The Knowledge Engineering Review, Vol. 13:2, 1998, 137–142
Printed in the United Kingdom
Copyright © 1998, Cambridge University Press

The First Autonomous Agents Conference

W. LEWIS JOHNSON[1] and BARBARA HAYES-ROTH[2]

[1]*USC / Information Sciences Institute, 4676 Admiralty Way, Marina del Rey, CA 90292, USA Email: johnson@isi.edu*
[2]*Computer Science Department, Stanford University, Palo Alto, CA 94305, USA. Email: bhr@cs.stanford.edu*

1 Why agents?

The *First International Conference on Autonomous Agents* brought together researchers concerned with implementing systems that perceive and act in dynamic, unpredictable environments, that coordinate interoperation among complementary component capabilities, and that perform significant jobs with a high degree of autonomy.

There are several reasons why it seemed appropriate to initiate a conference on the topic of autonomous agents at this time. The expansion of the Internet and World Wide Web is one indirect cause. As the Web has grown in popularity it has also grown in complexity. It can be extremely difficult at times for users to find useful information on the Web. Consequently, there is a demand for intelligent assistants that can help users to find relevant information. At the same time, the increasing availability of information on line makes it possible for such assistants to perform a range of tasks on behalf of users.

In a similar vein, other application environments have emerged that have created demand for agent technology. Interactive games incorporate animated characters, and so there is a need for investing these characters with lifelike behaviours. Autonomous vehicles are being employed in space and military applications.

There are common concerns that cut across this diversity of application areas. Many of these concerns deal in one way or another with managing complexity. One is the ability of an agent to produce robust behaviour in a complex, unpredictable environment. Algorithms that work fine in controlled conditions can fail miserably in realistic situations. Another problem is how to manage and control the functions within a complex agent so that the agent's overall behaviour is coherent. Just as it is inappropriate for an agent single-mindedly to perform the same action regardless of what happens in the environment, it may be inappropriate for an agent to constantly switch back and forth between unrelated behaviours. Many agent developers are investigating how to take teams of agents and make them work effectively together; these can be collections of softbots or teams of robots. Other commonalties can be found between the concerns of autonomous robots operating in real-world environments, and the concerns of animated characters operating in computer-generated worlds.

Because of these common concerns, it seemed appropriate to try and form a conference that would bring the various threads of agent research together and encourage interchange. Consequently, discussions began in 1994 concerning the creation of such a conference. We decided early on that the conference should be both broadly based and heavily oriented toward implemented systems. The broad focus was important to encourage the kinds of interactions between application areas that we sought. The orientation toward implemented systems was chosen in order to focus attention on the complex problems that arise when building real agents. Although we were not against theoretical papers *per se*, we wanted to make sure that the theoretical results were applicable,

and were not confined to abstract, idealized problems. In any event, there were already many conferences where theoretical aspects of agents may be addressed, such as the Distributed Artificial Intelligence (DAI) workshops, Agent Theories, Architectures, and Languages (ATAL), and Modeling Autonomous Agents in a Multi-Agent World (MAAMAW).

2 Technical papers

To reinforce the emphasis on complete agents, the call for papers was organized around four Application Areas: the relatively mature areas of Expert Assistants and Autonomous Robots and the relatively new areas of Software Agents and Synthetic Characters. Fifty-nine technical papers were included in the program, selected from a total of 168 submissions. Thirty-nine of these were from North America, fifteen were from the European Union, and five were from the Asia-Pacific region. A total of nine countries were represented. Thirty-four groups prepared posters for presentation at the conference. The poster session had an even broader international participation, representing fifteen countries on five continents. Remarkably, the accepted papers were nearly evenly divided among the four Application Areas, as summarized briefly below.

