



# Citizen Sensing, Social Signals, and Enriching Human Experience

Amit Sheth • Kno.e.sis Center, Wright State University

A computer which can calculate the Question to the Ultimate Answer, a computer of such infinite and subtle complexity that organic life itself shall form part of its operational matrix. And you yourselves shall take on new forms and go down into the computer to navigate its ten-million-year program! Yes! I shall design this computer for you. And I shall name it also unto you. And it shall be called ... The Earth." —Douglas Adams, *The Hitchhiker's Guide to the Galaxy*

**D**ouglas Adams's vision — the Earth transformed into a supercomputer powered by human intelligence — was fictional but reflects the potential of the most recent advances in science and technology to transform our planet into a powerful computing platform. With 6 billion human inhabitants acting as processing nodes, Earth could indeed become the computer that provides the best answers to life's most complex and difficult questions ([http://icsc.eecs.uci.edu/abstract\\_wed1.html](http://icsc.eecs.uci.edu/abstract_wed1.html)).

It might seem like science fiction at first blush, but with the Internet serving as the communication backbone that connects us all, we could reach this point sooner than we think. When *Time* magazine named "you" as its person of the year in 2006, it captured the infinite possibilities brought forth by connecting humans and providing a platform to harness their collective intellect, knowledge, and experiences. As much as we can't question the role technology has played in fostering this new era of computing, central to its success has been the participation of people from all walks of life. Through each of our small but significant and sustained contributions, we've created and maintained vast repositories such as Wikipedia. We're also helping machines organize the world's online resources by tagging and sharing various bits

of information. New tools are extracting and using the knowledge we've embedded into what we've created to improve searching, browsing, and decision-making, substantially improving on software that didn't previously use such a collective intelligence.

In this article, I introduce the exciting paradigm of citizen sensing enabled by mobile sensors and human computing—that is, humans as citizens on the ubiquitous Web, acting as sensors and sharing their observations and views using mobile devices and Web 2.0 services.

### Citizen-Sensor Networking

By contributing so much online content, many people have become "citizens" of an Internet- or Web-enabled social community; the use of Internet- or Web-enabled mobile devices to upload this data gives these devices the ability to act as sensors. Thus, the term *citizen-sensor network* refers to an interconnected network of people who actively observe, report, collect, analyze, and disseminate information via text, audio, or video messages.

This combination of human-in-the-loop sensing, Web 2.0, and mobile computing has led to the emergence of several citizen-sensor networks. In particular, Web 2.0 fostered the open environment and applications for tagging, blogging, wikis, and social networking sites that have made information consumption, production, and sharing so incredibly easy. However, two significant developments in mobile computing helped enable citizen-sensor networks as we know them today: enhanced features such as GPS capability and cameras became a standard part of most mobile devices, and large companies created open mobile operating systems, such as Apple's OS X for the iPhone and Google's Android.

