

Computing for Human Experience

Semantics-Empowered Sensors, Services, and Social Computing on the Ubiquitous Web

Amit Sheth • *Wright State University*

In his influential paper "The Computer for the 21st Century," Mark Weiser talked about making machines fit the human environment instead of forcing humans to enter the machine's environment.[1] He noted, "The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it." Weiser's vision, outlined two decades ago, led to ubiquitous computing. Now, we must again rethink the relationship and interactions between humans and machines — this time, including a variety of technologies, including computing technologies; communication, social-interaction, and Web technologies; and embedded, fixed, or mobile sensors and devices.

We're on the verge of an era in which the human experience can be enriched in ways we couldn't have imagined two decades ago. Rather than depending on a single technology, we've progressed with several whose semantics-empowered convergence and integration will enable us to capture, understand, and reapply human knowledge and intellect. Such capabilities will consequently elevate our technological ability to deal with the abstractions, concepts, and actions that characterize human experiences. This will herald *computing for human experience* (CHE).

The CHE vision is built on a suite of technologies that serves, assists, and cooperates with humans to nondestructively and unobtrusively complement and enrich normal activities, with minimal explicit concern or effort on the humans' part. CHE will anticipate when to gather and apply relevant knowledge and intelligence. It will enable human experiences that are intertwined with the physical, conceptual, and experiential worlds (emotions, sentiments, and so on), rather than immerse humans in cyber worlds for a specific task. Instead of focusing on humans interacting with a technology or system, CHE will feature technology-rich human surroundings that often initiate interactions. Interaction will be more sophisticated and seamless compared to today's precursors such as automotive accident-avoidance systems.

Many components of and ideas associated with the CHE vision have been around for a while. Here, I discuss some of the most important tipping points that I believe will make CHE a reality within a decade.

Bridging the Physical/Digital Divide

We've already seen significant progress in technology that enhances human-computer interactions; the iPhone is a good example. Now we're seeing increasingly intelligent interfaces, as exemplified by Tom Gruber's Intelligence at the Interface technology (http://tomgruber.org/news/sdforum-dec13.htm), which has demonstrated contextual use of knowledge to develop intelligent human–mobile-device interfaces. We're also seeing progress in how machines (devices and sensors), surroundings, and humans interact, enabled by advances in sensing the body, the mind, and place. Such research supports the ability to understand human actions, including human gestures and languages in increasingly varied forms. The broadening ability to give any physical object an identity in the cyber world (that is, to associate the object with its representation), as contemplated with the Internet of Things, will let machines leverage extensive knowledge about the object to complement what humans process.

Human-machine interactions are taking place at a new level, as demonstrated by Psyleron's Mind Lamp (www.psyleron.com/lamp.aspx), which shows connections between the mind and the physical world. Soon, computers will be

 1089-7801/10/$26.00 © 2010 IEEE

able to translate gestures to concrete actionable cues and understand perceptions behind human observations, as shown by MIT's Sixth Sense project (www.pranavmistry.com/projects/sixthsense).

In addition, interactions initiated in the cyber world are increasing and becoming richer. Examples range from a location-aware system telling a smart phone user about a sale item's availability nearby to advanced processing of sensor data and crowd intelligence to recommend a road rerouting or to act on behalf of a human. This bridging of the physical/digital divide is a key part of CHE.

Elevating Abstractions That Machines Understand

Perception is a key aspect of human intelligence and experience. Elevating machine perception to a level closer to that of human perception will be a key enabler of CHE. In 1968, Richard Gregory described perception as a hypothesis over observation.[2] Such hypothesis building comes naturally to people as an (almost) entirely subconscious activity. Humans often interpret the raw sensory observation before recognizing a conscious thought. On the other hand, hypothesis building is often cumbersome for machines. Nevertheless, to integrate human and machine perception, the convergence must occur at this abstraction level, often termed *situation awareness*. Therefore, regardless of the source, this integration requires a shared framework for communicating and comparing situation awareness.