Expert Assistants are software agents that assist humans in complex decision making or other knowledge processing tasks. These agents are often interfaced with other conventional software (e.g., database or legacy systems) as well as traditional control systems. Typical applications of such agents include medical monitoring, industrial control, business process management, computer-integrated manufacturing, and air-traffic management. These types of agents bring in related work from learning, planning, resource-bounded reasoning, and knowledge representation within a practical framework. For the first time, the Expert Assistants track of Autonomous Agents '97 conference brought together the practical aspects and complete implementations of such systems. The collection of papers in the conference proceedings shows the "mature" state of the Expert Assistants area. Anand Rao of the Australia AI Institute chaired the Expert Assistants Area.

Autonomous robots differ from other agents by having very close coupling with real environments, using sensors and actuators for interaction with the world. Some robots may work in inhospitable physical environments, interacting primarily with other robots and physical objects. Others may work among and assist people in everyday environments, such as offices. Hence a true robot (in contrast with a "softbot") must be able to operate when situated in the "real" world. Brady has defined robots as an intelligent connection between perception and action. Hence, they are clearly intelligent agents. The need for embodiment in the world is both a challenge and (to some) a burden to be avoided at all costs. In recent years, robots have been used increasingly as testbeds for ideas and algorithms in AI, as evidenced by the proliferation of robot contests, from the office-traversing robot at the AAAI conferences to robots which play soccer and other games. The papers dealing with robotics in Autonomous Agents '97 are a mixed collection. Many deal carefully with issues of visual perception or navigation or obstacle avoidance. Others try to escape from the difficult problems by studying only in simulation. This is a stimulating collection of papers, which addresses many of the current problems associated with autonomous behaviour in the physical world. George Bekey of the University of Southern California chaired the Autonomous Robots Area.

Software Agents are programs that interact with real-world software environments, such as a PC operating system, the Internet, or the World Wide Web. Their sensors observe features of these external environments, and their effectors can both alter the state of their environment directly and communicate with other agents. While AI programs that operate in software domains have appeared intermittently over the past decades, intelligent software agents have only become a major thrust of AI research in the last few years. This development has been occasioned by the emergence of the PC, the Internet, and the observation that software agents marry two desirable properties: they operate in complex real-world domains yet they are easy to build and convenient to explore as testbeds for AI research. Autonomous Agents '97 attracted high quality papers investigating a wide range of topics from automatically learning to extract information from Web

sites to using genetic algorithms to support personalized information filtering, and much more. The papers in the proceedings represent a broad cross section of a vibrant and rapidly growing field. Oren Etzioni of the University of Washington chaired the Software Agent Area.

Synthetic Characters operate in simulated environments, such as virtual worlds, MUDS, or video games. They typically interact with users and with other characters, as well as with the simulated environment. Synthetic characters have been developed for applications in the arts and entertainment, as well as in education. They emphasize qualities such as believability and personality, and often place less emphasis on deep intelligence and expertise. While there have been occasional efforts to develop synthetic characters during the history of AI, only recently have they been the topic of organized workshops, such as the 1990 AAAI workshop on Interactive Fiction, the1994 AAAI Spring Symposium on Believable Agents, and the Lifelike Computer Characters conferences. The area often draws on expertise in the arts, such as drama and character animation, to guide the efforts within AI. Autonomous Agents '97 attracted many more serious papers of this sort than prior conferences, and many of the papers made their connections with the arts explicit. While this approach is common at SIGGRAPH, and helps to drive the field forward, it is the first time an AI conference has so openly promoted such research. We believe the papers presented in the conference proceedings are the most substantive single collection of work on synthetic characters to date. Joseph Bates of Carnegie Mellon University chaired the Synthetic Characters Area.