Microblogging — in which users share short

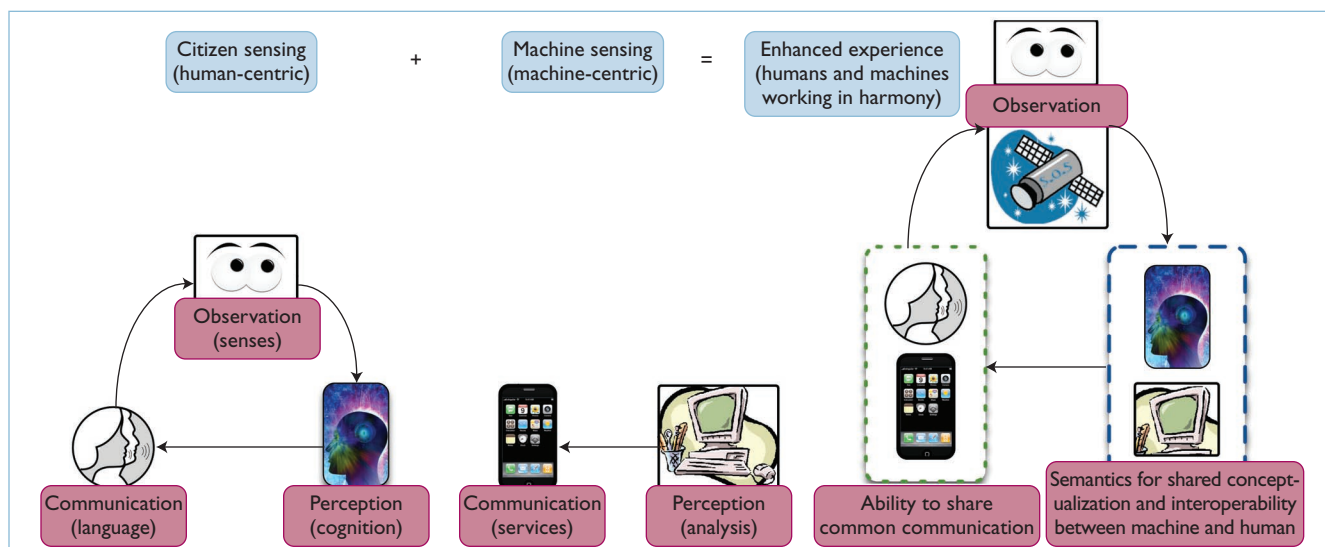


Figure 1. Model for sensing. The integration of machine-sensing with citizen-sensing provides for an enhanced experience and situational awareness that's more complete than either form of sensing could provide alone.

messages and pictures, typically over the Web – is of particular interest to citizen-sensors. This relatively new medium emerged on the Web in 2006 and achieved widespread adoption extremely quickly. Twitter, the most popular microblogging application, has nearly 6 million members who post almost 2 million messages per day (<http://twitterfacts.blogspot.com/2007/06/twitter-number-of-tweets-per-day.html>). Applications such as Twitterific and tweetie enable microblogging on mobile platforms, in which users can directly post photos and other digital captures of the events they observe onto the Web or social networking sites from their mobile devices.

Such applications have virtually eliminated the barriers of entry to participation and seem to have actively encouraged the emergence of citizen journalism and science. Other examples of citizen journalism include Wikinews, a growing number of sites and services such as CNN's IReport, Demotix, and Merinews. More recently, organizations such as the Boston police department have embraced citizen-sensors to assist in crime prevention ([www.cityofboston.gov/police/cristop.asp](http://www.cityofboston.gov/police/cristop.asp)). Several citizen science projects involve partici-

pants with mobile devices capturing observations and reports for environmental data collection, bird and animal counts, and more. One of the most visible uses of citizen-sensors occurred during the Mumbai terrorist attacks last November, when tweets (Twitter updates) and Flickr feeds by citizens armed with mobile phones reported observations of events in real time, often well before traditional media reports could do so ([www.informationweek.com/blog/main/archives/2008/11/twitter\\_in\\_cont.html](http://www.informationweek.com/blog/main/archives/2008/11/twitter_in_cont.html)).

The interesting twist that citizen-sensor networks bring to reporting a news story or scientific discovery is that they can record and report an event from multiple angles and perspectives. The messages that citizen-sensors send or upload come with a host of additional information, such as the spatiotemporal metadata provided in the devices used to capture them ([www.cnet.com.au/tag/camera-data-iphone-location.htm](http://www.cnet.com.au/tag/camera-data-iphone-location.htm)). Generally, an event has a time, location, and multiple thematic elements, which in turn become the basis of its semantic description. A collection of spatially, temporally, and thematically/conceptually (STT) related events define a situation; situational awareness, which represents “perception

of the environmental elements within a volume of time and space, and the comprehension of their meaning” ([http://en.wikipedia.org/wiki/Situational\\_awareness](http://en.wikipedia.org/wiki/Situational_awareness)) then leads to insight and actionable information.