Perceptual hypotheses represent the semantics, or meaning, of observation. Beginning with raw observation, we find such meaning by leveraging background knowledge of the interaction between observation and possible causes to determine the most likely hypothesis. Previous experience, schooling, and personality account for much of the background knowledge in a single mind. Machines also must leverage background knowledge for effective perception.[3] So, effective perception requires a framework for representing background knowledge.

From Perception to Semantics

John Locke, Charles Peirce, Bertrand Russell, and many others have extensively and wonderfully written about semiotics — how we construct and understand meaning through symbols. A key enhancement we're already seeing is the humanization of data and observation, including social computing extending semantic computing and vice versa. Metadata is no longer confined to structural, syntactic, and semantic metadata but includes units of observations that convey human experience, including perceptions, sentiments, opinions, and intentions.

Soon, we'll be able to convert massive amounts of raw data and observations into symbolic representations. We'll make these representations more meaningful through a variety of relationships and associations we can establish with other things we know. We'll then be able to contextually leverage all this to improve human activities and experience.

CHE will bring together many current technological advances in capabilities that are easy and natural for humans but harder for machines, fundamentally combining human sensing with machine sensing and processing. In CHE,

- pattern recognition,
- image analysis,
- casual text processing,
- sentiment and intent detection,
- using domain models to gather factual information, and
- polling social media to gather community opinions and build intelligence

will all come together to enable a system that makes conclusions and decisions with human-like intuition, but much more quickly than humans.

Semantics at an Extraordinary Scale

Semantic computing, aided by Semantic Web technology, is an ideal candidate framework for meaningful representation and sharing of hypotheses and background knowledge. Together with semantic computing, the large-scale adoption of Web 2.0 or social-Web technology has led to the availability of multimodal user-generated content, whether text, audio, video, or simply attention metadata, from a variety of online networks. The most promising aspect of this data is that it truly represents a population and isn't a biased response or arbitrary sample study. This means that machines now have at their disposal the variety and vastness of data and the local and global contexts that we use in our day-to-day processing of information to gather insights or make decisions.

We also see a move from document- and keyword-centric information processing that relies on search-and-sift to representing information at higher abstraction levels. This involves moving from entity- or object-centric processing to relationship- and event-centric processing. This, in turn, involves improving the ability to extract, represent, and reason about a vast variety of relationships, as well as providing integral support for information's temporal, spatial, and thematic elements.

With parallel advances in knowledge engineering, large-scale data analytics, and language understanding, we're able to build systems that can process, represent, and reason over data points much as humans do. In addition, we can provide extremely rich markups of all the observations available to a machine,



Figure 1. A contemporary architecture for semantic computing. It analyzes any form of Web, social, or sensor data by extracting metadata, resulting in comprehensive semantic annotation. This process is aided by conceptual models and knowledge and by a variety of information-retrieval, statistical, and AI (machine learning and natural-language processing) techniques, at the Web scale. Semantic analysis supported by mining, inferencing, and reasoning over annotations supports applications.

letting machines connect the dots (contexts) surrounding the observations (data) and draw conclusions that nearly mimic human perception and cognition. All these together are reducing the disparity between humans' perceptions and the conclusions that machines draw from quantified or qualified observations.

Semantics-empowered social computing, semantics-empowered services computing (smart mashups), and semantics-enhanced sensor computing (exemplified by the semantic sensor Web) are key building blocks of CHE.

The sidebar "Influential and Interesting Works That Lead to Computing for Human Experience" (available at www.computer.org/cms/Computer.org/dl/mags/ic/2010/01/extras/mic2010010088s.pdf) discusses such important research as the memex and trailblazing, ambient intelligence, the Semantic Web, experiential computing, the Relationship Web, humanist computing, and the PeopleWeb.

Semantic Computing as a Starting Point

At the center of the approach to achieving CHE is semantic computing. Figure 1 shows a contemporary architecture for semantic computing (simplified for brevity). It has four key components: data or resources, models and knowledge, semantic annotation, and semantic analysis or reasoning.