3 Other technical sessions

3.1 Tutorials

Four tutorials were presented at the conference. Katia Sycara of Carnegie Mellon University and Pattie Maes of MIT offered an introduction to software agents. Claude Touzet of the University of Aix-Marseille III and Sandip Sen of the University of Tulsa discussed how machine learning techniques can be incorporated into agent-based systems. Tuomas Sandholm from Washington University, St. Louis discussed techniques for creating electronic markets comprised of autonomous agents, in which agents coordinate via dynamically made contracts. Anand Rao of the Australian AI Institute and Klaus Fischer of DFKI presented an introduction to BDI (Belief-Desire-Intention) Agents and their application to business and industrial practice. Overall approximately 200 conference attendees signed up for one or more tutorials.

3.2 Live and video demonstrations

Agents '97 featured demonstrations of both robotic and software agent systems. Instead of organizing standardized competitions such as RoboCup, we held these demonstrations in a completely open fashion, allowing presenters to highlight the points of particular importance in their work.

The robot demonstrations included the following presentations. A group from DFKI in Germany demonstrated their use of the InteRRaP architecture in a loading dock application. They showed both physical robots and simulated robots moving goods through a simplified mock-up of a warehouse setting, communicating and cooperating as needed. Masahiro Fujita and Koji Kageyama of Sony's D21 Laboratory demonstrated their prototype robotic dog, which is able to respond to simple hand gesture commands. George Bekey's group at USC demonstrated a robot car that is able to employ reinforcement learning to learn to navigate through a walled environment. Robin Murphy and Dave Hershberger of the Colorado School of Mines demonstrated a mobile robot that is able to recover from sensor failures, obstructions impeding sensation, etc. Finally, Wei-Min Shen of the USC Information Sciences Institute gave a demonstration of their Yoda learning robot. R. James Firby of the University of Chicago organized the robot demonstration session.

The software demonstration session was particularly successful, with 29 groups giving live demonstrations of their system. The Information Sciences Institute at the University of Southern

California hosted the demonstrations. Unfortunately, the present authors were busy during this session demonstrating our own agents, Steve the Virtual Trainer and Erin the Bartender. Consequently, we cannot give a detailed review of all the demonstrations, but all of the demonstrations were well attended. Charles Rich of Mitsubishi Electric Research Laboratory organized the software demonstration session.

Two video presentations were given at the conference, and included in the video conference proceedings. Simone Strippgen of the University of Bielefeld showed some video footage from their simulation testbed used to study and evaluate different designs for behaviour-based agents. Ronan Boulic presented some work in Daniel Thalmann's laboratory at the Swiss Federal Institute of Technology on the use of real-time gesture recognition to permit interaction between humans and synthetic agents in virtual environments.

4 Plenary sessions

4.1 Invited talks

Danny Hillis of Walt Disney Imagineering gave the introductory presentation of the conference. In his talk on ''Hollywood Agents'', Hillis discussed some of the technical problems that autonomous agent technology must overcome in order to succeed in the entertainment arena, and talked about Disney's own efforts in this area. Disney of course has developed an extensive set of techniques over years to help make its animated characters lifelike, some of which may not be obvious to the casual viewer. Animators make careful use of body language to convey emotion and intentions in characters, and also make use of background lighting and colour to set the overall mood of a scene. The state of the art in automated animation and scene rendering is of course nowhere near the level of sophistication attained by expert Disney animators. Nevertheless, Hillis argued that there are important functions that autonomous agents can perform to improve animation productivity and provide effects that are difficult to achieve using hand drawn cell animation. For example, Disney has been exploring the use of agents to play the role of extras. The result of this work may be seen in the film *The Hunchback of Notre Dame*, in which autonomous agent technology is used to animate crowd scenes. Each extra in the scene is given a limited repertoire of behaviours, which they carry out autonomously. Another area where agents can play a role is in interactive computer-based fiction. Hillis cited an example of a computer-based story where the player assumes the role of a character in the story, and the trace of story line depends upon the actions of the player.