The human-in-the-loop aspect of citizen sensing offers several advantages to traditional (machine) sensing. Machines are good at symbolic processing but poor at perception, which is the act of converting sensory information into symbols or words that are meaningful to humans. Placing humans in the sensing loop greatly alleviates this deficiency: sensors or devices can perform continuous, long-term sensing, but humans are much better at contextualizing and discriminating (deciding what's interesting or important) data, filtering (reporting on things of interest and importance) it across multiple modalities, and capturing the resulting observations for future symbolic processing by machines or collectively with other humans. Humans are also better at using sensing and perception to adapt to subsequent activities, which in turn affect what they observe and report. What gives humans this distinct advantage is their ability to deal with semantics and leverage extensive

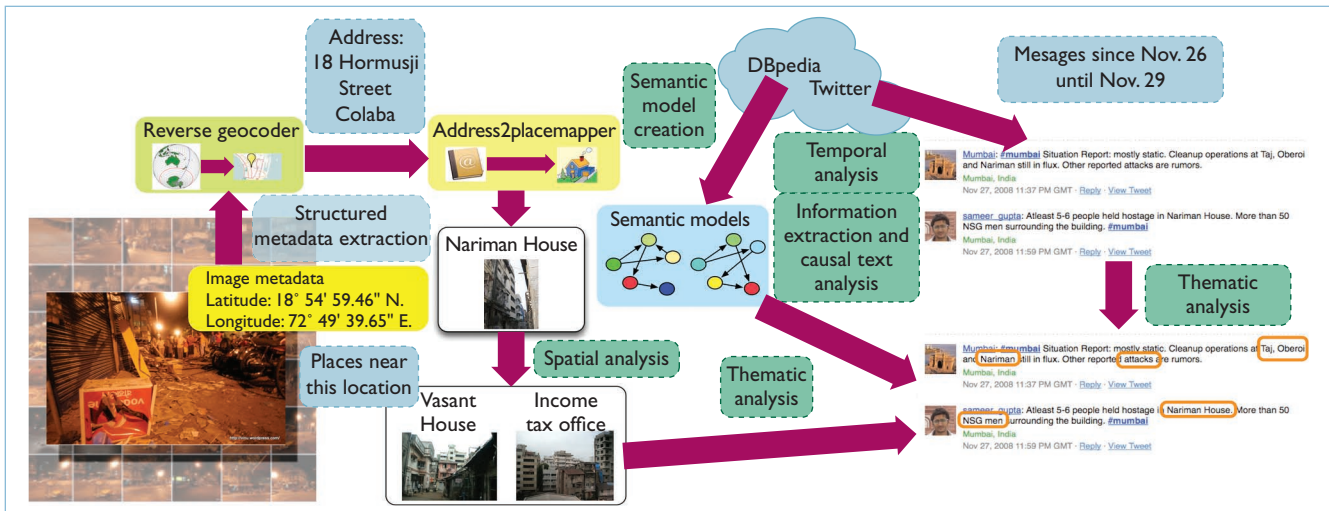


Figure 2. Citizen-sensor data. Semantic annotation integrates raw information from citizen-sensors and leads to situational awareness.

background knowledge, experience, common sense, and complex reasoning, even with fuzzy data or inconsistent information. Although traditional sensors merely report encoded observations, humans process observations via their intellect and available contextual knowledge.

A first step in a systematic approach to situational awareness is to model sensing as a cycle of operations involving observation, perception, and communication. Figure 1 shows both citizen- and machine-sensing in this general framework. Within the perception and communication operations, citizen- and machine-sensors can share information that might provide enhanced situational awareness that neither sensing system could offer alone. Two recent advances are noteworthy in this context: the ability to treat sensors as services on the Web (via standards such as Sensor Web Enablement) and the emergence of mobile sensing with humans in the loop (because humans are much better at reacting to observations).

