AI Magazine Volume 24 Number 1 (2003) (© AAAI)

AAAI-2002 Fall Symposium Series

Yukio Ohsawa; Peter McBurney, Simon Parsons; Christopher A. Miller; Alan Schultz; Jean Scholtz; Michael Goodrich; Eugene Santos, Jr.; Benjamin Bell; Charles L. Isbell, Jr.; and Michael L. Littman

The AAAI-2002 Fall Symposium Series was held Friday through Sunday, 15 to 17 November 2002 at the Sea Crest Conference Center in North Falmouth, Massachusetts. The five symposia in the 2002 Fall Symposia Series were

- Chance Discovery: The Discovery and Management of Chance Events
- Etiquette for Human-Computer Work
- Human-Robot Interaction
- Intent Inference for Users, Teams, and Adversaries
- Personalized Agents

The highlights of each symposium were presented at a special plenary session. American Association for Artificial Intelligence (AAAI) technical reports of most of the symposia will be made available to AAAI members.

Chance Discovery: The Discovery and Management of Chance Events

"Interesting keywords arose, such as serendipity, creativity, emergence, assertion, and...."

"You had a symposium on the creation of ideas by humans, did you?"

"Yes and no. We also talked about exploration, amplification, articulation, interaction, scenic information, subjectivity, and meaning."

"Hmmm, you considered the deepening of thoughts. I guess they are important for creation."

"That's right. In the first panel, we talked about prediction in dynamic environments, data mining, and...."

"So was it a conference on knowledge discovery inviting philosophers?"

"Not really. The first invited talk gave us deep insight into customer networks in the market, and the last panel extended to management, persuasion, communication, and trust, and so on. We found that the route, the context, and the timing of getting a piece of information during the process of discovery determines the value of the information."

"What was it about!?"

In this symposium, we had 17 papers, 2 invited lectures, and 14 other speakers. Six countries (Japan, United States, United Kingdom, Germany, Portugal, and the Czech Republic) were represented at the symposium, and the various research areas that were examined included psychology, management, economics, risk management, graph theory, biology, and AI. Through the discussions, we acquired keys to realizing chance discovery.

In a melting pot being heated up, things can adapt to the shape of the pot. A chance might mean no more than a sheer accident but also might be an event that might be rare but has significant impact on human decision. The occurrence of an event might be just an accident, but its significance to a human or an agent should have been determined in his/her/its mind not explicitly but lying with calm growth of desire and the shift of contexts. Such a prepared state of mind can be stimulated by information, sometimes given by a tool of data mining or a visualization technique.

However, even if you become aware of the value of a chance event, for example, with a new behavior of a customer in the market you are selling in, it is still hard to persuade your colleagues to make actions in response to the rare event. In the management of chance events, communication also becomes a significant issue. After the symposium, attendants are still continuing communications, seeking the next opportunity to meet and explore these ideas further. The symposium was a chance itself for all of us, without exception.

—Yukio Ohsawa,
University of Tsukuba
—Peter McBurney,
University of Liverpool
—Simon Parsons,
University of Liverpool

Etiquette for Human-Computer Work

The Symposium on Etiquette for Human-Computer Work began its meeting—with great propriety, of course—with a keynote address from Jeanne Comeau, an author, speaker, and teacher on etiquette and the director of the Etiquette School of Boston. Comeau taught us a great deal about etiquette's history and purposes as well as how to use our napkins and hold our forks.

Symposium participants wrestled with definitions of etiquette. Broadly speaking, two alternative but related

themes emerged: (1) Etiquette is a (frequently implicit) set of prescribed and proscribed behaviors that permits meaning and intent to be ascribed to actions between (human or machine) actors, thus facilitating group identification, streamlining communication, generating expectations, and so on. (2) Etiquette encodes "thoughtful consideration for others"; that is, etiquette operates (when obeyed) to make social interactions more pleasant, polite, and cooperative and (when violated) to make them insulting, exploitative, and unpleasant.

We heard research demonstrating that etiquette in machine behaviors can affect human plus machine performance in both good and bad ways. For example, the difference between a rude decision aid (intrusive recommendations and impatient queries) and a polite one is as large as the difference between 60- and 80-percent accuracy rates in the aid's recommendations in terms of users' trust and acceptance of these recommendations. That is, a polite but inaccurate aid was used and trusted about as a much as a rude but accurate one.