In a similar vein, Tom Porter of Pixar, Inc. gave a talk at the conference banquet on the topic of depicting perception, thought, and action in the movie *Toy Story*. Porter analysed in detail some of the animated sequences in the film, focusing on the use of body language and facial expression to portray the mental attitudes of the characters. He also described the process that Pixar professionals go through in creating animated characters. Porter entertained us enormously with his self-effacing account of the subordination of photo-realistic graphics and the brilliant supporting mathematics to the subjective off-the-cuff—but often correct—judgments of creative directors and marketing people in the film industry.

Yasuo Kuniyoshi gave an invited presentation on his work at MIT and at the Electrotechnical Laboratory in Japan in creating autonomous social robots, i.e., robots that are able to carry out tasks autonomously while interacting with other agents, including humans. He argued that autonomy and sociability should be addressed in tandem in order to create effective robotic systems. Dr. Kuniyoshi described two fascinating experiments in this direction. One, the Learning by Watching system, is able to watch a human instructor perform a task in the blocks world, and learns to imitate the task. It combines vision, plan recognition, plan instantiation, and motor control. The other experiment involved multiple cooperating mobile robots. One robot O1 was able to watch another robot O2 pushing an object across the floor, predict the object's trajectory, detect obstacles in the path of that trajectory, and cooperatively move the obstacles out of the way of O2.

Jan Eric Larsson of the Lund Institute of Technology gave a plenary talk on the Guardian system,

developed with Barbara Hayes-Roth. Guardian is an autonomous agent for monitoring and diagnosis of post-cardiac surgery patients in an intensive-care unit. It is built with an architecture that permits several different algorithms to cooperate in performing multiple loosely coupled tasks—data interpretation, fault detection and diagnosis, explanation, prediction, and control—under hard and soft real-time constraints. Although testing in a real clinical context must await rigorous certification, Guardian has been tested on realistic monitoring scenarios generated in real-time by a sophisticated patient simulator developed by a commercial third party. The simulator is routinely used to control a physical dummy (including both its manifestation of symptoms and its responsiveness to therapeutic interventions) in a realistic ICU environment in order to train physicians in monitoring skills and strategies. Larsson reported the results of a study comparing Guardian's performance against the performance of ten ICU physicians and nurses in this simulated environment. All but two of the medical professionals made serious, and in some cases life-threatening, errors in diagnosis and treatment of the simulated patient. By contrast, Guardian diagnosed and treated the patient correctly and in a timely fashion on all of 100 trials. Larsson suggested that Guardian's superior performance reflects two factors: (a) its consistently comprehensive consideration of alternative problems, diagnoses, and treatments; and (b) its objective willingness to consider extreme diagnoses and treatments that may cause naturally emotional human beings to hesitate. These qualities make Guardian and similar intelligent agents excellent additions to multi-person teams, such as the ICU team.

Daniel Weld of the University of Washington gave the final invited talk on the topic of softbots. His talk primarily was concerned with the use of planning technology to guide softbot behaviour, with some discussion of softbot perception and action as well. He cited as examples of tasks suited to softbots the problem of integrating information from multiple knowledge sources in order to answer a user's query. He took as an example a problem posed by his wife the week before, namely to locate on the Web reviews of all movies currently playing in the Seattle area that star Marlon Brando. A softbot performing this task successfully requires a variety of knowledge and skills: knowledge about what Web sites to search, knowledge about how to formulate queries on a given site and how to interpret the results, the ability to plan a series of Web queries in order to obtain useful results quickly, and the ability to synthesize the data obtained from the sites in order to answer the user's query. Weld described a system called OCCAM being developed in his laboratory that is able to perform these sorts of Web-based query tasks.

4.2 Panels

Two panels were featured at the conference. The first panel, chaired by Mark Wiegand of BT Laboratories, discussed efforts in various companies around the world to exploit agent technology in managing and exploiting telecommunications networks.