Whereas semantic computing started with primarily enterprise and then Web data, including business and scientific data and literature, it has expanded to include any type of data and massive amounts of Web-accessible resources, including services, sensor, and social data. Here are some impressive examples:

- the capture of comprehensive personal information (for example, MyLifebits; http://research.microsoft.com/en-us/projects/mylifebits),
- the collection of massive amounts of interlinked curated data (for example, Linked Data; http://linkeddata.org),
- data and information contributed by a community of volunteers (for example, Wikipedia) and the record of social discourse of millions of users on a vast number of topics (for example, Facebook and Twitter), and
- the collection of observations from sensors in, on, or around humans; and around the earth.

Semantic computing over such a bewildering variety of data is made possible largely by an agreement on what the data means. This agreement can be represented in a manner that's formal or informal; explicit or implicit; or static (through a deliberate, expert-driven process), periodic, or dynamic (for example, mining Wikipedia to extract a targeted taxonomy). The key forms in which such agreements are modeled include formal ontologies, folksonomies, taxonomies, vocabularies, and dictionaries.

Researchers have started making rapid strides in creating models and background knowledge from human collaborations (as exemplified by hundreds of expert-created ontologies). They've done this by

- selectively extracting or mining the Web for facts, as demonstrated by Voquette/Semangix,[4] or scientific literature and
- harvesting community-created content, as demonstrated by Yago (www.mpi-inf.mpg.de/yago-naga/yago) and Taxonom (http://taxonom.com).

Along with domain-specific or thematic conceptual models, temporal and spatial models (ontologies) have taken their rightful place for capturing meaning, especially as we seek to go from keywords and documents to objects, and then to relationships and events.

Such models and associated back-

ground knowledge have provided powerful ways to automatically extract semantic metadata or semantically annotate any type of data to associate meaning with the data. A variety of semantic computations aided by pattern extraction, inferencing, logic- and rule-based reasoning, and so on, then provide a range of applications including semantic search, browsing, integration, and analysis. Such applications can lead to insights, decision support, and situational awareness. One such application is the extraction of social signals as exemplified on Twitter (http://twitris.knoesis.org).



Figure 2. An example embodying some of the computing for human experience *(CHE) promises — helping a farmer in his natural context. The farmer sends a message requesting help with an unknown crop disease. The CHE system analyzes the message, contacts appropriate sources, and returns actionable information, while requiring minimal involvement or technology consciousness from the farmer.*

An Illustrative Example

To get an idea of some of the capabilities I've described, consider a scenario in which a farmer observes an unfamiliar disease on his crop and seeks information to manage it (see Figure 2). He clicks a picture, tags it with keywords "crop" and "blight," and sends a message seeking more information: "Looks bad, but I don't know what it is. Any help would be great."

A CHE system would analyze the image to help identify the exact crop (for example, sweet corn) and the disease. It would also use location coordinates and related contextual and background domain knowledge, including local weather and soil conditions (for example, to determine whether the disease is northern corn leaf blight). It would analyze the farmer's message to extract the information-seeking intent and detect an unfavorable sentiment associated with the content.

It would then broadcast this information to online forums and social networks of individuals whose profiles indicate a professional or scientific interest in farming, crop diseases, and plant pathology. It would track responses, prioritizing those from authoritative sources, pulling actionable information based on their suggestions, aggregating duplicate suggestions, filtering spam, and presenting summaries of crowd-contributed intelligence. In this case, the actionable information could be a list of suitable fungicides and pesticides and their prices, buying options, and action plans (for example, if the farmer has a sweet corn variety, he could spray fungicide and then use hybrid seed in the future to provide blight resistance). The CHE system would also have a feedback mechanism, prompting the farmer for progress and informing the community when metrics deviate from known specifications.

The system would do all these things while requiring minimal involvement from the farmer.

