

Reports on the AAAI Fall Symposia

23–25 October 1998

The American Association for Artificial Intelligence (AAAI) held its 1998 Fall Symposium Series on 23 to 25 October at the Omni Rosen Hotel in Orlando, Florida. This article contains summaries of seven of the symposia that were conducted: (1) Cognitive Robotics; (2) Distributed, Continual Planning; (3) Emotional and Intelligent: The Tangled Knot of Cognition; (4) Integrated Planning for Autonomous Agent Architectures; (5) Planning with Partially Observable Markov Decision Processes; (6) Reasoning with Visual and Diagrammatic Representations; and (7) Robotics and Biology: Developing Connections.

Cognitive Robotics

Research in cognitive robotics is concerned with the theory and implementation of robots that reason, act, and perceive in changing incompletely known, unpredictable environments. Such robots must have higher-level cognitive functions that involve, for example, reasoning about goals, actions, the cognitive states of other agents, and time as well as when to perceive and what to look for. In short, cognitive robotics is concerned with integrating reasoning, perception, and action within a uniform theoretical and implementation framework.

Cognitive robotics is at the crossroad between knowledge representation and robotics. In a sense, it is a reaction to pure deliberative knowledge representation in favor of principled theories of how deliberation relates to reactivity, perception, and so on. Robots implemented on these principles provide a criterion for success.

The symposium proposed several overviews and invited talks to get the feeling of the state of art and point out similarities and differences between

the various research projects. Specific issues were addressed by panels. Technical papers were presented in a poster session (available as a AAAI Press technical report). Finally, a joint session with the Symposium on Integrated Planning for Autonomous Agent Architecture took place to allow interaction with a different community.

Most participants were researchers working in cognitive robotics. Some participants were working in robotics and were looking for deliberative-level techniques to resolve perception-actuation uncertainty. Others were working in planning and reasoning about actions and were willing to test their work on actual robots.

A lot of discussion was generated. One of the hot topics was competitions for cognitive robots. Interesting

The American Association for Artificial Intelligence (AAAI) held its 1998 Fall Symposium Series on 23 to 25 October at the Omni Rosen Hotel in Orlando, Florida. This article contains summaries of seven of the symposia that were conducted: (1) Cognitive Robotics; (2) Distributed, Continual Planning; (3) Emotional and Intelligent: The Tangled Knot of Cognition; (4) Integrated Planning for Autonomous Agent Architectures; (5) Planning with Partially Observable Markov Decision Processes; (6) Reasoning with Visual and Diagrammatic Representations; and (7) Robotics and Biology: Developing Connections.

proposals on a specific AAAI robot competition were made. More information is available at www.dis.uniroma1.it/~cogrob98.

Giuseppe De Giacomo
Università di Roma La Sapienza

Distributed, Continual Planning

This symposium brought together researchers from classical planning, intelligent agents, Markov decision processes, machine learning, and mixed-initiative systems to explore possibilities for extending AI planning systems beyond the traditional closed-world, static view. The goal of the symposium was to share current approaches and discuss future research directions for planning applications in dynamic, changing environments in which the planning process is distributed across multiple agents and sites and in which planning is interleaved with execution.

The symposium was small but lively and productive. The participants included researchers from diverse fields, many of whom had not met before. Much of the discussion centered on coming to a common understanding of the problem, which was a useful exercise for all.

Although there was consensus that members of the community have a long way to go, the research presented at the symposium represents important first steps on the path toward true distributed, continual planning. The presentations included representations for user guidance, the planning of context, the monitoring of classes, and the planning of stances; techniques for goal transformation, plan merging, interagent communication, the incorporation of user guidance, and the localization of planning decisions; and frameworks for planning at variable levels of abstraction, incorporating context into planning and replanning, and formalizing notions of collaboration, intentions, and commitments.

A significant part of the symposium was set aside for open discussion. These sessions raised a number of recurring themes that participants believe will be central for focusing future research in this area. The themes included flexible use of abstraction; the interaction of time, abstraction levels, and uncertainty; communication and intentions in multiagent planning; and incorpo-

The final “moral” could be summarized as, When engineering artifacts, we should not put in them more emotion than what is required by the complexity of the system-environment interaction.

rating utilities, priorities, and soft constraints in planning and execution.

Marie desJardins
SRI International

Emotional and Intelligent: The Tangled Knot of Cognition

This multidisciplinary symposium brought together researchers in artificial and human emotions in an effort to investigate the roles that emotions play in grounding intelligent behavior and the possibilities for cross-fertilization between theoretical and engineering approaches to the study of emotion.

Presentations were organized around seven main topics: (1) emotion, cognition, and learning; (2) architectures for emotion; (3) emotion in social interaction; (4) emotion expression and recognition; (5) psychological models of emotion; (6) emotion in arts and entertainment; and (7) theoretical issues. Lively discussions took place at the end of each session and were furthered in small working group meetings.

The last session was devoted to identifying open problems and issues to be taken into account when designing artificial systems endowed with (some aspects of) emotions. Consensus was reached on several ideas, including the following:

First, the degree of complexity of emotional phenomena that is meaningful to model within an AI system strongly depends on whether our concern is an engineering or a scientific one.

Second, the designers of artificial systems cannot assume the adaptive character of emotions only on the grounds of biological evidence. The contributions of emotions must be shown by our results and carefully evaluated.

Third, the design must be guided by the requirements on the system and on the interactions with its particular environment, and our architectures must be related with the desired functions.

The final “moral” could be summarized as, When engineering artifacts, we should not put in them more emotion than what is required by the complexity of the system-environment interaction.

