

Doug Riecken
Guest Editor


Intelligent

Recently, the term *intelligent agent* has become quite popular. Announcements of products like the Apple Newton with its *agent software* and General Magic's *messaging agents* are evidence of significant interest in agent research and development. The critical issue in the emerging study of agents is assuring that their development is not oversold by the needs of marketing, but is instead based on a scientific approach.

All photos by Jeff Harris

Agents





his idea of agents is not new. Over the decades numerous researchers such as Marvin Minsky, Oliver Selfridge, Alan Kay, Nicholas Negroponte, Rodney Brooks, and Pattie Maes have studied problems that demonstrate some type of agent behavior. One would hope that by now we would be able to identify, to some degree, what an agent is. At best, there appears to be a rich set of emerging views. Even the terminology is a bit messy. Various researchers have invented numerous new words to better describe their agents—intelligent agents, intelligent interfaces, adaptive interfaces, knowbots, knobots, softbots, userbots, taskbots, personal agents, and network agents—just to name a few.

This past March, at the annual American Association for Artificial Intelligence (AAAI) 1994 Spring Symposium Series, there were three separate symposia each addressing some flavor of agent research: Software

Agents, Believable Agents, and Intelligent Multimedia Multimodal Systems. It is here that I would like to begin our story on agents. At both the Software Agents and Believable Agents symposia, the participants represented a rich diversity of research areas and the research focus had a direct impact on all end users. Let's examine this further:

First, the participants consisted of researchers from such fields as software engineering, robot-

ics, knowledge representation, knowledge-based systems, databases, problem solving, planning, machine learning, cognitive science, psychology, computer graphics, art, music, film, and human-computer interaction. This diversity motivates an inter-



esting question. When was the last time such a collection of researchers, at a conference, all intersected for the same topic? Based on this intersection of research interests, is it possible that agent research will eventually draw from the results of many different disciplines? I suspect the science fiction writers already know the answer to this question. In essence, the study of agents presents a unique opportunity to integrate many significant results from many diverse research areas.

Agent research also presents us with the opportunity to put our technical results directly in the hands of the end users. In the past, many AI technologies have resulted in published papers and laboratory experiments. The application of numerous efforts, such as expert systems research, has only benefited specialized groups of users. The basic idea of agent research is to develop software systems which *engage and help* all types of end users. This is clearly a unique situation. Agent research needs the integration of many different research talents, and the results serve to directly benefit everyone.

In composing this special issue on intelligent agents, I attempted to look for issues that could drive agent research, rather than select articles that discuss various prototype agents. Together, the 15 articles composing this issue provide a diverse set of views and insights.



They are intended to raise questions for future work. Therefore, if agents are a tractable research area, we might consider the following issues as addressed by the contributing authors:

- How might people think about agents?
- Will they want agents?
- How might agents *feel* about people?
- Should agents have emotions? If not, how should they measure their performance relative to human feedback?
- How will people and agents communicate with each other?
- How will people trust their agents and instruct them if they make a mistake?
- How will different agents work together and how will they communicate?
- Do agents need some common corpus of *shared knowledge* in order to communicate?
- Will agent research result in



many specialized agents—one for my email, one to manage my calendar, one to find things, one to buy tickets to a Bruce Springsteen concert, and so forth? Or, will my personal agent be like “Phil the Agent” in the Apple Computer “Knowledge Navigator” video? Phil is a very smart human-like agent.

- Should an agent be human-like or just be a computer program?
- How about common sense? How will an agent obtain enough knowledge to be useful so that people will trust it?
- What is an agent software architecture? Is it possible that agent intelligence will emerge from many diverse system components?
- How will agents learn?
- What types of surrogate tasks will agents perform?
- How might agents teach people new things?
- How can agents improve human

performance?

- Can they improve the creative performance of people as well?
- Can agent technology improve collaboration between groups of people?
- What types of agents might we find available today in industry or in emerging software products?

When we talk about intelligent agents, where do we begin? As a starting point, let us assume an



agent is a computer program. As a computer program, it is possible to measure the agent based on such properties as its correctness, completeness, efficiency, and reliability to determine if the program might be a *trusted agent*. It is from this basic computing science perspective that we might begin the study of agents as a new research area.

As Bob Hall (a colleague at Bell Labs) and I pursue our study of agents, we have found the following research question to be quite direct as a bootstrap: How can we design applications to be better surrogates while requiring less control over the environments in which these applications perform? In the end, the future of agents will come from the science of building real software programs and the empirical study of these programs with end users. To achieve this we will need to learn from many diverse examples and disciplines. **G**

