Wright State University

CORE Scholar

Kno.e.sis Publications

The Ohio Center of Excellence in Knowledge-Enabled Computing (Kno.e.sis)

3-2016

Semantic, Cognitive, and Perceptual Computing: Paradigms That Shape Human Experience

Amit P. Sheth

Wright State University - Main Campus, amit@sc.edu

Pramod Anantharam

Wright State University - Main Campus, anantharam.2@wright.edu

Cory Henson

Wright State University - Main Campus

Follow this and additional works at: <https://corescholar.libraries.wright.edu/knoesis>

Part of the [Bioinformatics Commons](#), [Communication Technology and New Media Commons](#), [Databases and Information Systems Commons](#), [OS and Networks Commons](#), and the [Science and Technology Studies Commons](#)

Repository Citation

Sheth, A. P., Anantharam, P., & Henson, C. (2016). Semantic, Cognitive, and Perceptual Computing: Paradigms That Shape Human Experience. *Computer*, 49 (3), 64-72.
<https://corescholar.libraries.wright.edu/knoesis/1089>

This Article is brought to you for free and open access by the The Ohio Center of Excellence in Knowledge-Enabled Computing (Kno.e.sis) at CORE Scholar. It has been accepted for inclusion in Kno.e.sis Publications by an authorized administrator of CORE Scholar. For more information, please contact library-corescholar@wright.edu.

Keywords: *semantic computing, cognitive computing, perceptual computing, human-centric computing, computing for human experience, experiential computing, Internet of Things, IoT, Internet of Everything, IoE, Web evolution, Web 3.0, semantic Web, big data*

Research Feature

Semantic, Cognitive, and Perceptual Computing: Paradigms That Shape Human Experience

Amit Sheth and Pramod Anantharam, Kno.e.sis—Wright State University

Cory Henson, Bosch Research and Technology Center

Unlike machine-centric computing, in which efficient data processing takes precedence over contextual tailoring, human-centric computation provides a personalized data interpretation that most users find highly relevant to their needs. The authors show how semantic, cognitive, and perceptual computing paradigms work together to produce actionable information.

While Bill Gates, Stephen Hawking, Elon Musk, Peter Thiel, and others engaged in OpenAI discuss whether or not AI, robots, and machines will replace humans, proponents of human-centric computing continue to extend work in which individuals and machines partner to enhance human experience and better inform decision making. Figure 1 shows how these respective views fall into a spectrum from machine-centric work—including intelligent machines¹ and ubiquitous computing²—to more human-centered views such as the augmentation of human intellect³ and man-machine symbiosis.⁴

Figure 1. Views along the spectrum of machine-centric to human-centric computing. At the far end is our work on Computing for Human Experience⁵, which explores paradigms such as semantic, cognitive, and perceptual computing.

Among efforts to explore human-centric paradigms are computing for human experience (CHE)⁵ and experiential computing,⁶ both of which aim to empower individuals while keeping them in the loop. In this way, machines make humans more productive by supporting highly informed and timely decision making and by generally enriching people's quality of life. CHE encompasses the elements of a human-centric computing paradigm, while experiential computing uses the symbiotic relationship between computers and people and exploits their complementary strengths, including machines' ability to manipulate symbols and perform logical reasoning and humans' ability to recognize complex patterns. These efforts are inspired not by virtual machine efficiency but by human perception and cognition.

For nearly three decades, the Web has been an important vehicle for human-centric computing, as it carries massive amounts of multimodal and multisensory observations about situations pertinent to people's needs, interests, and idiosyncrasies. The Internet of Things (IoT) is poised to substantially expand this information infrastructure. Maturing paradigms such as semantic computing (SC) and cognitive computing (CC), and the emerging perceptual computing (PC) paradigm provide a continuum through which to exploit the ever-increasing amount and diversity of data that could enhance people's daily lives. SC and CC sift through and personalize raw data according to context and individual users, creating abstractions that move the data closer to what humans can readily understand and apply in decision making. PC, which interacts with the surrounding environment to collect data that is relevant and useful in understanding the outside world, is characterized by interpretative and exploratory activities. The sidebar "[Underlying Concepts](#)" explains the role of semantics, cognition, and perception in more detail.

