

# Agents in the Web

Michael N. Huhns and Munindar P. Singh  
(ic-agents@ieee.org)

November 15, 1996

**“There are agents everywhere!”** A decade ago this claim would evoke images from the Cold War (and probably get this column classified!). Today, it describes one of the hottest technical areas on the Internet.

The increasing use of intelligent software entities for advertising, finding, fusing, using, and presenting information would be no news to the readers of *Internet Computing*.

Agents are being used on the web for such tasks as information brokering, information filtering, electronic commerce, workflow management, and intelligent manufacturing. Agents are the "glue" that makes a system as large as the web manageable and viable.

And, in the months and years ahead, we expect agents will become an essential part of most web-based applications.

**“Charter.”** Our title pretty much says it. We will cover agent technology as it applies in the context of webs—both internet and intranet. We will review

- agent products and services
- research projects at universities and industrial labs
- impact of standards (e.g., KQML and CORBA) on agent techniques, (another column in *Internet Computing* will cover the rest of the Internet standards)
- trends
- our opinions on the above.

To keep our feet on the ground, we will try to feature at least one working system in each column.

**“Agents.”** So, what exactly is an agent? Must it be intelligent? Adaptive? Itinerant? Rule-based? There are almost as many opinions on this as there are agents themselves, leading to frequent debates flaring up on several Internet forums (see below).

To fan the flames, we offer our own definition: *Agents are active, persistent software components that perceive, reason, act, and communicate.*

Note that they are not required to be intelligent, which would then require its own very debatable definition. Also note that persistence—operating for extended periods without explicit guidance from their user or creator—presupposes some degree of autonomy.

Some researchers add further properties, such as being goal-directed with the ability to accept high-level tasks, being reactive, or being programmed declaratively. Other researchers limit agents to a role of representing a user or a database. We shall talk about these properties and roles in later columns, usually with discussions of systems that instantiate them to varying degrees.

**“None but the lonely.”** One property, which we feel strongly about, is *sociability*. Although there are lots of agents in the Web, they are almost universally asocial. Most agents today are typically designed to perform some task for a user—such as filter e-mail, find the prices for music CDs, and buy flowers—but they are completely unaware of each other. If one agent finds the best price for a bouquet by querying or visiting all of the Internet florist sites, there is no way for this information to be communicated to another agent engaged in the same quest.

Having social agents that interact with each other will move the Web from a pure client-server paradigm to a distributed, or better yet, *cooperative* paradigm. You will be hearing a lot on this.

**“System of the bimonth.”** An impressive system is WARREN <<http://www.cs.cmu.edu/~softagents/warren/>> developed by Dr. Katia Sycara and colleagues at CMU's Robotics Institute.[\*] WARREN is a system of intelligent agents for helping you manage your financial portfolio. It coalesces market data, financial report data, technical models, analysts' reports, and breaking news with current prices from a stock ticker.

For example, while one agent finds and plots the current price of your favorite stock, another will monitor the newswire for anything that mentions the company. Later, you might see that a sharp drop in the price occurred shortly after a news report about a brokerage that downgraded the stock.

All of the information is already available in some form on the Web—WARREN simply integrates it by having a specialized agent responsible for each resource, and then presenting it to, or alerting, a user. The agents operate on your behalf for months, whether you are logged on or not.

Check it out!

**“Resources.”** There is so much going on in agents that this column couldn't hope to cover it all. For those interested in probing deeper, there are two lively mailing lists: DAI-List and agents. To learn more about these, contact [DAI-List-request@ece.sc.edu](mailto:DAI-List-request@ece.sc.edu) and [agents@cs.umbc.edu](mailto:agents@cs.umbc.edu). There is also a Usenet newsgroup, [comp.ai](mailto:comp.ai), which sees a lot of traffic. We promised ourselves not to mention [comp.ai.philosophy](mailto:comp.ai.philosophy), so we won't.

**“Feedback.”** We welcome it. Send it to the alias above. Tell us about web sites you encountered that we should check, or those we recommended that don't work any more. Tell us about great agent applications; tell us about duds. We might not always respond to your messages, but we will always read them.

\* K. Decker and K. Sycara, "Designing Reusable Behaviors for Information Agents." In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1996.