AAAI-96 Workshop on Agent Modeling

Milind Tambe and Piotr Gmytrasiewicz

■ Agent modeling—the ability to model and reason other agents' knowledge, beliefs, goals, and actions—is central to intelligent interaction. The Workshop on Agent Modeling, held as part of the Thirteenth National Conference on Artificial Intelligence, was organized to bring together researchers working in these areas to assess the state of the art and discuss the common issues in representation and reasoning with models of agents.

An intelligent agent operating in a realistic environment will often need to interact with other agents to achieve its goals. *Agent modeling*—the ability to model and reason about other agents' knowledge, beliefs, goals, and actions—is central to intelligent interaction, and it is being investigated in a variety of research areas, including distributed AI and multiagent systems, plan recognition, natural language discourse, intelligent tutoring, and user interfaces, as well as in related areas, such as game theory and cognitive science and psychology. The Workshop on Agent Modeling, held as part of the Thirteenth National Conference on Artificial Intelligence, was organized to bring together researchers working in these areas to assess the state of the art and discuss the common issues in representation and reasoning with models of agents.

The workshop succeeded in drawing together researchers from a surprising variety of backgrounds and diverse concerns about agent modeling. The workshop presentations can broadly be divided into five categories in terms of modeled entities and areas of application: (1) self-modeling of an agent's own capabilities, (2) modeling of agent groups or teams, (3) agent modeling in natural language discourse, (4) modeling in intelligent tutoring–expert-critiquing systems, and (5) modeling agents in economic terms. Although agent-modeling applications were strongly represented, theories of agent modeling were underrepresented.

The task of weaving together the participant diversity was superbly accomplished by Ed Durfee's (University of Michigan) invited talk "Agent Models or Model Agents." As the title indicates, the talk referred to the continuum of applications of agent modeling: Agents can construct models of other agents or agent-groups or organizations (*agent models*), and these models serve to guide the agent's own behavior or those of others (*model agents*).

Two presentations at the workshop focused on an agent's self-modeling rather than on modeling others. Keith Decker, Katia Sycara, and Mike Williamson (all of Carnegie Mellon University [CMU]) discussed the capability representation of information agents—for advertising behaviors—for exploitation in organization development and plan coordination. G. Wickler and L. Pryor (both of University of Edinburgh) defined the problem of an agent's competence self-assessment; for example, answering, "Am I capable of solving a problem without actually solving the problem first?"

Among work on modeling agent groups, Sandip Sen and Thomas Haynes (both of University of Tulsa) discussed the impact of learning on group behavior, using the predator-prey domain for illustration. Individual agents evolve their models of other agents within a group by learning exceptions to their generic models. These exceptions are represented as cases that override the generic model. Milind Tambe's (University of Southern California Information Sciences Institute) work focused on one agent's modeling of other agent teams in dynamic multiagent environments such as RoboCup Soccer or real-world fighter-jet combat simulators. As well recognized in theories of collaboration, a team's joint goals and plans are not simply the union of individual agent's actions. This observation is key when inferring and modeling a team's goals from dynamic observations of team members' behaviors. Amol Mali (Colorado School of Mines) explored the impact of social laws—constraints on agent behaviors—on agent modeling.

The work reported on modeling agents for communication and natural language discourse underscored the relationship between discourse and modeling for recognizing plans and intentions in action. A paper by Nancy Green and Jill Lehman (both of CMU) related their modeling approach in NL-SOAR to modeling in a fighter-pilot agent domain. They point out that both problems amount to attempting to comprehend the stream of observational input and organizing it and processing it to arrive at an appropriate response. It turns out to be useful to model agents in terms of their beliefs, goals, and intentions for both physical interaction and communication. The unique distinction comes from the character of the changes brought about by action, which changes the physical environment, and the communicative acts, which change the state of beliefs of the agents involved.