To better understand how these paradigms complement one another within the continuum, we reviewed the literature and explored the paradigms' use in an asthma-management system (<http://wiki.knoesis.org/index.php/Asthma>), which involves converting volumes of heterogeneous multimodal data (big data) into actionable information (smart data). This application is an excellent context for examining the computational contributions of SC, CC, and PC in synthesizing actionable information.

[Sidebar](#)

Underlying Concepts

The semantic, cognitive, and perceptual computing paradigms are based on the broader concepts of semantics, cognition, and perception, which interact seamlessly. Semantics attaches meaning to observations by providing a definition within a system context or the knowledge that people possess. Semantics allows an individual to process observations by relating them to other observations and data. Cognition enables the individual to understand his or her environment and paves the way for the application of perception to explore and deepen that understanding. Perception enables individuals to focus on the most promising course of action by incorporating background knowledge that provides a comprehensive contextual understanding.

These three concepts, which are fundamental to human-centric computing, have multiple definitions and applications, but our concern is how they define computing paradigms inspired by human cognition and

perception: how they complement one another as they mimic the human brain and how observational data relates to them. We thus ignore their use in other contexts, such as in programming languages (semantics) or in how people interact with computing peripherals (perception).

Semantics

Semantics is the meaning attributed to concepts and their relationships within the mind. The network of concepts and relationships is used to represent knowledge about the world, which in turn enables the cognition and perception needed to interpret daily experiences. Semantic concepts represent, unify, subsume, or map to various data patterns and thus provide a conceptual abstraction that hides the data's syntactic and representational differences. Generally, this involves mapping observations from various physical stimuli, such as visual or speech signals, to concepts and relationships as humans would interpret and communicate them. For example, an observer might recognize a person by face (visual signal) or voice (speech signal). Together, these signals represent a single semantic concept, such as "Mary," which is meaningful to the observer.

Perception and Cognition

Perception is an active, cyclical process of exploration and interpretation,^{1,2} with the two phases constantly exchanging places. Individuals might start with interpretation of data from their senses and background knowledge and then switch to exploration to collect more data, which must then be interpreted for relevance to a particular context. Conversely, they might start with exploration of a particular idea and then interpret the data they have collected in terms of a real-world event.

In this process, perception involves recognizing and classifying patterns from sensory inputs that physical stimuli have generated and using those patterns to recognize facts and form feelings, beliefs, and ideas. Collected data helps individuals use cognition—the act of combining data from perception with existing knowledge to understand their environment. The interpretation of observations leads to abstractions (some concept in perceptual or background knowledge), and exploration leads to actuation to seek the most relevant next observation (to disambiguate candidate abstractions).

During perception, the individual applies expectations or predictions based on cognition while receiving sensory inputs and then constantly tries to match input with domain and background knowledge and reasoning (cognition).³ The process becomes increasingly contextual and personalized.

In this sense, perception interleaves the processing that occurs in the bottom and top parts of the brain. According to cognitive-model theory, the bottom brain organizes the received signals from sense organs resulting in an individual's perception of the real world, while the top brain deals with planning, goal setting, and even dynamically changing goals and outcomes.⁴ Interpretation is analogous to the bottom-brain operation of processing observations from the senses; exploration is akin to the top-brain processing of making and adapting plans to solve problems.⁵ This type of interaction—often involving focused attention and physical actuation—enables the perceiver to collect and retain relevant data from the vast ocean of all possible data, and thus facilitates a more efficient, personalized interpretation or abstraction.

References

1. U. Neisser, *Cognitive Psychology*, Prentice-Hall, 1967.
2. R. Gregory, "Knowledge in Perception and illusion," *Philosophical Trans. Royal Soc. London, Series B, Biological Sciences*, vol. 352, no. 1358, 1997, pp. 1121–1127.
3. D.S. Modha et al., "Cognitive Computing," *Comm. ACM*, vol. 54, no. 8, pp. 2011, pp. 62–71.
4. S. Kosslyn and G.W. Miller, *Top Brain, Bottom Brain: Harnessing the Power of the Four Cognitive Modes*, Simon and Schuster, 2015.
5. S. Michael and S. Kosslyn, *Top Brain, Bottom Brain: Surprising Insights into How You Think*, Simon and Schuster, 2013.

