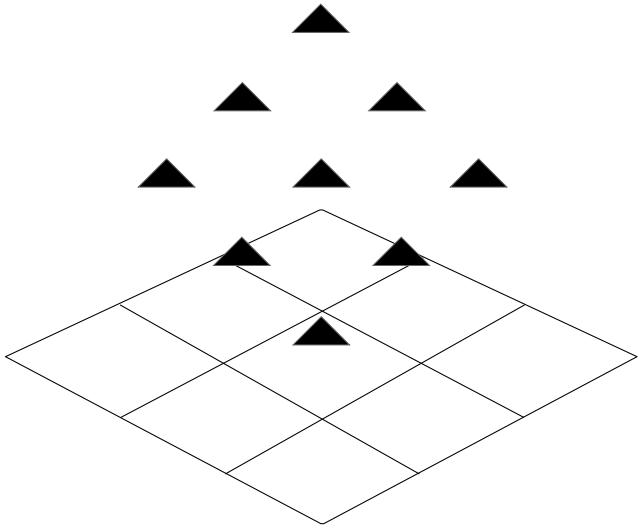

2000 Fall Symposium Series

Call for Participation

NOVEMBER 3–5

Sea Crest Conference Center, North Falmouth, Massachusetts

Sponsored by the American Association for Artificial Intelligence

The American Association for Artificial Intelligence is pleased to present its 2000 Fall Symposium Series, to be held Friday through Sunday, November 3-5, 2000 at Sea Crest Conference Center in North Falmouth, Massachusetts. The topics of the five symposia in the 2000 Fall Symposia Series are:

- 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

An informal reception will be held on Friday, November 3. A general plenary session, in which the highlights of each symposium will be presented, will be held on Saturday, November 4.

Symposia will be limited to between forty and sixty participants. Each participant will be expected to attend a single symposium. Working notes will be prepared and distributed to participants in each symposium. In addition to invited participants, a limited number of interested parties will be able to register in each symposium on a first-come, first-served basis. Registration information will be available in early July.

To obtain registration information, write to:
AAAI Fall Symposium Series
445 Burgess Drive
Menlo Park, CA 94025-3442
Voice: 650-328-3123
Fax: 650-321-4457
fss@aaai.org
www.aaai.org/Symposia/symposia.html

Submission Dates

- Submissions for the symposia are due on March 29, 2000
- Notification of acceptance will be given by May 5, 2000
- Material that will be included in the working notes of the symposium must be received by August 25, 2000.

See the appropriate section in each symposium description for specific submission requirements.

Building Dialogue Systems for Tutorial Applications

Studies of human tutoring have argued the importance of conversation between the student and the tutor in making tutoring interactions successful, suggesting that intelligent tutoring systems will be more effective if they can engage in dialogues with students. Building dialogue systems for tutorial applications presents a wide range of new computational challenges. Recent advances in computational linguistics have made it possible to make significant strides towards the development of highly interactive dialogue-based systems. The goal of this symposium is to bring together researchers working on all aspects of this diverse and interdisciplinary area of research.

We expect participation from many areas of artificial intelligence, including intelligent tutoring system development, computational linguistics, planning, user or student modeling, knowledge representation and probabilistic reasoning. We also encourage participation from the related fields of cognitive science, applied cognitive psychology and educational technology.

Relevant topics include, but are not limited to, the following areas of research: system architecture, dynamic curriculum planning, conversation planning, dialogue management, robust language understanding, discourse comprehension, text planning and generation, coordination of language and other modalities, user/student modeling, system evaluation, and field reports.

Submissions

The symposium will include paper presentations, panel discussions, a poster session, and time for demos and videos. Potential participants may submit a technical paper (up to 3,500 words) in any of the areas above, or a short paper (up to 1,500 words). Short papers can be an extended abstract, a position statement, or a description of a proposed demo, video or poster. Electronic submission in PostScript, PDF or Microsoft Word format is requested. Further information is available at www.pitt.edu/~circle/its-symp.html.

