

A Robotic Web

Michael N. Huhns • University of South Carolina

Q. How is the Web like a robot?
A. Well, the Web is not yet very robotic, but it is becoming more so, and its progression along this path provides a guide to its evolution.

Let's first think about the Web in robotic terms. A common, generic, cognitive architecture for a robot's behavior is PRA: *perceive, reason, act* (see Figure 1). A robot with a PRA architecture would use its sensors to perceive its environment – which is typically physical but could include online information – reason about its perceptions by applying whatever knowledge it's received or learned, and then decide to take an action. This PRA sequence then iterates indefinitely. The Web is on its way to encompassing this cognitive architecture, and the result will be a far more useful Web.

Web Evolution

The initial Web was little more than a database of documents that users could query. It was essentially passive, in that it could neither perceive nor reason nor act (see Figure 2). Human users perceived the Web, thought about it, and acted to update it.

Spiders, crawlers, and search histories, along with the indexes they produce, gave the Web a reasoning capability: it could infer pages' semantics and relative importance, as well as which pages were like others. The spiders themselves were similar to robots: they could perceive a page, reason about its content, and act to update search engine indexes. The Web as a whole could reason about its own contents, although the only action it could take was to update those contents. Still, this made the Web active.

The Web is now poised to ascend to the next step in increasing its utility as a societal assistant, acquiring more PRA capabilities

characteristic of a robot (see Figure 3). Let's look at how these capabilities might transpire in the future Web.

Perceiving

How well does the Internet perceive the world – both physical and informational – in which it resides? The Web can perceive the contents of databases and document stores very well. Its spiders crawl billions of websites, pages, and data stores periodically, after which it makes available searchable indexes of everything it finds. But it currently does a poor job of perceiving the physical world. You can search for the information your organization has stored on its Web servers, but not information about the buildings where the Web servers are located.

The Internet of Things (IoT) is changing this. Predictions for the IoT are that a significant fraction of real-world objects will be connected to the Internet, with the ability to make their state accessible. Buildings, their contents, trees, automobiles, highways – all will be accessible, along with the interpretation of their state. Initially, a variety of sensors will connect and stream their state to the Web. Eventually, the Web will be able to query these sensors to seek the information it needs to respond to user requests. Such active sensing is a simple form of action that the Web can take, and the result will be increased or improved content.

The goal is to enable anyone, anywhere to perceive the state of the world to an appropriate level of granularity (with permission). The problem is to make sense out of billions of percepts.

Reasoning

The Internet has *a lot* of knowledge, information, and data. It iterates in that spiders/crawlers continually locate new or changed pages, and users conduct searches, both of which help improve

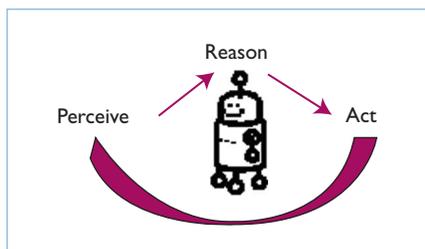


Figure 1. A common architecture for robot behavior. A robot with this PRA architecture would use its sensors to perceive its environment, reason about its perceptions, and then decide on an action.

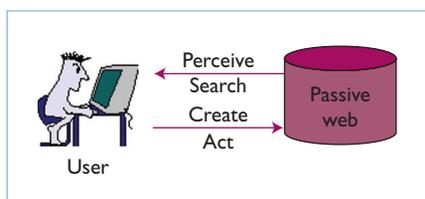


Figure 2. The passive Web. In this initial incarnation, human users perceived the Web (typically by searching), thought about it, and acted to update it.

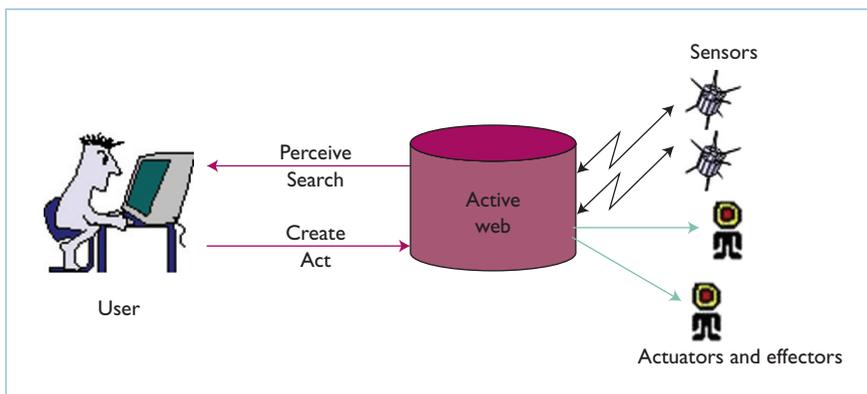


Figure 3. Active Web. Using sensors and effectors, the Internet of Things is changing the Web, enabling it to sense its physical environment.

