeting information during a specific event.

For example, in the case of a devastating typhoon such as the one that struck the Philippines on 8 November 2013, microworkers labeling tweets should be familiar with the impacted country's geopolitical organization—such as the names of regions, provinces, cities, and municipalities—and institutions. This is not to say that outsiders without deep knowledge of the country would be automatically disqualified from interpreting tweets broadcast during the typhoon, but such individuals would need to do more work—in the form of quick searches on Wikipedia, for example—to first obtain the necessary background information for the task.

The importance of common ground in pragmatic processing can be seen in another example from the Twitter datasets,<sup>7</sup> collected during the flooding of the Red River in North Dakota and Minnesota in March–April 2009:

The river is now over 40 feet and there has been a breach in the dike; neighborhood a few blocks south of us is evacuating.

The linguistic concept of *markedness* connotes the tacit understanding of certain places, landmarks, people, items, and so on when referred to in general terms. In this tweet, “the river” is an unmarked, assumed entity—the phrase itself does not specifically indicate the Red River, but the reference is understood within the local context.

Conversely, during the Red River flood, a tributary—the Sheyenne River—was also at risk for flooding. When communicating about this river, it was necessary for Twitter users to “mark” the Sheyenne, as seen in the following example:

The Sheyenne River went up 10 feet in 12 hours ...

Thus, microworkers who have greater familiarity with an area impacted by a disaster are in a better position to effectively interpret tweets sent by those affected or others with knowledge of the locale. The importance of common ground in interpreting tweets likewise extends to cultural knowledge, norms, and shared contextualization cues.

In short, all microworkers are not created equal when it

comes to pragmatic interpretation. A crucial question for designers of systems that harness human computation is how to train microworkers to more effectively carry out an assigned task. In the case of a disaster response platform like AIDR, this may involve selecting individuals who are knowledgeable about the part of the world where a disaster has occurred or providing microworkers with a synopsis of the disaster along with links to news summaries, Wikipedia articles, and other Web resources they can consult to gain a foundational understanding of the affected region. It is also helpful if microworkers spend time learning about past disasters, which provides valuable background information, to help interpret future communications. Overall, the more common ground there is between microworkers and the authors of tweets (or other microblogging data), the richer the interpretive possibilities.

**W**e argue for an understanding of context that more closely aligns with the way humans draw upon resources to make sense of communicative interactions. Context-aware systems that treat context simply as enumerated facts, rules, or axioms about the surrounding physical and social environment will always have trouble handling information requiring human pragmatic interpretation. One way to overcome such systems' limitations is to integrate human computation with tasks at which machines excel, such as big data extraction. Our work shows that during mass emergencies, hybrid systems that draw upon the strengths of humans as well as machines can more quickly and effectively process situational awareness information in microblogging sites such as Twitter. **□**

## References

1. P. Lukowicz, A. Pentland, and A. Ferscha, “From Context Awareness to Socially Aware Computing,” *IEEE Pervasive Computing*, vol. 11, no. 1, 2012, pp. 32–41.
2. T.P. Moran and P. Dourish, “Introduction to This Special Issue on Context-Aware Computing,” *Human-Computer Interaction*, vol. 16, no. 2, 2001, pp. 87–95.
3. P. Dourish, “Seeking a Foundation for Context-Aware Computing,” *Human-Computer Interaction*, vol. 16, no. 2, 2001, pp. 229–241.
4. L.A. Suchman, *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge Univ. Press, 1987.
5. J.C.R. Licklider, “Man-Computer Symbiosis,” *IRE Trans. Human Factors in Electronics*, Mar. 1960, pp. 4–11.
6. S. Vieweg et al., “Microblogging during Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness,” *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI 10)*, 2010, pp. 1079–1088.
7. S. Vieweg, “Situational Awareness in Mass Emergency:

- A Behavior and Linguistic Analysis of Microblogged Communications,” PhD dissertation, Alliance for Technology, Learning and Society (ATLAS) Inst., Univ. of Colorado Boulder, 2012.
8. M.R. Endsley, B. Bolté, and D.G. Jones, *Designing for Situation Awareness: An Approach to User-Centered Design*, Taylor & Francis, 2003.
  9. S. Verma et al., “Natural Language Processing to the Rescue? Extracting ‘Situational Awareness’ Tweets during Mass Emergency,” *Proc. 5th Int’l AAAI Conf. Weblogs and Social Media (ICWSM 11)*, pp. 385–392.
  10. V. Bellotti and K. Edwards, “Intelligibility and Accountability: Human Considerations in Context-Aware Systems,” *Human-Computer Interaction*, vol. 16, no. 2, 2001, pp. 193–212.
  11. A.K. Dey and G.D. Abowd, “Towards a Better Understanding of Context and Context-Awareness,” *Proc. 1st Int’l Symp. Handheld and Ubiquitous Computing (HUC 99)*, LNCS 1707, Springer, 1999, pp. 304–307.
  12. P. Mehra, “Context-Aware Computing: Beyond Search and Location-Based Services,” *IEEE Internet Computing*, vol. 16, no. 2, 2012, pp. 12–16.
  13. G.D. Abowd et al., “Context-Aware Pervasive Computing,” *IEEE Pervasive Computing*, vol. 1, no. 3, 2002, pp. 22–23.
  14. R. Bauman and C.L. Briggs, “Poetics and Performances as Critical Perspectives on Language and Social Life,” *Ann. Rev. Anthropology*, vol. 19, 1990, pp. 59–88.
  15. L. von Ahn, “Human Computation,” PhD dissertation, School of Computer Science, Carnegie Mellon Univ., 2005.

*Sarah Vieweg* is a scientist in the Social Computing group at the Qatar Computing Research Institute. Her research focuses on the use of social media and microblogging during mass emergency and humanitarian crises in an effort to inform the development of tools that help stakeholders gain relevant information in time-sensitive situations. Vieweg received a PhD in technology, media, and society from the University of Colorado Boulder. Contact her at [svieweg@qf.org.qa](mailto:svieweg@qf.org.qa).

*Adam Hodges* is a visiting assistant professor of English at Carnegie Mellon University in Qatar. His research draws from sociolinguistics, linguistic anthropology, and discourse analysis to examine the social, cultural, and political dimensions of language use. Hodges received a PhD in linguistics from the University of Colorado Boulder. Contact him at [adam.hodges@cmu.edu](mailto:adam.hodges@cmu.edu).



Selected CS articles and columns are available for free at <http://ComputingNow.computer.org>.