Organizing Committee

Carolyn Penstein Rose (Cochair), University of Pittsburgh; Reva Freedman (Cochair), University of Pittsburgh; Vincent Aleven, Carnegie Mellon University; Sandra Carberry, University of Delaware; Michael Glass, Illinois Institute of Technology; Art Graesser, University of Memphis; Nancy Green, University of North Carolina/Greensboro; Pamela W. Jordan, University of Pittsburgh; James Lester, NCSU; Susan McRoy, University of Wisconsin/ Milwaukee; Ronnie Smith, East Carolina University; Ingrid Zukerman, Monash University.

Parallel Cognition for Embodied Agents

The aim of this symposium is to draw together researchers working in the society of mind, behavior-based robotics, connectionism, cognitive science, and neuroscience to discuss advances in, and prospects for, a theory of high level cognition that is compatible with computational and neurophysiological constraints, and with grounding in an environment through sensors and actuators.

How can high-level cognition be efficiently grounded in interaction with the environment? Many of the most successful robot architectures are based on parallel collections of experts (behaviors, schemas, agents, etc.). Can these techniques be extended to more "cognitive" tasks, or does symbolic reasoning require a fundamentally different model of computation? If so, how do we ground that model in sensors and actuators?

How can parallel systems be extended to more expressive representations? The binding problem limits most parallel systems to propositional reasoning. How much binding (and therefore how much predicate inference) can we plausibly implement in parallel?

What kinds of operations can plausibly be parallelized? Many common AI algorithms, such as unification, are believed to be unparallelizable. Does this mean we should reject them, use limited versions that are parallelizable, or simply accept the seriality of these computations?

When is parallelism really necessary? While human cognition is clearly parallel over short time scales, it is often argued that it looks serial over medium-to-long time scales (>100 ms). Is this where the real action is?

Should parallel inference systems even look like classical AI systems? Many systems try to simulate the operation of classical AI inference systems using parallel hardware. Other systems (e.g. society of mind) adopt architectures and semantics that are fundamentally different from conventional serial inference systems.

What does the biological evidence really say about high-level cognition in humans and other animals?

Researchers interested in participating should send a statement of interest (1–2 pages) or a full paper (less than 8 pages) to Ian Horswill ian@cs.nwu.edu

Symposium Chairs

Ian Horswill, Northwestern University and Alan Schultz, Naval Research Laboratory.

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 is a serious bottleneck in the development of these kinds of systems.

The focus of this symposium is on 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. The techniques for acquiring this knowledge may depend on many variables, including size of the domain (e.g., number of recipes); amount of input data; number of steps in a typical task; type of tasks (e.g., analysis vs. synthesis); number of agents involved (e.g., one, two, or many); type of agents involved (e.g., human versus computer); intended use of the knowledge (e.g., acting, critiquing, etc.); degree of supervision (e.g., teaching versus unsupervised learning); level of abstraction (e.g., primitive operations versus high-level goals); and degree of initiative (e.g., learning by experimentation versus passively).

Because of this problem diversity, we hope to include participants in the workshop from a number of research areas, including:

- Programming by demonstration (highly supervised, small amount of input data)
- Data mining (unsupervised, large amount of input data)
- Case-based problem solving (cases are like recipes, especially if abstracted)
- Machine learning (range of techniques)
- Cognitive and social sciences (e.g., studies of human instructional dialogues)
- Instructable agents

Submissions

Potential participants should submit a short position paper (maximum three pages) containing the following elements. (1) Primary contact: name, affiliation, postal and email addresses, telephone and fax numbers. Invitations to secondary authors will be made only if they are also listed on this submission; (2) Statement and discussion of two or three important research questions that could be presented and discussed at the symposium; (3) Statement and discussion of a domain that could serve as a shared example for the symposium. Explain how this particular domain would help make our discussion more concrete and productive; A short summary of authors' relevant work, including references (please supply URLs if available).

Please e-mail submissions (plain ascii text only) to learninghow@dfki.de. Confirmation of receipt will be returned by e-mail.

Organizing Committee

Mathias Bauer, DFKI, Cochair (bauer@dfki.de); Charles Rich, Mitsubishi Electric Research, Cochair (rich@merl. com); Andrew Garland, Brandeis; Abigail Gertner, University of Pittsburgh; Eric Horvitz, Microsoft Research; Tessa Lau, University of Washington; Neal Lesh, Mitsubishi Electric; James Lester, North Carolina State; Henry Lieberman, MIT; Jeff Rickel, USC/ISI; Candace Sidner, Lotus Development.