[End sidebar](#)

The Vision of Personalized Computing

To enable the development of technologies that “disappear into the background,”² human-centric computing research is tackling how to endow the Web with sophisticated, human-like capabilities to reduce information overload. The vision is to produce computers that can process and analyze data in a highly contextual and personalized manner at a scale much larger than the human brain can accommodate. Human-centric computing systems will be highly proactive in that they will sense, measure, monitor, predict, and react to the user’s physical, cyber, and social environment and thus more intimately support any decision or action.⁷

The physical sphere encompasses reality as measured by sensors and devices that make up the IoT, the cyber sphere encompasses all shared data and knowledge on the Web (Wikipedia and linked open data, for example), and the social sphere encompasses the many forms of human interaction. Observation data might represent events of interest to a population (such as a global warming), a subpopulation (such as city traffic), or an individual (such as an asthma attack). These observations help shape human experience, which we define as the materialization of feelings, beliefs, facts, and ideas that can be acted upon.⁸

CHE uses the Web to manage and share massive amounts of multimodal and multisensory observations that capture moments in people’s lives, whether those situations are pertinent to the individual’s immediate needs and interests or simply reflections of his or her idiosyncrasies. From the collected data emerges a contextual and personalized data interpretation that people can more readily grasp than raw data. For example, a person with severe asthma would rather receive a recommendation to take preventive medication than data about conditions that make an asthma attack more likely, such as high pollen level, extreme humidity, and poor air quality.

From Raw Data to Actionable Information

Figure 2 shows how SC, CC, and PC work together toward the ultimate goal of human-centric computing: exploiting volumes of increasingly diverse data to provide information about things that a particular individual in a specific context would need.

Figure 2. Conceptual distinctions among semantic computing (SC), cognitive computing (CC), and perceptual computing (PC). SC gives meaning to raw data that is needed to interpret and integrate observations from the physical, cyber, and social (PCS) spheres. CC interprets observations with SC annotations, and PC seeks observations from the environment to collect relevant data that will enhance an understanding of the outside world. PC is cyclical, interpreting and exploring observations to tailor background knowledge to the context and individual of interest.

Semantic computing's role

SC encompasses the technology required to represent concepts and their relationships in an integrated semantic network that loosely mimics the brain's conceptual interrelationships. This knowledge, represented formally in an ontology, can be used to annotate data and infer new knowledge from interpreted data. SC supports horizontal operators that semantically integrate multimodal and multisensory observations from diverse sources.⁷ To support semantic integration, these operators often explicitly model the domain of interest as an ontology or a knowledge graph.

SC's 15-year history⁹ includes annotation standards for social and sensor data that are still in use.¹⁰ Figure 2 shows SC as a long vertical rectangle that PC's interpretation and exploration pass through. SC also provides a formal representation of background knowledge.

Cognitive computing's role

When the Defense Advanced Research Projects Agency (DARPA) launched a CC project in 2002, it defined CC as the ability to "reason, use represented knowledge, learn from experience, accumulate knowledge, explain itself, accept direction, be aware of its own behavior and capabilities, [and] respond in a robust manner to surprises."¹¹

Our definition of CC is consistent with DARPA's, as the diagram in Figure 2 shows. CC interprets annotated observations obtained from SC, or raw observations from diverse

sources, and presents the results to humans. Humans, in turn, can use the interpretation to perform an action such as taking medication to prevent an asthma attack, which forms additional input for CC. Using background knowledge, CC works with PC to interpret and understand observations.