This example doesn't capture all the promise of advances in machines interacting with humans at higher abstraction levels. Perhaps the CHE system, using machine perception, could also detect the disease-ridden crops before the farmer even notices and initiate the search I described, providing actionable intelligence to the farmer. More than just provide automation, CHE would work around the farmer's natural work pattern and environment, making technology interactions minimal and natural.

Acknowledgments

I acknowledge the contributions of Kno.e.sis researchers, especially Karthik Gomadam, Cory Henson, and Meena Nagarajan, and funding from the US National Science Foundation, the US National Institutes of Health, the US Air Force Research Laboratory, IBM, Hewlett-Packard, and Microsoft.

References

1. M. Weiser, "The Computer for the 21st Century," *ACM SIGMOBILE Mobile Computing and Communications Rev.*, vol. 3, no. 3, 1999, pp. 3–11.
2. R. Gregory, "Perceptual Illusions and Brain Models," *Proc. Royal Soc. London B*, vol. 171, 1968, pp. 279–296.
3. R. Bajcsy, "Active Perception," *Proc. IEEE*, vol. 76, no. 8, 1988, pp. 996–1005.
4. A. Sheth et al., "Semantic Content Management for Enterprises and the Web," *IEEE Internet Computing*, July–Aug. 2002, pp. 80–87.

Amit Sheth is director of the Wright State University's Kno.e.sis Center and is the university's LexisNexis Ohio Eminent Scholar. He's on the Web at http://knoesis.org/amit.

Semantics & Services

Influential and Interesting Works That Lead to Computing for Human Experience

Amit Sheth, Wright State University

V. Bush, "As We May Think," *The Atlantic*, July 1945; www.theatlantic.com/doc/194507/bush. A true visionary, Vannevar Bush presented interesting ideas on the memex (a theoretical computer system) and trailblazing. This is one of the early and important writings emphasizing the role of relationships, which has inspired much of my research. Cartic Ramakrishnan and I further explored an aspect of this vision in "Relationship Web: Blazing Semantic Trails between Web Resources," *IEEE Internet Computing*, vol. 11, no. 4, 2007, pp. 84–88.

E. Zelkha and B. Epstein, "From Devices to 'Ambient Intelligence: The Transformation of Consumer Electronics,'" talk given at the 1998 Digital Living Room Conference (PowerPoint presentation available at www.epstein.org/brian/ambient_intelligence/DLR%20Final%20Internal.ppt). This vision of digital environments that are sensitive and responsive to the presence of people goes beyond ubiquitous computing and advanced human-computer interaction. In this vision, devices work together to help people easily and naturally perform everyday activities. To do this, the devices employ information and intelligence hidden in the network connecting the devices, using context-aware, anticipatory, and adaptive technologies. Computing for Human Experience (CHE) also seeks to enhance this vision beyond its device roots. CHE leverages additional progress in perceptual, semantic, and social capabilities, emphasizing experiential attributes. It spans the scope from an individual and his or her surroundings to the collective social consciousness.

T. Berners-Lee and M. Fischetti, *Weaving the Web*, Harper, 1999. Chapter 12 of the book introduced the Semantic Web, which uses metadata to extensively associate meaning with data. A 2000 keynote I gave built on the data/information/knowledge-centric work performed in the 1990s (see http://knoesis.wright.edu/library/resource.php?id=735). In that keynote, I discussed uses of ontologies and semantic metadata extraction and annotation that led to commercial applications in semantic categorization, cataloging, search, browsing, personalization, and targeting. Berners-Lee (along with James Hendler and Ora Lassila) followed up *Weaving the Web* with a more detailed intelligent-agent, AI-centric vision in the highly cited May 2001 *Scientific American* article "The Semantic Web." The Semantic Web initiative carried out in the W3C's context has led to the development of standards (for example, RDF and OWL) resulting in Web-scale data integration and sharing, and knowledge aggregation (for example, using ontologies) and application. CHE envisages extending these unprecedented semantic data and computing capabilities to sensory, social, and experiential aspects.