Simulating Human Agents

Simulated human agents are a key software component in many kinds of applications including, for example, simulation-based training, games and other forms of 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.

This symposium will address practical questions about the incorporation of existing AI and human performance modeling technologies into applications such as those listed above. Questions to be addressed during the symposium include:

- What AI technologies are most relevant for simulating human behavior? How should these be improved or adapted?
- What aspects of existing human modeling architectures are most / least helpful for building new applications? How can they be improved to become more useful to applications developers?
- Which aspects of human behavior are most worth capturing in a human modeling architecture, generally or for a given application area?
- What relevant scientific findings are "ready" 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?

The symposium will consist of invited talks, individual presentations, and group discussion. Those interested in attending should submit a research paper of no more than 10 pages, a 3-page position paper, or 1-page statement of interest to mfreed@mail.arc.nasa.gov. Only electronic submissions will be accepted. We also invite demonstrations of existing systems; email a brief description of the proposed demo to the above address.

Organizing Committee

Michael Byrne, Rice University; Ron Chong, Soar Technology; Michael Freed (chair), NASA Ames Research Center; Randy Hill, USC/ISI; Lewis Johnson, USC/ISI; John Laird, University of Michigan; Frank Ritter, Pennsylvania State University

Socially Intelligent Agents
The Human in the Loop

The highly interdisciplinary area of socially intelligent agents has attracted a number of active researchers who model, design and analyse agents (software or robotic) which behave socially. Much of this work is strongly inspired by forms of natural social intelligence characteristic of humans. This symposium will address recent technological, methodological and theoretical developments in the field of socially intelligent agents (SIA's), as well as discuss social and cultural issues, and limitations and problems of socially intelligent agents. A focus will be the issue of the "human-in-the-loop."

Both agents and humans can have different roles during agent-human interaction, for example, as designers, users, observers, assistants, collaborators, competitors, customers, or friends. The symposium will concentrate primarily on socially intelligent agents that are either directly interacting with humans, showing aspects of human-style intelligence, supporting interaction among humans and/or modelling explicitly aspects of human social intelligence.

The symposium will focus on four key themes for which considerations of the "human-in-the-loop" are crucial. Interdisciplinary approaches are particularly encouraged.

The symposium will comprise keynote talks, panel discussions and individual paper presentations, addressing one or several of the following four key themes:

Connecting to SIA's: architectures and design spaces for SIA's; innovative user-interfaces, novel environments and new methodologies for software and robotic agents interacting and collaborating with humans and facilitating communication and collaboration between humans; hot approaches (emotional, empathic aspects) and cold approaches (intention and plan ascription, reasoning etc.); synchronisation in human-agent dialogue; the role of embodiment in human-agent interaction; exploiting anthropomorphism; believability and degrees of agent complexity

Learning and playing with SIA's: new applications of social agent technology in rehabilitation and education; SIA's as instructors, guides, teachers, assistants and friends; SIA's which support human creativity and imagination; SIA's in living environments (for example, at school, at home, at work, on holiday, at meeting points).

Living with SIA's: social agent technology which influences attitudes/ opinions /behaviour; issues of social relationships between human and agent such as helping, competition and cooperation, autonomy and control, predictability, deception, manipulation, initiative, delegation, responsibility, conflicts.

Growing up and evolving with SIA's: social agent technology which empowers humans, addressing the cognitive and emotional needs of humans; impact of SIA's on human society and culture; agents adapting to and supporting cultural diversity; ethical considerations

For more detail see www.cyber.rdg.ac.uk/people/kd/WWW/SIA-2000.html or contact Kerstin Dautenhahn (K.Dautenhahn@cyber.reading.ac.uk).

Organizing Committee

Elisabeth Andre, DFKI GmbH, Ruth Aylett, University of Salford, Cynthia Breazeal, MIT AI Lab, Cristiano Castelfranchi, Italian National Research Council, Justine Cassell, MIT Media Lab, Kerstin Dautenhahn (Chair), University of Reading, Francois Michaud, Universite de Sherbrooke, Fiorella de Rosis, University of Bari.

American Association for Artificial Intelligence
445 Burgess Drive
Menlo Park, California 94025 USA
650-328-3123 • 650-321-4457 (fax)
fss@aaai.org
www.aaai.org/