

AAAI 2000 Fall Symposium Series Reports

*Carolyn Penstein Rose, Reva Freedman, Mathias Bauer,
Charles Rich, Ian Horswill, Alan Schultz, Michael Freed,
Alonso Vera, and Kerstin Dautenhahn*

The American Association for Artificial Intelligence presented the 2000 Fall Symposium Series was held on Friday through Sunday, 3 to 5 November, at the Sea Crest Oceanfront Conference Center. The titles of the five symposia were

- Building Dialogue Systems for Tutorial Applications
- Learning How to Do Things
- Parallel Cognition for Embodied Agents
- Simulating Human Agents
- Socially Intelligent Agents: The Human in the Loop

Building Dialogue Systems for Tutorial Applications

It was a pleasure to participate in the symposium entitled Building Dialogue Systems for Tutorial Applications along with 44 researchers from a variety of relevant application areas. A number of tutorial dialogue systems at the cutting edge from across the United States and abroad were represented, including AUTOTUTOR (University of Memphis), THE GEOMETRY TUTOR (Carnegie Mellon University [CMU]), MISS LINDQUIST (CMU), CIRCSIM-TUTOR (Illinois Institute of Technology), BEE-TUTOR (University of Edinburgh), CALL (Universite du Maine, France), and ATLAS-ANDES (University of Pittsburgh).

To encourage as much interaction as possible, the entire symposium consisted of seven panel sessions, two demo sessions, and a poster session. Each symposium attendee played an active role in making this symposium a great success. Many thanks to all!

Each panel was composed of researchers from a variety of backgrounds, each with his/her own unique experience and perspective to share. In the initial and final panel discussions, we covered high-level issues such as what natural language potentially has to offer tutoring systems and what the most important target areas are for current research. In the remaining five panels, we covered meaty issues relating to system architecture, tutorial dialogue analysis, language generation, and language understanding. We addressed specific questions such as "What will give us the most bang for our buck?" "How

The American Association for Artificial Intelligence presented the 2000 Fall Symposium Series was held on Friday through Sunday, 3 to 5 November, at the Sea Crest Oceanfront Conference Center. The titles of the five symposia were

- Building Dialogue Systems for Tutorial Applications
 - Learning How to Do Things
 - Parallel Cognition for Embodied Agents
 - Simulating Human Agents
 - Socially Intelligent Agents: The Human in the Loop
-

can we develop tutoring systems that are engaging without getting too much in our students' faces?" "Under what circumstances should dialogue interfaces for tutoring systems offer students a chocolate in response to their input?" "Is 'hintify' an acceptable morphological construction in standard English?"

Nine working systems were showcased in two lively demo sessions.

Whether we were navigating a Navy ship using the COVE virtual environment or discussing the inner workings of a central processing unit with the AUTOTUTOR talking head, the demo sessions were both informative and entertaining. The poster session was stimulating and colorful, displaying some recent results and inviting more detailed discussion with individual authors. All in all, it was a weekend event well worth participating in.

— Carolyn Penstein Rose
Reva Freedman
University of Pittsburgh

Learning How to Do Things

Knowing how to do things is an important category of knowledge underlying many kinds of intelligent behavior in artificial agents, such as critiquing, advice giving, tutoring, collaboration, and delegation. In the current state of the art, most of this procedural knowledge is encoded "manually" by a single person (or a small team) who needs to be expert in both the task domain and the appropriate knowledge representation formalisms. This bottleneck is serious in the development of these kinds of system.

The focus of this symposium was how to automate or partially automate the acquisition of procedural knowledge, namely, indexed collections of what are variously called macros, plans, procedures, or recipes for action. Participants presented current research related to this goal from a number of different perspectives.

Andrew Garland (Brandeis University) et al. described their work in the context of an application-independent collaborative tool, called COLLAGEN. Good task models are a prerequisite for any kind of general approach to collaborative behavior. The authors discussed the various techniques that could, in principle, be applied to this knowledge-acquisition task, the general representational issues to be taken into account, and typical problems with the various learning approaches.