The second panel, which closed the conference, was chaired by Barbara Hayes-Roth, and brought together the four Area Chairs—Anand Rao (Expert Assistants), George Bekey (Autonomous Robots), Oren Etzioni (Software Agents), and Joe Bates (Synthetic Characters). Hayes-Roth invited the panelists to speak briefly, in round-robin fashion, to each of a series of questions designed to elicit the shared themes and opportunities for cross-fertilization among the four areas: What are the big questions in your area? What concepts or techniques can your area contribute to the others? What can your area learn from the other areas represented at the conference? Following a period of open discussion with the audience, Hayes-Roth invited each of the panelists to address one final question: ''When we return for the Agents '98 Conference, what new and exciting results will we hear about from your area?''

5 The future

By most measures, Agents '97 was a great success. The papers, demonstrations, and plenary sessions were all of extremely high quality and so was the audience participation. The conference benefited

greatly from the efforts of the sponsors, the presenters, the reviewers, and the organizing committee[1]. In addition, the conference succeeded in attracting attention to the potential role of autonomous agents from people outside of the field. Around 50 reporters and other press representatives showed up, and filed stories that appeared soon afterwards in mass media such as *The New York Times, Computer World, Hot Wired*, and CNN.

Planning for Agents '98 to be held in Minneapolis in May, is well underway. The 1998 conference has received approximately the same number of submissions as Agents '97. The conference program will be similar in structure to that of the 1997 conference, but will be further expanded to include a workshop track.

In the longer term, the prospects for the conference are closely tied to the prospects for the autonomous agent field itself. Will interest in agent technology continue to increase, or will the field fall victim to exaggerated expectations, as artificial intelligence did in the 1980s? There is good reason to be optimistic. There have been a number of commercial successes involving autonomous agents in the past year, particularly in the synthetic agents area. Computer game players increasingly demand games with lifelike characters, and synthetic agent developers have been able to meet the demand. Commercial Web applications such as Internet shopping services incorporate autonomous agents. Meanwhile, NASA has committed itself to autonomous agent development, in the aftermath of the success of the Mars Pathfinder mission. Meanwhile funding agencies such as DARPA and NSF have been increasing their support for agent research.

The agent community is an umbrella for a number of loosely related technical disciplines: softbots, mobile agents, animated characters, expert assistants, and autonomous robots. Of course, each group has a rather different set of technical concerns. However, regardless of the particular subfield, agent research tends to be broad and interdisciplinary, integrating a variety of techniques and algorithms into complete systems. The Autonomous Agents Conference serves the critical function of bringing agent researchers together to exchange information on shared problems and to better understand important differences.

6 Further reading

The proceedings and other conference materials from the Agents '97 conference may be ordered from the conference Web site at `http://www.isi.edu/isd/Agents97/info.html`. Also, look for the upcoming volume 8 issue of the *SIGART Bulletin*, which includes several expanded articles from the Agents '97 conference, and the first volume of the new *Journal of Autonomous Agents and Multi-Agent Systems*, which is edited by Katia Sycara, Conference Chair of the upcoming Agents '98 conference.

[1]We would like to acknowledge the contributions and efforts of a number of organizations and individuals in supporting the Agents '97 conference. Nick Jennings and Pattie Maes served as general co-chair and program co-chair, respectively. Joseph Bates, George Bekey, Oren Etzioni, and Anand Rao acted ably as technical program area chairs. Sheila Coyazo and Kary Lau contributed endless hours to the organization and management of the conference. We also wish to acknowledge James Lester, treasurer, Milind Tambe, local arrangements chair, Michael Wooldridge, publicity chair, Jöerg Müeller, tutorials chair, Charles Rich, software demonstrations chair, and R. James Firby, robotic demonstrations chair. The sponsorship and assistance of the following corporations and organizations is gratefully acknowledged: ACM SIGART, SIGCHI, and SIGGRAPH; AAAI; The University of Southern California; British Telecommunications plc, Intel, Microsoft, MERL – A Mitsubishi Electric Research Laboratory, Sun Microsystems, Silicon Graphics, and DigiLink Network Services.