the indexes that form the basis for search engines. The Internet also infers the semantics of webpages by using or guessing at keywords, and then decides which pages are like other pages based on those keywords and user patterns. The current Web, however, does a poor job of relating

webpage contents to each other: search engines retrieve a set of pages that match a query's terms, but they don't aggregate the contents. For example, a million pages might have values for a country's population, but the Web can't compute the average of those values. Aggregation

IEEE Internet Computing

Editor in Chief

Michael Rabinovich • michael.rabinovich@case.edu

Associate Editors in Chief

M. Brian Blake • m.brian.blake@miami.edu
 Barry Leiba • barryleiba@computer.org
 Maarten van Steen • steen@cs.vu.nl

Editorial Board

Virgilio Almeida • virgilio@dcc.ufmg.br
 Elisa Bertino • bertino@cerias.purdue.edu
 Fabian Bustamante • fabianb@cs.northwestern.edu
 Yih-Farn Robin Chen • chen@research.att.com
 Vinton G. Cerf • vint@google.com
 Siobhán Clarke • siobhan.clarke@cs.tcd.ie
 Fred Douglass • f.douglass@computer.org
 Schahram Dustdar • dustdar@dsg.tuwien.ac.at
 Stephen Farrell • stephen.farrell@cs.tcd.ie
 Robert E. Filman* • filman@computer.org
 Carole Goble • cag@cs.man.ac.uk
 Michael N. Huhns • huhns@sc.edu
 Anne-Marie Kermarrec • anne-marie.kermarrec@inria.fr
 Anirban Mahanti • anirban.mahanti@nicta.com.au
 Cecilia Mascolo • cecilia.mascolo@cl.cam.ac.uk
 Pankaj Mehra • pankaj.mehra@ieee.org
 Peter Mika • pmika@yahoo-inc.com
 Dejan Milojčić • dejan@hpl.hp.com
 George Pallis • gpallis@cs.ucy.ac.cy
 Charles J. Petrie* • petrie@stanford.edu
 Gustavo Rossi • gustavo@lifia.info.unlp.edu.ar

Amit Sheth • amit.sheth@wright.edu
 Munindar P. Singh* • singh@ncsu.edu
 Oliver Spatscheck • oliver@spatscheck.com
 Torsten Suel • suel@poly.edu
 Craig W. Thompson • cwt@uark.edu
 Doug Tygar • tygar@cs.berkeley.edu
 Steve Vinoski • vinoski@ieee.org
 * EIC emeritus

CS Magazine Operations Committee

Paolo Montuschi (chair), Erik R. Altman, Nigel Davies, Lars Heide, Simon Liu, Cecilia Metra, Shari Lawrence Pfleger, Michael Rabinovich, Forrest Shull, John R. Smith, Gabriel Taubin, George K. Thiruvathukal, Ron Vetter, and Daniel Zeng

CS Publications Board

Thomas M. Conte (chair), Alain April, David Bader, Angela R. Burgess, Greg Byrd, Koen DeBosschere, Frank E. Ferrante, Paolo Montuschi, Linda I. Shafer, and Per Stenström

Staff

Editorial Management: Rebecca Deuel-Gallegos
 Lead Editor: Linda World, lworld@computer.org
 Publications Coordinator: internet@computer.org
 Contributors: Keri Schreiner and Joan Taylor

Director, Products & Services: Evan Butterfield
 Senior Manager, Editorial Services: Robin Baldwin
 Manager, New Media & Production: Steve Woods
 Senior Business Development Manager: Sandy Brown
 Membership Development Manager: Cecelia Huffman
 Senior Advertising Supervisor: Marian Anderson, manderson@computer.org

Technical cosponsor:



IEEE Internet Computing
 IEEE Computer Society Publications Office
 10662 Los Vaqueros Circle
 Los Alamitos, CA 90720 USA

Editorial. Unless otherwise stated, bylined articles, as well as product and service descriptions, reflect the author's or firm's opinion. Inclusion in *IEEE Internet Computing* does not necessarily constitute endorsement by IEEE or the IEEE Computer Society. All submissions are subject to editing for style, clarity, and length.

Submissions. For detailed instructions, see the author guidelines (www.computer.org/internet/author.htm) or log onto *IEEE Internet Computing's* author center at ScholarOne (<https://mc.manuscriptcentral.com/cs-ieee>). Articles are peer reviewed for technical merit.

Letters to the Editors. Email lead editor Linda World, lworld@computer.org

On the Web. www.computer.org/internet/
Subscribe. Visit www.computer.org/subscribe/.

Subscription Change of Address. Send requests to address.change@ieee.org.

Missing or Damaged Copies. Contact help@computer.org.

To Order Article Reprints. Email internet@computer.org or fax +1 714 821 4010.

IEEE prohibits discrimination, harassment, and bullying. For more information, visit www.ieee.org/web/aboutus/whatis/policies/p9-26.html.