Moreover, researchers have made several computational advances in terms of the Semantic Web and its derivatives<sup>1</sup> and in the corresponding ability to develop domain models (ontologies) and knowledge bases, se-

mantically annotate all types of data (specifically, to extract STT metadata), and computationally exploit data along these three dimensions.<sup>2</sup> As Semantic Web proponents know, annotation is the key to making data more meaningful, both for human consumption and for machine computation. Semantically annotated sensor data is more easily integrated, interpreted, and combined with databases, knowledge bases, and advanced computing capabilities. Although I've discussed semantic annotation of (machine) sensor data as part of this column before,<sup>1</sup> let's shift the focus here to semantic annotation of messages submitted in citizen-sensors. Both of these capabilities share characteristics with the semantic annotation of casual text, such as that used in social networking content.<sup>3</sup>

### Semantic Annotation of Citizen-Sensor Data

The high level of citizen participation in disseminating information during last year's terrorist attacks in Mumbai, India, demonstrated the growing power of citizen journalism. Using Flickr and Twitter, ordinary people such as Vinu Ranganathan shared their views of the events as they unfolded

(<http://www3.flickr.com/photos/vinu/sets/72157610144709049>). Although user contributions played an invaluable role in disseminating news, we can realize significant additional value through their integration with semantic analysis, which leads to situational awareness (<http://knoesis.wright.edu/library/resource.php?id=00702>).

The example depicted in Figure 2 shows metadata gathered from Twitter updates and Flickr images posted during the Mumbai attacks. Metadata can be used to extract spatial information about a resource (such as geo-coordinates for where a picture was taken or from where a message was posted) to determine the closest street address. From the image information in Figure 2, for example, we can identify the closest street address as 5, Hormusji Street, Colaba, Mumbai. When given to an "address to location" service, this information yields prominent locations near this address, including the Nariman House, Vasant Vihar, and the Income Tax Office. Next, by using temporal information from the image, we can get Twitter messages posted around the time it was taken; spatial information helps restrict the geography to just where these messages originated. The location information in

conjunction with semantic models that describe a particular domain of interest (terrorism, in this context) let us connect tweets that describe the event to images found in Flickr. Such integration provides a richer description of the event and lets us create trails of various events.

The bursty and high-throughput nature of citizen-sensor data, the thematic differences between messages, and the text's unmediated and casual nature pose several interesting research challenges, such as determining the trustworthiness of information sources ([http://news.yahoo.com/s/ap/20090511/ap\\_on\\_re\\_eu/eu\\_ireland\\_wikipedia\\_hoaxer](http://news.yahoo.com/s/ap/20090511/ap_on_re_eu/eu_ireland_wikipedia_hoaxer)), creating semantic models for general-purpose domains, and integrating application-specific semantic metadata across information sources.

### Thematic Analysis and Casual Text

The problem of semantically integrating citizen-sensor data is nontrivial. On one hand, the social context surrounding the production of such data offers exciting opportunities, but on the other, this same social context introduces challenges in terms of the content's informal nature. Off-topic discussions are common, making it difficult to automatically identify context. Moreover, the content is often fragmented, doesn't always follow grammar rules, and relies heavily on domain- or demographic-specific slang, abbreviations, and entity variations (using *skik3* for SideKick 3, for example). Content from microblogging sites is rather terse by nature, so all these factors combined make the process of automatically identifying what a message is actually about that much harder.

We can define the semantic metadata extracted from citizen-sensor content as *thematic information* – that which tells us more about the topic or theme underlying the content. In addition to the

metadata encoded in citizen-sensor messages, we can extract semantic metadata from the messages themselves. In light of various reported events, integrating potentially multimodal data from different citizen-sensor sources using spatiotemporal and thematic information can significantly enhance situational understanding and awareness, which in turn plays a vital role in our response to such events.

Semantic annotation of content refers to the process of making data more meaningful through labels (via marking up, tagging, or annotating) that conform to an agreed-upon reference model, be it a common nomenclature, dictionary, taxonomy, folksonomy, or ontology that models a specific domain. Annotations with these vocabularies make Web-based documents and data understandable to machines as well as easier to integrate and analyze. When applications use ontology rules, whether they range from simple to complex or are explicitly stated or inferred from the ontology's class properties and relationships, such applications can realize powerful reasoning over annotated data.