Cognitive algorithms, a major focus of current CC research, interpret data by learning and matching patterns in a way that approximates human cognition. CC systems, another research area, use machine learning and other AI techniques without explicit programming. CC systems learn from their experiences and improve when performing repeated tasks. A CC system acts as a prosthetic for human cognition by analyzing a massive amount of data and answering questions humans might have when making decisions. An example is IBM's Watson supercomputer, which won the *Jeopardy!* game show against human contestants in early 2011. Watson's approach (although not its technology) has since been extended to aid medical doctors in clinical decisions.

Socrates taught that knowledge is attained through the careful and deliberate process of asking and answering questions. With data mining, pattern recognition, and natural-language processing, CC is rapidly progressing toward developing technology that supports and extends people's ability to answer complex questions. By implementing vertical operators to rank answers to an explicit question, or hypotheses, from unstructured data, a CC system facilitates cyclical interaction between question asker and machine.⁷ Vertical operators use background knowledge to create abstractions from raw data, making that data easier for humans to interpret. The CC system can then constantly learn and refine the generated hypotheses in the context of those abstractions. Interaction is explicit and defines the symbiotic relationship between person and machine.

Perceptual computing's role

PC supports the ability to ask contextually relevant and personalized questions,¹² and complements SC and CC by providing the machinery to ask the next question or derive a hypothesis. PC conducts iterative cycles of exploration and interpretation on the basis of observations relevant to possible hypotheses. Additional facts and observations assist in evaluating or narrowing the list of candidate hypotheses—all of which aids decision makers in gaining actionable insights.

In asthma management, for example, PC determines which data is most relevant, disambiguating multiple possible causes of an asthma condition. Through focused use of sensing and actuation technologies, PC seeks data from the physical, cyber, and social domains that will add to its store of observations and domain knowledge. Thus, a PC monitoring device might synthesize various fine-grained patient observations, such as coughing, reduced activity, and sleep disturbances, into a larger symptom set, such as a particular asthma control level, and then invoke iteration to seek more information. As this example illustrates, PC's final action is to match the symptom set (data about the patient) with a semantic abstraction (control level).

Although PC efforts to date have investigated data interpretation, they have yet to adequately address the relationship between data interpretation and environmental exploration or interaction. PC implements horizontal operators to integrate heterogeneous and multimodal observations and vertical operators, transforming massive amounts of multimodal and multisensory observations into SC abstractions that are intelligible to people.⁷ In addition to SC and CC technologies, PC uses machine perception¹³ together with available background knowledge to explore and interpret observations. Exploration and interpretation are implicit when background knowledge is available but can become explicit when incorporating inputs from people.

Asthma Management Application

Asthma is a multifaceted, highly contextual, and personal disease. It is multifaceted because it is characterized by many aspects, including environmental triggers and the patient's sensitivity to those triggers. It is highly contextual because events of interest, such as the patient's location and the triggers at that location, are crucial for timely alerts. It is personal because patients have varying responses to triggers, and their actions are based on the severity of their condition.

Patients are diagnosed in terms of two levels. *Severity* level indicates the degree to which asthma affects the patient, which can be mild, mild persistent, moderate, or severe. *Control* level indicates the extent to which asthma exacerbations are managed, which can be well controlled, moderately controlled, or poorly controlled. Severity level seldom changes, but control level can vary drastically, depending on triggers, environmental conditions, medication, and symptomatic variations.

Consider Anna, a 10-year-old diagnosed with severe asthma and well-controlled exacerbations because she takes her medication consistently and avoids exposure to triggers. However, suppose Anna receives an invitation to play soccer in a few days. She and her parents must balance Anna's desire to participate in the soccer game with the need to maintain well-controlled exacerbations.

The solution to this balancing problem is not straightforward and cannot be found using only facts on the Web or in electronic medical records (EMRs). The knowledge available on the Web is not specific to Anna's context; websites describe general symptoms, triggers, and tips for better asthma management, but Anna's symptomatic variations for environmental triggers could be unique. Although the EMRs' medical domain knowledge of asthma might contain symptomatic variations for triggers, Anna's parents will not find any guidance specific to their daughter's case. For example, symptomatic variations do not account for Anna's unique environmental and physiological dynamics or her quality-of-life choices, which would factor into any recommendation about playing in the soccer game.