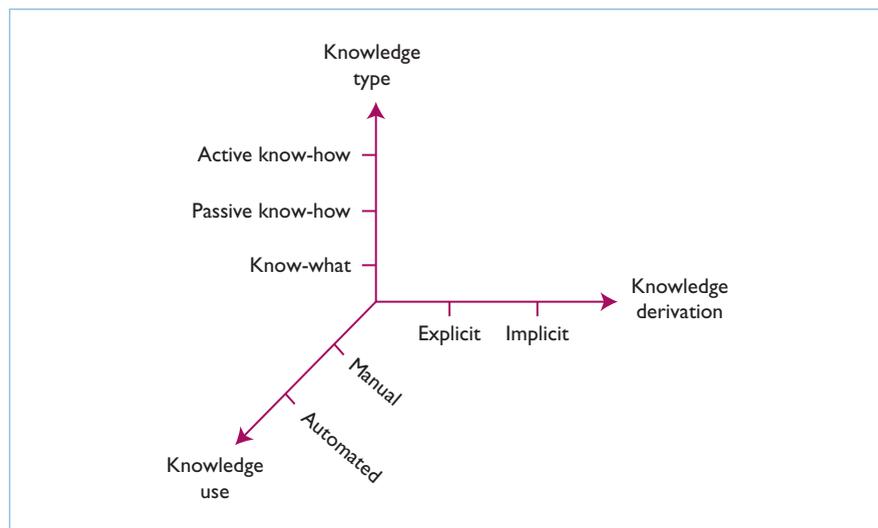


Figure 4. Three dimensions of Web knowledge. The Web must organize its knowledge into a form that lets it decide which actions are appropriate for a perceived state.

is an important inferencing mechanism that the Web will soon have, thus enhancing its reasoning power. The goal is to enable the Web to take better advantage of the information and data it has available.

Acting

The Web can take very few actions. Most of them involve improving both its contents and, to a greater extent, the metadata describing those contents. It can also engage in actions involving forms of inaction or the transfer of incorrect information. In the same vein, viruses and other malware perform a type of action, albeit in the form of corrupting information or preventing access to it during denial-of-service attacks.

The largest future change to the Web will be its connection to various actuators and effectors. The Web will be able to not only perceive temperatures, but also change them. It will be able to sense the power needs in homes, businesses, and devices, and control smart grids to move the right amount of power to the right location.

The goal is to enable anyone, anywhere to take appropriate action in the world with a sufficient level of precision (with permission). The

problem is to ensure that those actions' results are observable, controllable, and stable.

When actuators connected to a smart grid alter the distribution of electric power, will oscillations or other instabilities arise that damage the grid or the devices connected to it? Will the Web's sensors let it perceive the Web's state (its observability) with sufficient accuracy for taking the correct actions?

Managing Web Knowledge

To manage all possible actions that the Web might take will require that it organize its knowledge into a form that's amenable for deciding which actions are most appropriate for a given perceived state of the world. Perceiving, reasoning, and acting involve different kinds of knowledge, so it would be useful for Web knowledge to be partitioned in ways that facilitate these three features. We can identify three dimensions of Web knowledge: type, use, and derivation, as Figure 4 shows.

Knowledge Derivation

All large data and information networks, no matter how well they are indexed and reverse-indexed, fail

to make most of their knowledge explicit. Even if a network indexes n items, implicit knowledge exists in the relationships among the n items, and the number of relationships is potentially of order 2^n . For example, Google searches will return explicit information on the payloads of both the Chinese Long March rocket and the Indian GSLV, but no search will provide an answer as to which has greater lifting capability or how much greater one is (that is, the relationships between the two payloads). Further evolution toward a robotic Web will require progress in its ability to derive explicit knowledge.

Knowledge Types

Three kinds of knowledge are available on the Web: know-what, passive know-how, and active know-how. Consider the following example: If you search for "balance check-book," you will find pointers to sites that define a balanced checkbook, but don't describe how to balance one; sites that describe how to balance one; and sites that ask for the amounts of checks and actively perform the balancing (as a service).

Intended Use for Knowledge

The information on the Web is primarily intended for human use. The Semantic Web, by associating ontologies with webpages, provides information for machine use.

A goal for the Web as a whole is to uncover implicit, passive knowledge and make it explicit and active, for both human and machine use. A robotic Web would dramatically affect our everyday lives and activities; would grow into a "natural" extension of our capabilities, both physical and mental; and would even help guard against malware by making the Web more difficult to fool. I look forward to the Web of actions incorporating actuators, controllers,

manipulators, effectors, and robots, and in the process transforming itself as a whole into a robotic entity. □

Michael N. Huhns holds the NCR Professorship and is Chair of the Department of Computer Science and Engineering at the University of South Carolina. He also directs the Cen-

ter for Information Technology. Huhns has a PhD in electrical engineering from the University of Southern California. He's the author of nine books and more than 200 papers in machine intelligence, including the coauthored textbook *Service-Oriented Computing: Semantics, Processes, Agents*. He serves on the editorial boards for 12

journals, is a senior member of ACM, and is a fellow of IEEE. Contact him at huhns@sc.edu.

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