User-generated content (UGC) and other observations from citizen-sensor networks have unique characteristics that set them apart from the traditional content found in news or scientific articles. Coupled with the issues associated with social media content mentioned earlier (such as textual informality), the task of annotation becomes even more challenging when entities named with English language-words (Stephen King's novel *It*, Madonna's album *Music*, or Why, Arizona, one of the state's smaller cities) must be identified within informal text. This is an important challenge that Web 3.0 applications will consistently face – the process of automatically creating accurate markups or annotations from UGC to common referenced models.

The key to semantically annotating content is the process of identifying and disambiguating named entities. In short, semantic annotation transforms unstructured data into a structured representation that lets applications search, analyze, and aggregate information. When looking for information about General Motors, for example, semantically annotated content can return analyses on all its variations, such as GM, GenlMotors, and so on. Clearly, the roles of ontologies and knowledge bases in creating markups will be even more important than they were before the social Web's explosive growth – not only can they act as common reference models, but they'll also play a crucial role in inferring semantics behind UGC while supplementing well-known statistical and natural language processing (NLP) techniques. Consider this tweet from the Mumbai terror attacks: “mumbai taj 4th floor left wing fire, live on desitv.” Although natural language understanding is hard in itself, the noncapitalization of key entities such as Mumbai and the Taj Hotel makes for inaccurate natural language parse structures (compare Figures 3a and 3b, generated using the Berkeley Natural Language Parser at <http://nlp.cs.berkeley.edu/>). In such scenarios, knowing from a domain model that the Taj Hotel is a landmark in the city of Mumbai can offer meaningful support to the statistical strength of a corpus's entities.

Additional metadata that situational-awareness applications can exploit is the availability of spatial information, typically obtained from the device generating the content. An important area of investigation for such imminent applications will be how to effectively supplement existing statistical and NLP-based content-analysis frameworks with available domain knowledge and the spatial and social context surrounding the generated content.



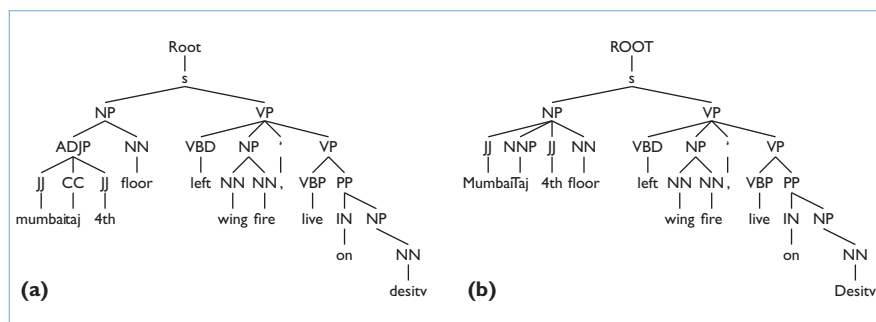


Figure 3. Natural language parser. The structure for casual text can lead to inaccuracies without the help of a natural language parser: (a) “mumbai taj 4th floor left wing fire, live on desitv” should read (b) “Mumbai Taj 4th floor left wing fire, live on Desitv.”

### Creating Semantic Models

Integrating data on the basis of thematic information is a harder problem owing to the rich vocabulary people use when describing a particular situation. As a simple example, think of how to relate two pieces of content that talk about an explosion and a blast in the same spatiotemporal setting.