One panel debated whether taking an "etiquette perspective" offers anything different from traditional human-computer interaction design. Most participants believed it did, suggesting that it emphasized politeness and user consideration or the soft and often implicit constraints of social or work settings and modes. Another session conducted a post mortem of a highly familiar system whose design quite explicitly considered etiquette: the Microsoft OFFICE ASSISTANTS, such as the ubiquitous PAPERCLIP. The discussion produced several specific suggestions about how etiquette considerations might fine-tune office interactions.

In all, we had 30 attendees who gave 21 separate presentations on etiquette's implications for human-computer, computer-mediated human-human, and even computer-computer interactions. Subthemes emphasized (1) the role of etiquette or politeness in "good citizenship" in the digital age through implied or explicit social contracts and the ability to abuse etiquette's expectations in digital "cons"; (2) the ability of embodied agents to convey richer elements of etiquette; (3) the particular role and challenges of etiquette in pedagogical systems, where subtle nuances in how a tutoring system presents itself can have profound impacts on its effectiveness; and (4) the role of etiquette in the design of high-criticality systems, such as aviation, power generation, automobile driving, and military systems. Here, the role of the polite forms of etiquette can be less important, but the role of etiquette in conveying accurate indications of system intent, accuracy, and trustworthiness is extremely important. Although the participants were diverse in their interests and backgrounds, the consensus opinion was that etiquette forms a common, unifying theme that will take on increasing importance as computers become more complex and autonomous social actors.

—Christopher A. Miller
Smart Information Flow Technologies

Human-Robot Interaction

Robots are making their way out of the laboratory and into the real world. Therefore, researchers need to consider how to make these robots usable by real users. This symposium brought together 35 researchers representing a number of diverse domains: speech and language technology, systems architectures, simulations, AI researchers, and human-computer–interaction researchers.

Fifteen short presentations were given by participants in the symposium. After hearing these presentations, the group collectively outlined some of the challenges faced by various research approaches to human-robot interaction. To gather more perspective on these issues, the participants worked in groups and outlined research scenarios in the domains of healthcare, search and rescue, and military, leading to an elaborated set of challenges classified by the following seven research domains: (1) cognitive models, (2) natural language, (3) user interface issues, (4) user-centered design, (5) robot design issues, (6) the use of models in design and evaluation, and (7) team-related issues.

Cognitive Models

Issues include when it is useful for a robot to model human behavior: the level of cognitive model that is needed for various domains and tasks. Robots have to understand the meaning of behaviors; communication between human and robot is grounded in behaviors.

Natural Language

When is it appropriate to use a constrained vocabulary? What are the domains for which speech is the appropriate interaction modality? How much does the robot need to understand, and how much can the designer constrain the situation to effect communication? Can combining human speech and other information, such as that obtained through sensors, make inroads into speech issues? Mutlimodal interactions should be the goal, combining speech and other modalities as appropriate.

User Interface Issues

It is important to consider the roles of the robots and the users in designing the interface. Robots in some instances can be the interface, and users need to understand the causes and effects of the robots' behaviors. Tool kits for human-robot–interface design are needed.

User-Centered Design

Baseline studies are needed to produce user requirements. Collaboration between robot and user interface designers, both in the research world and in industry, is needed to produce usable interfaces. Context switching and situational awareness are key usability issues that arise in human-robot interfaces. Human-robot–interface research should concentrate on future robot capabilities, not current.

Robot Design Issues

Barriers to effective interaction with people include the difficulty in understanding peoples' intent, communicating the robot's intent to a hu-

man, and developing sensors to facilitate interaction.

The Use of Models in Design and Evaluation

There are many types of models in the design of human-robot systems that need to be integrated. Important models include models of the collaboration process, models of cognition, models of individuals and models of teams, and models of effect.

Team-Related Issues

Modeling individual activities and coordinating these activities in a specific context can be a key to effective teaming. In a team of humans and robots, both the humans and the robots can have multiple roles. How should these roles be selected, how are they communicated, and how can trust be managed?

The participants decided to put together a resource page where information from the various disciplines could be posted. A number of conferences and journals were mentioned that would be advantageous for presentation of human-robot–interface research.