Introne et al. described VESSEL-WORLD, a scenario in which three users try to solve a common goal by coordinating their activities—a prob-

lem that is made more difficult by the fact that each participant has only limited perception capabilities. The system's task is to observe the users' actions, identify segments that serve one particular goal, and create a case base for future use, thus enabling the system to facilitate future coordinated problem solving.

Soller et al. dealt with the problem of supporting collaborative learning among humans. Their main focus was on training the system to recognize (and support) effective collaboration strategies using observed interactions as input. First experiments provided evidence that neural networks, or HMMs, can serve the purpose of acquiring the knowledge necessary to support a human learning group.

Mathias Bauer (DFKI) et al. discussed the problem of how effective communication can be achieved in cases where the agent to be instructed and its instructor share little knowledge. The application considered involves information agents that are to be programmed by demonstration to identify and extract relevant information on their users' behalf. A crucial aspect of this work is the exact identification of the problem source, which enables the system designer to enhance the agent's reasoning capabilities in a goal-directed way.

Another approach to programming by demonstration was present by Lau et al. The most important aspect of their work is the use of a version-space algebra to represent the various possible combinations of primitive actions. The system learns repetitive sequences of text-editing commands that can later be executed automatically when an appropriate situation occurs.

Scholer et al. dealt with the problem of acquiring the procedural knowledge needed by an intelligent tutoring system. In a virtual reality environment, the first acquisition step involves the trainer demonstrating how to perform certain tasks. Once the system has learned a stable action sequence, it starts experimenting on its own by leaving out one step at a time and checking the result of doing so in the simulated environment. The representation is thus refined and abstracted by identifying necessary precondition

tions and removing unnecessary actions.

Boylan et al. described a case-based system designed to help users draft effective business letters in English. Their approach combined a constructivist, self-learning philosophy with case-based reasoning technology.

Boicu et al. looked at the problem of how a subject matter expert, who does not have any knowledge engineering experience, can effectively "teach" an agent. The work used a civil engineering (for example, damaged bridge) domain and was evaluated as part of the Defense Advanced Research Projects Agency's High-Performance Knowledge Base Program.

The EXPECT Project, presented by Gil et al., is a "classical," broad knowledge-acquisition project aimed at facilitating the creation of knowledge-based systems in general. Procedural knowledge is just one aspect of this endeavor. EXPECT focuses on developing tools that enable the user to specify knowledge in English-like phrases, that is, without having to learn a formal language, with the system automatically identifying gaps and inconsistencies in the knowledge.

Alissandrikis et al. discussed a biologically inspired approach to learning in which agents try to acquire new skills by imitating other agents. Interesting effects can be observed especially in cases where both agents have dissimilar embodiments resulting in strongly diverging basic action repertoires. In these cases, the key factors affecting the learning result are a metric that allows the goodness of the imitated attempt to be assessed and a subgoal strategy that allows the imitating agent to concentrate on the vital aspects of an observed action sequence.

Finally, Marques and Pain examined some of the philosophical and psychological underpinnings of the question of what kinds of task model can be learned.

The symposium was organized by cochairs Mathias Bauer (DFKI) and Charles Rich (MERL) and a committee comprised of Andrew Garland (MERL), Abigail Gertner (MITRE), Eric Horvitz (Microsoft Research), Tessa Lau (University of Washington), Neal Lesh (MERL), James Lester (North Carolina

State University), Henry Lieberman (MIT Media Lab), Jeff Rickel (USC/ISI), and Candace Sidner (MERL).

— Mathias Bauer
DFKI

— Charles Rich
Mitsubishi Electric Research
Laboratory

Parallel Cognition for Embodied Agents

The AAAI Symposium on Parallel Cognition for Embodied Agents brought together researchers from a broad range of disciplines, ranging from behavior-based robotics to knowledge representation, to discuss the design of distributed parallel reasoning and control systems. The group was modest in size, allowing for in-depth presentations and discussions.