We can integrate data based on thematic information by using a variety of statistical NLP and knowledge-intensive techniques. An intriguing possibility is to create and exploit semantic domain models to supplement traditional techniques, but creating a semantic model to describe thematic information for general-purpose domains, such as disaster management, is a challenging problem in itself. The most important aspect in the creation of domain models is the agreement required to define the domain. Although domain experts can come to an agreement in specialized domains (as in biomedical and healthcare domains, for example), the same isn't true for the Web. Broader and less specific areas require fewer agreements, and a clique- or committee-driven approach won't help reach mass consensus for the larger areas. For dynamic real-world events that rely on very narrow contexts, high-level concept models might be less useful. Instead, we might need to rely on more community-driven

sources of information to generate domain models. One such class of recent efforts used Wikipedia as a source for extracting an ontology because it's a community created and maintained source that reflects a degree of agreement.<sup>4,5</sup>

Wikipedia's all-encompassing scope isn't only a strength – it can be a weakness when we're only interested in a specific domain. Recent work<sup>6</sup> assists the user in carving small and focused domain descriptions out of Wikipedia based on a seed-category, article, or query. The resulting domain model contains concepts of immediate interest to the task at hand. To use it in classification tasks or as a starting point for more formal ontology development, users can export it to OWL, RDF-S or XML. Representation of domain knowledge is another important consideration. Unlike narrower, more constrained domains such as business and science, in which formal domain modeling can enable powerful reasoning for search, aggregation, and integration purposes, lightweight knowledge representation is both adequate and desirable when applications don't need to exploit all of a domain model's features.

Using domain models to semantically annotate unstructured data is a well-known research area. But to annotate citizen-sensor observations, we must effectively complement

spatial, temporal, and thematic data processing with available domain knowledge. Consider the example in Figure 3, which mentions the Taj Mahal Palace in Mumbai but refers to it only as the Taj. A domain model of landmarks in Mumbai along with other contextual information about hotels that have wings and DesiTV being an Indian TV channel supports the annotation of Taj with the concept Taj Mahal Palace in the domain model. Disambiguating such casual mentions by referencing a common domain model facilitates citizen-sensor data aggregation.

### Empowering Situation-Aware Applications

In an ongoing effort, Kno.e.sis researchers have built a system called Twitris (<http://twitris.dooduh.com>) to gather real-time citizen-sensor observations from Twitter that support STT analytics. The goal is to preserve social signals and present event indicators that lead to situational awareness. Let's review some illustrative examples of how Twitris can be used to identify social signals by analyzing tweets from around the world.

**What's new and interesting?** Consider a scenario in which two event descriptors “mumbai attacks” and “hawala funding” appear in citizen-sensor observations – specifically, the term “mumbai attacks” has occurred every day in the past week whereas “hawala funding” is a new descriptor for today. In most circumstances, users are more likely to be interested in perspectives and experiences that differ from yesterday. Looking at spatial contexts, we also find that “hawala funding” doesn't appear in any other country on the same day, whereas “mumbai attacks” occurs in almost all of them, which implies that “hawala funding” is unique and a stronger descriptor local to the US; “mumbai attacks” is

a weaker descriptor in terms of its uniqueness to this local region.

Combining the STT components of event descriptors during their analysis can offer several opportunities for presenting and using specific observations. Besides being able to find information about known facts and discover new ones, an STT analysis also allows situational-awareness applications to effectively preserve local and global social signals pertaining to any real-world event. Consider this year's G20 financial summit – by using appropriate spatial conditions, we can quickly assess what's being said about it in Asia versus in North America. Thus, the meaning and importance of entities found in citizen-sensor observations not only depends on their distribution in a corpus of related observations, but also on how they're discussed in other spatial and temporal settings. Such analysis into isolating spatial and temporal social signals lets us ask a range of questions.

**What's a region paying attention to today? What are people most excited or concerned about?** For any particular event, an STT slice of citizen-sensor observations will readily tell us today's prevalent descriptors or entities. Figure 4, for example, culls keywords and phrases out of citizen-sensor observations pertaining to the Mumbai terrorist attacks from different parts of the world. Twitris has weighted the words by their distribution within the country, their local versus global importance as relevant to the event, and the descriptor's recent popularity. Such summaries are most helpful for aiding situational-awareness applications, rather than simply viewing a list of observations themselves.