—Alan Schultz,
Naval Research Institute
—Jean Scholtz, *National Institute of Standards and Technology*
—Michael Goodrich,
Brigham Young University

Intent Inference for Users, Teams, and Adversaries

This symposium was held in part as a continuation of the 2001 AAAI Fall Symposium on Intent Inference for Collaborative Tasks. Motivation for the 2002 meeting came from advances in AI that have enabled decision support systems to assume substantive roles in supporting human operators of complex systems. In particular, this symposium explored the broader context in which such systems are applied. As intent-aware applications become more capable of autonomous performance, they must engage more fully with human operators, negotiating tasks assignments; anticipating near-term needs; and proactively providing information,

analysis, and alerts. The symposium addressed work in intent inference that has shed much light on how automation systems can be given some measure of understanding of their users' tasks and needs. In the broadened scope of this year's topic, the symposium extended the discussion to systems with multiple operators and to research into team and adversarial intent inference. The notion of a team or crew is central to applications involving complex systems and organizations ranging from transportation systems to command-and-control centers. For adversarial intent inference, decision support for teams facing an intelligent opponent (hostile force) are limited in their utility without an understanding of the adversary's goals and actions. Thus, adversarial intent inference is a key capability in applications including counterterrorism, counterdrug, and asymmetric warfare.

The objective of this symposium was to bring together researchers and practitioners from various areas such as plan recognition, cognitive science, interface agents, gaming agents, and knowledge representation in an attempt to continue to promote the development of intent-aware decision support for multioperator complex systems. The symposium began with an invited talk on intent inference in the game of Go by Tim Huang (Middlebury College) that provided a look at how intent inference can be leveraged in adversarial situations. The remainder of the two-day meeting was devoted to presentations and group discussions on how intent inference from users, teams, and adversaries can be ascribed and used in a wide variety of problem domains. On the first day, participants focused on how the intent-inference process is determined and used in agent and cognitive models. On the second day, participants continued the discussion of intent inference with multiagent systems and the role of intent inference in plan recognition and task analysis. Finally, the last two hours of the second day were used to define the important concepts and processes of intent inference from various perspectives. Sheila Banks (Air Force Research Laboratory) summarized the symposium results in a plenary address, observing that there are many parts and pieces to intent inferencing that are being explored in a wide variety of contexts and applications. Successfully addressing intent inference involves a multidisciplinary look across different domains to best find our solutions.

—Eugene Santos, Jr.,
University of Connecticut
—Benjamin Bell, *CHI Systems, Inc.*

Personalized Agents

Although the term *agent* has come to mean many things, it perhaps has the most traction when identified with an anthropomorphized and autonomous program that acts as a personal assistant to a specific user (or set of users). In this model, the agent usually "lives" in a virtual world; can have access to data about its user; and is empowered to act on its user's behalf in a variety of computer-based tasks, including scheduling appointments, getting messages, engaging in negotiation with other users, discovering items of interest, and even initiating contact with other users and agents.

The state of the art in this nascent field is still quite primitive, although selected applications are growing in complexity and scope. Our symposium drew a set of enthusiastic and creative researchers from a variety of countries and academic backgrounds. We heard about exciting existing work such as sensor-laden homes, smart tutors, arbitration facilitators, network optimizers, social robots, bibliographic assistants, visualization apprentices, and planning aids. We also learned a bit about how personalized agents can be evaluated.

In our discussions, we identified a core set of issues that need to be addressed in the design of any personalized agent system: task, preference elicitation, interaction, action selection, and autonomy.

Task

What is the personalized agent responsible for doing? Is this a task that might be better addressed using some other human-computer–interaction paradigm? If the task requires another entity that might have access to complementary information, could take independent action for the user, and can benefit from adapting its decisions to a particular user, a personalized agent can be the right tool for the job.

Preference Elicitation

For the personalized agent to be effective, it needs to know what its user wants. Systems can vary in the format with which this information is provided (natural language, numbers, logical expressions), the channel by which it arrives (implicit inferences, explicit statements, data transferred from related tasks), and the initiator of this exchange (the agent, the user, or some combination). We demonstrated conclusively that a personalized agent must behave in a way that doesn't annoy the user.

Interaction

By what means do the user and personalized agent interact? What is the effective location of the agent, its outward appearance, its ability to initiate interactions, and its degree of context sensitivity?

Action Selection

How does the agent decide what to do? Is it hardwired, or can it use learning and adaptation to improve its performance? Does it undergo unilateral behavior optimization, or are its decisions primarily collaborative?

Autonomy

What can the system do on its user's behalf? How does it gain the user's trust (a spotless track record, transparency of decision making, proof of reliability)?

Although research into personalized agents is still maturing, the area is poised to take advantage of increases in computer power, sensing capabilities, and storage volume to help us fight the growing complexity of our daily lives.

—Charles L. Isbell, Jr.,
Massachusetts Institute of Technology
—Michael L. Littman,
Rutgers University