Scientists have recognized semantics' importance in dealing with data heterogeneity since the 1980s, when conceptual or semantic models emerged. The 1990s through the early 2000s saw the use of conceptual models or ontologies for semantic metadata extraction and annotations, for faceted or semantic search, and subsequently for semantic analytics. Some early examples of semantics in computing appear in J. Heflin and J. Hendler, "A Portrait of the Semantic Web in Action," *IEEE Intelligent Systems*, vol. 16, no. 2, 2001, pp. 54–59, and A. Sheth and C. Ramakrishnan, "Semantic (Web) Technology in

Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion.
Inclusion does not necessarily constitute endorsement by the IEEE or the IEEE Computer Society.

Action: Ontology Driven Information Systems for Search, Integration and Analysis," *IEEE Data Engineering Bull.*, vol. 26, no. 4, 2003, pp. 40–48. An example of the extension of semantic computing to today's data explosion appears in C. Bizer, T. Heath, and T. Berners-Lee, "Linked Data—the Story So Far," *Int'l J. Semantic Web & Information Systems*, vol. 5, no. 2, 2009, pp. 1–22.

R. Kurzweil, *The Age of Spiritual Machines: When Computers Exceed Human Intelligence*, Penguin, 2000. That the pace of technological innovation is accelerating is believable. However, the prospect of computers exceeding human intelligence, as Ray Kurzweil postulates, is (whether likely or not) certainly less exciting to us humans, in my opinion.

R. Jain, "Experiential Computing," *Comm. ACM*, vol. 46, no. 7, 2003, pp. 48–55. This article focuses on incorporating insight and experiences into computing, supported by users applying their senses to all forms of data related to events. Also relevant are Ramesh Jain's views on EventWeb; see "EventWeb: Developing a Human-Centered Computing System," *Computer*, vol. 41, no. 2, 2008, pp. 42–50. CHE is partly a continuation of the experiential-computing vision. It builds on advances in information and communication technology, sensors, and embedded computing, incorporating social and semantic computing and emphasizing higher-level abstractions. Jain and I coedit the Springer book series Semantic and Beyond: Computing for Human Experience, in which we first conceived the term "computing for human experience."

J. Gemmell, G. Bell, and R. Lueder, "MyLifeBits: A Personal Database for Everything," *Comm. ACM*, vol. 49, no. 1, 2006, pp. 88–95. MyLifeBits presages an ability to record virtually everything in a person's life. Such an ability to capture personal and related social information is complemented by the increasing ability to observe or sense a vast variety of physical as well as social phenomena and events through the emerging sensor Web and social Web. CHE will understand, integrate, analyze, and utilize all these to enhance human experience.

J. Rossiter, "Humanist Computing: Modelling with Words, Concepts, and Behaviours," *Modelling with Words*, LNCS 2873, Springer, 2003, pp. 124–152. This paper is relevant to CHE's focus on raising the level of abstraction. It proposes modeling with words, concepts, and behaviors to define a hierarchy of methods that extends from low-level data-driven modeling with words to the high-level fusion of knowledge in the context of human behaviors.

R. Ramakrishnan and A. Tomkins, "Toward a PeopleWeb," *Computer*, vol. 40, no. 8, 2007, pp. 63–72. This article and numerous other works (including Wikipedia's Collective Intelligence entry and my article "Citizen Sensing, Social Signals, and Enriching Human Experience," *IEEE Internet Computing*, vol. 13, no. 4, 2009, pp. 87–92) discuss how connecting billions of people, generating content based on massive numbers of user contributions, citizen or participatory sensing, and using the Web to capture social consciousness are changing how we collect, share, and analyze human knowledge and experience. To realize CHE, we'll need to combine physical, personal, social, collective, and environmental scopes.

This Web extra accompanies the article, "Computing for Human Experience: Semantics-Empowered Sensors, Services, and Social Computing on the Ubiquitous Web, by Amit Sheth, IEEE Internet Computing, *vol. 14, no. 1, 2010, pp. 88–91; http://doi.ieeecomputersociety.org/10.1109/MIC.2010.4.*

cn *Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.*