Much of the work presented involved attempts to combine behavior-based systems with more traditional knowledge representation and planning techniques. Eyal Amire presented work on reformulating the subsumption architecture in terms of first-order logic theorem proving. Monica Nicolescu presented an architecture based on abstract behaviors that can be assembled dynamically into task-specific networks. Ian Horswill presented a behavior-based architecture that allows for explicit reasoning about the states and capabilities of behaviors.

Other presenters described robots that were patterned on existing biological systems. Metta, Panerai, and Sandini presented work on a robot that learns to orient a reaching behavior using techniques suggested by human infant development. Brian Scassellati presented a partial implementation of current theories of shared social attention and theory of mind. Alan Schultz presented work on combining a reactive control system with the ACT-R cognitive model.

The lively discussions allowed authors to receive detailed feedback with a range of viewpoints. A good time was had by all.

— Ian Horswill
Northwestern University

— Alan Schultz
Naval Research Laboratory

Simulating Human Agents

Simulated human agents are a key software component in many kinds of application, including simulation-based training, interactive entertainment, and simulation-based tools for analyzing human-machine system designs. Creating sufficiently powerful and realistic human agents presents several challenges. To get the agent to behave capably in dynamic, time-pressured, and otherwise demanding application environments requires adapting state-of-the-art AI techniques. Making the human model accurate or believable requires identifying and incorporating relevant human performance data. Finally, reusable, well-documented software architectures are needed to reduce the time and expertise needed to construct new human-agent simulations.

The Symposium on Simulating Human Agents gathered the top researchers working in the area of Cognitive Modeling. SOAR, ACT-R, GOMS, FORR, AMBR, OMAR, BRAHMS, and APEX were among the modeling approaches and architectures represented. Symposium attendees included researchers from academia, industry, and government research-funding groups such as ONR, NRL, AFRL, and ARI. Marvin Minsky and Pat Hayes, the keynote and plenary speakers respectively, contributed to the presentations and discussion sessions. The symposium was designed to address practical questions about the incorporation of existing AI and human-performance modeling technologies into applications such as those listed previously. For example, what AI technologies are most relevant for simulating human behavior? What aspects of existing human-modeling architectures are most or least helpful for building new applications? Which aspects of human behavior are most worth capturing in a human modeling architecture, and which are well enough understood scientifically to be incorporated into general-purpose human-simulation tools? How should one best go about filling in the gaps where appropriate scientific findings do not yet exist?

These questions were motivated by an initial belief that bringing together researchers doing basic scientific mod-

eling with those doing or funding applied modeling would begin a dialog leading to both better research goals and more useful models. As it turned out, most participants had already determined that a closer coupling of applied and scientifically motivated modeling was needed. Many had begun bridging the gap in their own research. Perhaps the most significant consequence of the symposium was an increased awareness among participants of a nascent community of like-minded researchers with complementary methods.

The meeting achieved a broad critical evaluation of the state of the art in cognitive modeling, including a characterization of what is right and what is wrong with current modeling architectures and approaches. In particular, the shortcomings of current approaches were seen to be that models behave predictably (and are therefore, "gameable"), reactive planning is absent or highly restricted, important performance modulators (for example, stress, risk aversion, limited situation awareness) are inadequately accounted for, models are not robust, and models need to be far more adaptable. The symposium also resulted in the definition of achievable short- and long-term goals, based on these concerns, required to take cognitive modeling efforts to the next level of performance. Critical domains, target cognitive processes, and computational components were identified that need to be explored to achieve these goals. Given the broad representation of cognitive modeling approaches among the attendees, the symposium achieved a remarkable degree of conceptual convergence on the current concerns and goals in the area of simulating human agents.