**How is an entity's perception changing over time in any region?** Allowing temporal aspects into analysis also lets us observe how the percep-

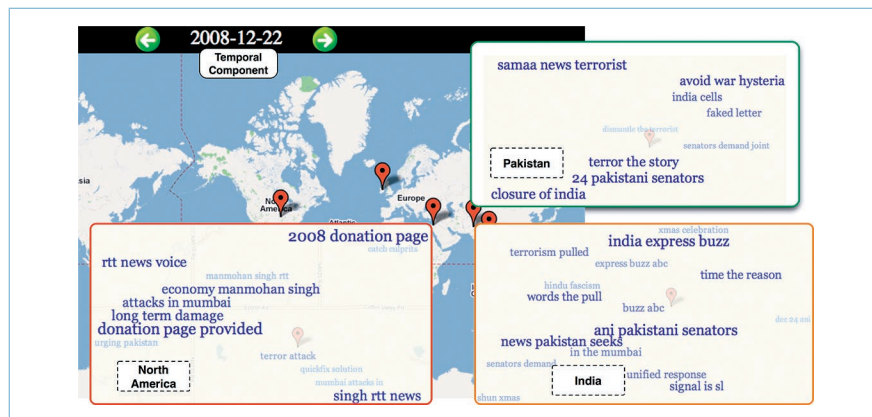


Figure 4. Twitris. Example shows how Twitris preserves social signals from the Mumbai terrorist attacks around the world via spatio-temporal-thematic analysis of citizen-sensor observations. (Application available at <http://twitris.dooduh.com>.)

tion of an event or entity changes over time. This ability is critical in understanding how an event progressed, the key descriptors involved at important timelines, which perceptions originated in what specific regions, and so on.

Looking beyond today's primitive yet compelling capabilities for understanding and analyzing the data reported by citizen-sensors, the future holds a much bigger promise for addressing more challenging problems and improving the human experience. Specifically, it will involve using semantics and social computing to exploit what tens of billions of machine sensors and more than 3 billion citizen-sensors produce on a regular basis. Before long, most of our work will focus on computing for the human experience – perhaps we'll even witness the Earth turn into a supercomputer during our lifetimes. □

**Acknowledgments**

This article incorporates research partially funded by the US National Science Foundation (IIS award #071441) and performed by Kno.e.sis members (aggregation and integration of social data by Karthik Gomadam, analysis of user-generated content by Meena Nagarajan, extraction of a domain model

from Wikipedia by Christopher Thomas, and Semantic Web sensors by Cory Henson).

**References**

1. A. Sheth, C. Henson, and S. Sahoo, "Semantic Sensor Web," *IEEE Internet Computing*, vol. 12, no. 4, 2008, pp. 78–83.
2. A. Sheth and M. Perry, "Traveling the Semantic Web through Space, Time, and Theme," *IEEE Internet Computing*, vol. 12, no. 2, 2008, pp. 81–86.
3. A. Sheth and M. Nagarajan, "Semantics-Empowered Social Computing," *IEEE Internet Computing*, vol. 13, no. 1, 2009, pp. 76–80.
4. F. Suchanek, G. Kasneci, and G. Weikum, "Yago: A Core of Semantic Knowledge," *Proc. 16th World Wide Web Conf. (WWW 2007)*, ACM Press, 2007, pp. 697–706.
5. F. Wu and D. Weld, "Automatically Refining the Wikipedia Infobox Ontology," *Proc. 17th World Wide Web Conf. (WWW 08)*, ACM Press, 2008, pp. 635–644.
6. C. Thomas et al., "Growing Fields of Interest: Using an Expand and Reduce Strategy for Domain Model Extraction," *Proc. 2008 IEEE/WIC/ACM Int'l Conf. Web Intelligence and Intelligent Agent Technology*, Sydney, ACM Press, 2008, pp. 496–502.

Amit Sheth is an IEEE fellow, LexisNexis Ohio Eminent Scholar, and director of the Kno.e.sis Center at Wright State University. Contact him via <http://knoesis.wright.edu/amt>.