Figure 3 shows how SC, CC, and PC work together to balance these concerns and provide the optimal solution to asthma control in this specific context. The paradigms

basically play the same role as in Figure 2, except their application is specific to asthma management. SC enables PC to synthesize tailored abstractions like control level, which are much more intelligible to doctors who must recommend corrective actions than a list of general symptoms, such as high cough, disturbed sleep, and reduced activity.

Figure 3. Roles of SC, CC, and PC in providing actionable information for asthma management. As in Figure 2, SC provides an abstraction to make raw data meaningful; in this case, control level. CC reviews asthma literature and lists disturbed sleep, high cough, and reduced activity as control-level contributors. PC interprets those contributors in terms of a specific context, such as environment, by including conditions such as temperature, humidity, and pollen. The result is actionable information that the doctor can use to make a well-informed recommendation.

Figure 4. Operations performed over raw data by SC, CC, and PC for an asthma-management scenario. The CC system gleans statistics from patient observations and asthma literature, and the PC system iteratively interprets and explores them to learn personalized and context-specific normalcy for the patient.

Figure 4 shows operations that demonstrate how SC, CC, and PC can synthesize actionable information from raw data. Even annotated data is not sufficiently valuable to Anna’s doctor in making a recommendation to Anna’s parents. Although the CC system can provide the number of nights of disturbed sleep, days of coughing, and days of reduced activity per week from its asthma literature analysis,¹⁴ these statistics only give rise to other questions: What does reduced activity mean in terms of steps taken per day? What does disturbed sleep mean in terms of duration of rapid-eye-movement (REM) sleep per night?

The PC system enables the personalization of information provided by the CC system, which can answer these two questions because the PC system can learn personalized and context-specific normalcy for Anna through its iterative cycle of interpretation and exploration.

SC’s role

The semantic network of general medical domain knowledge related to asthma and its symptoms define asthma control levels in terms of symptoms. This general knowledge can be integrated with knowledge of Anna’s specific case as documented in her EMR.

SC makes raw data more meaningful by annotating data with semantic concepts defined in an ontology, thus enhancing data consumption as well as reasoning about and sharing data. The Semantic Sensor Network (SSN) ontology^{10,15} defines concepts

and relationships for modeling sensors and their observations. Figure 4 shows three observation types: sleep quality, number of steps, and number of coughs per day. Raw data points (1 hour and 17 minutes of sleep, 672 steps, and 20 coughs) do not carry much meaning, but linking them through annotation to concepts defined in an ontology produces much richer data (REM sleep, steps, and cough incidents) that is amenable to knowledge-aware interpretation. Abstractions such as annotation are extremely useful in knowledge-rich domains like medicine. Annotation represents the complex patterns in sensor data in a form that is easier for decision makers, such as doctors, to understand and act on in a timely manner—abilities that are critical in practical medicine.

Annotations span multiple modalities, general domain knowledge, and context-specific knowledge (Anna's asthma severity and control level). Asthma triggers such as weather, pollen count, and air-quality index can be researched through Web services. However, without background knowledge, interpretation can be challenging, and Anna's parents do not have such asthma-related knowledge. They are left with no particular insights at this stage, as manually interpreting all observations is not practical.

CC's role

With no background knowledge of asthma management, Anna's parents contact her pediatrician, Dr. Jones. In our scenario, Dr. Jones has access to a CC system similar to Watson¹⁶ that is designed for asthma management. The CC system looks at asthma articles and medical journals (unstructured data) and historical data in asthma patients' EMRs, which provide information on treatment regimes, medications, risks, and patient outcomes, revealing valuable information that is otherwise hidden in massive amounts of medical literature. Thus, the CC system minimizes the time the doctor must spend reviewing research and clinical outcomes.

Dr. Jones discovers from this analysis that people with well-controlled asthma, like Anna, can engage in physical activities as long as they are using appropriate preventive medication. However, Dr. Jones is still uncertain about how vulnerable Anna's asthma would be to fluctuations in the weather and air-quality index because she lacks personalized and contextualized knowledge about Anna's day-to-day environment. She can use only her experience with other patients and their symptomatic variations.