Similarly, the work on turn taking in discourse, by Toby Donaldson and Robin Cohen (both of University of Waterloo), although it approaches the fundamental question of what should be said, by whom, and when, is rooted in the work on the rational agency approach to planning and acting by Bratman, Cohen, and Levesque as well as Georgeff and Rao. An interesting element in Donaldson and Cohen's work is their use of constraint satisfaction as a uniform

Intelligent Multimedia Information Retrieval

Edited by Mark T. Maybury

The expansion of the information highway has generated requirements for more effective access to global and corporate information repositories. These repositories are increasingly multimedia, including text, audio, graphics, imagery, and video. Now attention has turned toward the problem of processing and managing multiple and heterogeneous media in a principled manner, including their creation, storage, indexing, browsing, search, visualization, and summarization. Providing machines with the ability to interpret, generate, and support interaction with multimedia artifacts will be a valuable facility for a number of key applications such as video teleconference archiving, custom on-line news, and briefing assistants. This edited collection addresses fundamental issues in multimedia processing, including content-based retrieval of graphics, audio, imagery, and video; intelligent hypermedia access, and empirical studies.

ISBN 0-262-63179-2
590 pp., index, $40.00 softcover

To order, call toll-free:
1-800-356-0343 or 617-625-8569
The AAAI Press / The MIT Press
5 Cambridge Center
Cambridge, MA 02142-1493

The work reported on modeling agents for communication and natural language discourse underscored the relationship between discourse and modeling for recognizing plans and intentions in action.

method to manage the agent's own beliefs, desires, and intentions (BDIs) as well as the other agent's BDIs. David Pautler (Northwestern University Institute for Learning Sciences), in his work on automatically generating e-mail messages, concentrates on the issue of the appropriateness of responses, modeled as a network of abstract psychological, social, and practical effects.

Work on expert assistants and critiquing systems is one of the more traditional applications of agent modeling, which requires an agent to model the problem-solving processes of the interacting human to provide appropriate feedback. Ole Mengshoel and D. C. Wilkins's (both of University of Illinois at Urbana-Champaign) presentation focused on recognizing erroneous agent actions through plan recognition to provide such feedback. Fu and Hayes apply agent modeling to track a user's focus of attention in problem solving to provide appropriate feedback.

Work describing economic approaches to agent modeling takes the common perspective of an agent as a self-interested, payoff-maximizing entity. This perspective, although seemingly at odds with the spirit of the traditional paradigm of symbolic AI, is of particular interest for agent modeling because it precisely formalizes the notion of agent's rationality, used informally in other work. The paper by Sunju Park (University of Michigan) and Durfee described how agents can model others while they decide on how to contract their tasks in the digital library environment. The models of the other agents are needed to predict others' behavior; based on this information, the agents determine their optimal behavior, for example, the payment amount that maximizes the expected payoff. J. M. Vidal (University of Michigan) and Durfee's paper put forth the issue of the usefulness of modeling. For example, is it the case that the agents that make the effort to model other agents perform better than those not using the models of others? Further, does it pay to model how the other agents model others, and so on? It turns out that in a dynamic economic environment in which agents compete with one another while they buy and sell goods, the agents equipped with models of others generally get higher rewards. Interestingly, the advantage of modeling others is diminished when the volatility of the domain is low. Thus, the models of other agents are more useful in variable domains, when they are a particularly valuable guide to predict what the other rational agents are going to do.

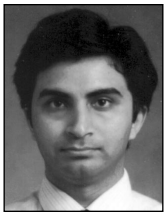

Milind Tambe is a research computer scientist at the Information Sciences Institute, University of Southern California (USC), and a research assistant professor with the Computer Science Department at USC. He received his Ph.D. in 1991 from the School of Computer Science at Carnegie Mellon University. His interests are in the areas of multiagent systems, specifically multiagent collaboration and agent modeling, and parallelism and real-time performance of AI programs, especially rule-based systems. His e-mail address is tambe@isi.edu.

Piotr Gmytrasiewicz is an assistant professor in the Computer Science and Engineering Department at the University of Texas at Arlington. His interests are in multiagent systems, treatments of uncertainty, and decision-theoretic reasoning in AI. His research centers on designing intelligent agents capable of coordination and rational communication with other agents. His e-mail address is piotr@huckle.uta.edu.