— Michael Freed
NASA Ames Research Center


— Alonso Vera
University of Hong Kong

Socially Intelligent Agents: The Human in the Loop

The Symposium "Socially Intelligent Agents: The Human in the Loop" gave an overview of the state of the art of theory and applications in the active

and highly interdisciplinary area of socially intelligent agents (SIA). Much of this work is strongly inspired by forms of natural social intelligence characteristic of social animals, in particular, human beings (for example, communicating verbally and nonverbally, expressing and recognizing emotions, reading another agent's "mind"). Different from meetings in multigent systems (MAS) or distributed AI (DAI), this symposium discussed the design and evaluation of socially intelligent agents with the "human in the loop," that is, systems that can establish, maintain, and develop social relationships with human beings. Here, the human can find himself/herself, for example, in the role of user, observer, assistant, collaborator, competitor, customer, patient, or friend of such agents. The importance of such work is demonstrated in application areas such as electronic commerce; agents for training, learning, and therapy environments; agents for entertainment; and others. In all these application areas, the human user's attitudes toward the agent, in terms of believability, credibility, trust, and so on, are important factors that determine the acceptance and success of such a system and its utility in real-world applications. Therefore, an in-depth study is required of theories and models originally developed in areas such as psychology, brain research, ethology, and other fields not traditionally linked to the domains of AI and software engineering. Also, human-agent interactions need to be studied and evaluated carefully: Making agents "just like we are" is not necessarily desirable and/or feasible in particular application domains; compare discussions on autonomy versus control in human-machine interface design.

Environments and devices that interface humans (preferably nonintrusively) with agents and computers was another main theme of the symposium. It was generally felt that the more information an agent can acquire about intentions, emotions, beliefs, and other internal and mental states of a human, the better the agent can adapt to the human and predict his/her behavior and, possibly, changing attitudes. This is particularly rele-



Diagrammatic Reasoning

Cognitive & Computational Perspectives

Edited by Janice Glasgow, N. Hari Narayanan, & B. Chandrasekaran

“Understanding diagrammatic thinking will be of special importance to those who design human-computer interfaces, where the diagrams presented on computer screens must find their way to the Mind’s Eye.... In a society that is preoccupied with ‘Information Superhighways,’ a deep understanding of diagrammatic reasoning will be essential to keep the traffic moving.” — *Herbert Simon*

800 pp, index. ISBN 0-262-57112-9, softcover

Published by The AAAI Press / The MIT Press

To order call toll free: (800) 356-0343 or (617) 625-8569 or fax (617) 258-6779. *MasterCard and VISA accepted.*

vant in scenarios where agents are supposed to have “long-term” relationships with human beings and requires adaptation and learning abilities for agents, that is, being able to individually identify the human, as well as learning and acting on interaction histories with humans.

Particular research challenges presented at the symposium were projects that study heterogeneous agent societies, that is, how communities of agents (hardware or software) can be integrated and perform useful tasks in human societies. In other projects, the application area itself poses a particular challenge, for example, agent systems that are used in social problem solving for adults or therapy for children with autism. In these application

areas, it is not only desirable but an explicit goal that the agents be “persuasive,” that is, change behavior and attitudes of the humans interacting with them, and therefore requires careful consideration of ethical issues.

Presentations, working groups, and general discussions at the symposium demonstrated (1) significant advancements in the field, compared to the 1997 AAAI Fall Symposium “Socially Intelligent Agents”; (2) the importance of interdisciplinary work that could advance the development of a “theory of social minds,” both natural and artificial; and (3) the identification of particular research challenges (for example, unconstrained scenarios, heterogeneous scenarios) and challenge scenarios that can help to

further develop the field and systematically explore design spaces and spaces derived from application areas with particular requirements. Future developments in social, emotional, and narrative intelligence research might lead to “truly” socially intelligent agents, for example, agents that have “natural” (nontrivial) conversations with us, can recognize people as individuals and “mental agents” with distinct emotions and personality, and ultimately be agents that truly “care about us.” For more information on the symposium and the field of SIA, see homepages.feis.herts.ac.uk/~comqkd/aaai-social.html.

— Kerstin Dautenhahn
University of Hertfordshire