Moreover, Dr. Jones does not have direct access to what is normal for Anna in terms of sleep, activity, and symptoms and must therefore elicit that information by querying Anna's parents. There is no way to evaluate the accuracy of the information provided by Anna's parents because it lacks supporting evidence. In short, Dr. Jones must deal with so much uncertainty that any recommendation would be based on educated guesses. Because asthma symptoms vary considerably with context and patient, Dr. Jones chooses not to make a definitive recommendation to Anna's parents.

PC's role

A PC system minimizes that uncertainty by providing a personalized and contextualized understanding of Anna's environmental and symptomatic variations. It explores Anna's EMR data to derive normalcy information related to her asthma history and then translates that into a specification of her unique symptomatic normalcy. It uses this specification to interpret current observations, categorizing them into disturbed sleep, low activity, and high-coughing incidents.

Interpretation is based not only on symptomatic normalcy but also on asthma severity level, so the corresponding recommended action is likewise conditioned on both these aspects. The PC system interprets contextual observations to provide personalized and contextualized abstractions, which Dr. Jones can use to make a more informed recommendation. The decisions and recommendations are based on evidence provided by the PC system rather than on educated guesses.

However, Dr. Jones must still contend with variations in weather and air quality. With the rise of mobile computing and IoT technologies, a PC system might be implemented as an *intelligence at the edge technology*—technology that makes sense of data on resource-constrained devices¹⁷—as opposed to a logically centralized system. Computation would be carried out on a mobile device, allowing better control of data access, sharing, and privacy. In Anna's case, the PC system would possibly run as a mobile application with inputs from multiple sensors, such as kHealth (<http://wiki.knoesis.org/index.php/Asthma>).

From its exploration and interpretation, the PC system knows that last month Anna exhibited reduced activity during a soccer practice (interpreted as an instance of asthma exacerbation). It then seeks information about how weather and air quality affect Anna's asthma symptoms, exploring generic background knowledge and adding contextual and personalized knowledge. Generic knowledge might include how poor air quality could exacerbate asthma; contextual and personalized knowledge might include how Anna's exposure to poor air quality could exacerbate Anna's asthma and the predicted environmental conditions on the day of the soccer game.

Dr. Jones will have access to this information as well as observations from the CC system, which empowers her to provide a well-informed recommendation that Anna should forego the soccer match because air quality will be poor on that day. Her advice is both personalized to Anna and contextualized to the event.

Clearly, SC, CC, and PC are both complementary and synergistic. With technological support, SC can deal with big-data challenges, CC can use relevant knowledge to improve data understanding for decision making, and PC can provide personalized and contextual abstractions over massive amounts of multimodal data from the physical, cyber, and social realms.

Using these paradigms synergistically, machines can provide answers to the complex questions posed to them as well as ask the right follow-up questions and interact with physical, cyber, and social aspects to collect relevant data. As PC evolves, personalization components will extend to include temporal and spatial contexts and other factors that drive human decisions and actions, such as emotions and cultural and social preferences. With this extension, machines will be able to provide effective answers and enable decisions and timely actions tailored to each person.

We believe PC will eventually enable fully personalized and contextualized background knowledge. Often hailed as the next Web phase, the IoT—with its emphasis on sensing and actuation—will exploit all three computing paradigms. Over the next decade, their development, both individually and in concert, along with their integration into the Web’s fabric, will likely enable the emergence of a far more intelligent human-centric Web.

References

1. J. McCarthy, “What Is Artificial Intelligence?,” 2007; www-formal.stanford.edu/jmc/whatisai.pdf.
2. M. Weiser, “The Computer for the 21st Century,” *Scientific American*, vol. 265, no. 3, 1991, pp. 94–104.
3. D.C. Engelbart, *Augmenting Human Intellect: A Conceptual Framework*, summary report, Stanford Research Institute, SRI Project 3578, 1962.
4. J.C.R. Licklider, “Man-Computer Symbiosis,” *IRE Trans. Human Factors in Electronics*, vol. 1, 1960, pp. 4–11.
5. A. Sheth, “Computing for Human Experience: Semantics-Empowered Sensors, Services, and Social Computing on the Ubiquitous Web,” *IEEE Internet Computing*, vol. 14, no. 1, 2010, pp. 88–91.
6. R. Jain, “Experiential Computing,” *Comm. ACM*, vol. 46, no. 7, 2003, pp. 48–55.
7. A. Sheth, P. Anantharam, and C. Henson, “Physical-Cyber-Social Computing: An Early 21st Century Approach,” *IEEE Intelligent Systems*, vol. 28, no. 1, 2013, pp. 78–82.
8. P. Lutus, “The Levels of Human Experience,” 2012; http://arachnoid.com/levels/resources/levels_of_human_experience.pdf.
9. A. Sheth, “Semantics Scales Up: Beyond Search in Web 3.0,” *IEEE Internet Computing*, vol. 15, no. 6, 2011, pp. 3–6.
10. M. Compton et al., “The SSN Ontology of the W3C Semantic Sensor Network Incubator Group,” *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 17, 2012, pp. 25–32.
11. R.C. Johnson, “DARPA Puts Thought into Cognitive Computing,” *EETimes*, 9 Dec. 2002; www.eetimes.com/document.asp?doc_id=1227413.
12. C. Henson, A. Sheth, and K. Thirunarayan, “Semantic Perception: Converting Sensory Observations to Abstractions,” *IEEE Internet Computing*, vol. 16, no. 2, 2012, pp. 26–34.
13. C. Henson, “A Semantics-Based Approach to Machine Perception,” PhD thesis, Wright State Univ., 2013.
14. S. Spangler et al., “Automated Hypothesis Generation Based on Mining Scientific Literature,” *Proc. 20th ACM Int’l Conf. Knowledge Discovery and Data Mining (SIGKDD 14)*, 2014, pp. 1877–1886.
15. L. Lefort et al., “Semantic Sensor Network XG Final Report,” W3C Incubator Group Report, vol. 28, 2011; www.w3.org/2005/Incubator/ssn/XGR-ssn-20110628.
16. L.F. Friedman, “IBM’s Watson May Soon Be the Best Doctor in the World,” *Business Insider*, 22 Apr. 2014; www.businessinsider.com/ibms-watson-may-soon-be-the-best-doctor-in-the-world-2014-4.

17. C. Henson, K. Thirunarayan, and A. Sheth, "An Efficient Bit Vector Approach to Semantics-Based Machine Perception in Resource-Constrained Devices," *Proc. 11th Int'l Semantic Web Conf. (ISWC 12)*, 2012, pp. 149–164.

Amit Sheth is the LexisNexis Ohio Eminent Scholar and executive director of Wright State University's Ohio Center of Excellence in Knowledge-Enabled Computing (Kno.e.sis). His research interests include semantics-empowered integration and analysis; applications of enterprise, Web, social, sensor, and Internet of Things (IoT) data; smart data; semantic sensor Web; citizen sensing; and semantic perception. Sheth received a PhD in computer and information sciences from Ohio State University. He is an IEEE Fellow. Contact him at amit@knoesis.org or <http://knoesis.org/amit>.

Pramod Anantharam is a PhD candidate in computer science at Kno.e.sis. His research interests include information extraction, applied machine learning to solve sustainability and healthcare problems, predictive analytics, information extraction, and the semantic Web. Anantharam received a bachelor's degree in electrical and electronics engineering from Visvesvaraya Technological University. He is a member of the Society for Industrial and Applied Mathematics. Contact him at pramod@knoesis.org.

Cory Henson is a senior research scientist at Bosch Research and Technology Center in Pittsburgh. His research interests include the application of semantic technologies to enable interoperability and data integration for IoT applications. Henson received a PhD in computer science and engineering from Wright State University. He is a member of the Open Geospatial Consortium. Contact him at cory.henson@us.bosch.com.