Dolores Canamero
Technical University of
Catalonia and IIIA,
Spanish Scientific Research
Council

Integrated Planning for Autonomous Agent Architectures

The Integrated Planning for Autonomous Agent Architectures symposium was intended to address two primary questions: (1) what do robots need to plan and (2) how should this planning be integrated with their architectures?

Regarding the first question, participants decided that for many domains, planning is not necessary for a robot in part because most existing robots do not have a particularly rich set of interactions with the world. We speculated that planning will be more useful for future robots with larger action sets. We sought to identify characteristics of domains where planning does appear to be useful. These domains can possess low observability, irreversible actions, multiagents, or conflicting goals. The question of a domain's observability was particularly interesting because it asserts that a richer sensor suite reduces the need for planning, which was demonstrated by Carnegie Mellon University's tour-giving robot SAGE. Also of interest was the idea that experimentation (trial and error) can be a viable alternative to

planning in certain situations.

Regarding the second question, we decided that in domains where planning is important, it should not necessarily be done the “classical” way. That is, planning is a process by which an execution architecture can be provided with imperfect but useful information. Planning needs to be evaluated by trade-offs between the information gained and the computational effort required. Ideas we discussed for integrating planning into an agent architecture included using limited planning horizons, using anytime planning, accepting suboptimal plans, and maintaining planning abstractions for as long as possible.

Glenn Wasson
University of Virginia

Planning with Partially Observable Markov Decision Processes

In recent years, there has been growing interest among researchers in AI, psychology, engineering, and other areas in a mathematical model for planning. The model, developed in the 1960s and 1970s by the operations research community, is an elegant formalism for describing and solving sequential decision problems in stochastic domains where there is hidden state. The main purpose of the symposium on Planning with Partially Observable Markov Decision Processes (POMDPs) was to gather interested researchers from these different disciplines to discuss current and future research issues.

A highlight of the symposium was the presence of researchers from outside AI who shared their unique views and perspectives. Two researchers involved in the early ground-breaking work on the POMDP model provided comments and suggestions that were extremely valuable for placing recent POMDP developments into perspective; the increased speed of computers and our improved understanding of computational and algorithmic issues are two of the main new tools AI researchers bring to the table. From the field of psychology, we heard how the POMDP model relates to classical

animal-learning experiments (for example, pigeons pecking for food). The ensuing discussion was one of the highlights of the symposium and one that clearly warrants further exploration.

From within the computer science community, there was representation from a broad range of research areas. We discussed how generalizations and specializations of the POMDP model, driven by the concerns of robotics and other applications, affect planning, learning, and computational complexity. Interesting POMDP variations, including models with a continuous representation of time, models with a risk-sensitive optimality criterion, and reinforcement learning models were discussed and found to be far less complex than many of us originally thought. This realization will likely influence future research in POMDP applications and algorithms.

Overall, it was satisfying that many connections between researchers in different disciplines and subfields were made. Because this was the first organized meeting of POMDP-oriented researchers, many of the attendees were grateful for the opportunity to meet and learn from their colleagues. A series of exciting and fruitful collaborations from the attendees can be expected in the years to come.

Michael Littman
Duke University

Reasoning with Visual and Diagrammatic Representations

Diagrammatic languages are ubiquitous in almost every human endeavor: Virtually all technical disciplines use diagrammatic notations, and non-technical fields use their own diagrammatic systems, for example, choreography notation.

Consequentially, visual human-computer interfaces have become standard. However, these interfaces rarely exploit the rich structure and complex semantics of diagrammatic languages. To envision intelligent diagrammatic interfaces, new ways of interpreting and manipulating diagrammatic notations on a high level

of abstraction need to be devised.

This challenge was taken up by the Symposium on Reasoning with Visual and Diagrammatic Representations. Previously, work on these topics has been carried out by largely unconnected groups in such diverse fields as AI, visual languages, spatial reasoning, and other neighboring disciplines. The primary aim of the symposium, therefore, was to strengthen the dialogue among these communities, and as expected, it attracted contributions and participation from a number of different fields.

Among the most actively discussed questions were foundational issues; for example, What distinguishes diagrammatic representations from other types of visual representations? Does diagrammatic reasoning require different, new modes of processing that are essentially diagrammatic in nature? More technically oriented questions regarding the analysis of diagrammatic systems and of formal methods for interpreting and manipulating diagrammatic notations were discussed in 10 contributed papers, 2 invited tutorials, and 2 invited papers. Nine system demonstrations and a panel discussing how visual language and diagrammatic reasoning research can benefit from each other completed the program.

It is our sincere hope that the fruitful exchange initiated here will be continued in subsequent events.

Gerard Allwein
Indiana University

Kim Marriott
Monash University

Bernd Meyer
University of Munich

Robotics and Biology: Developing Connections

There is a growing body of work developing at the intersection of biology and robotics. In particular, robotic systems are being built that are not merely biologically inspired but designed as accurate models of particular biological systems. The performance of these systems can thus be used directly to evaluate biological hypotheses as well as to test how well biological mecha-

nisms perform as robot controllers.

This symposium was a gathering of researchers working in this area. There was substantial variety in research backgrounds and systems studied (from ants to humans). Nevertheless, there emerged a consensus about the research methods leading to the most productive advances:

First is modeling a specific animal, for example, a cockroach-based walking machine, as opposed to generic animal-like designs. A target animal usefully constrains and focuses both the design of the robot and the experiments to be performed. Thus, the potential for mapping robot results back to the biology is increased, setting up a research interaction rather than a one-way flow of ideas.

Second is building a physical system, for example, an underwater robot to model lobster chemotaxis. Simulation can be a useful preliminary step but cannot replace real-world testing in evaluating sensorimotor systems that must deal with complex environments. Physical systems are useful for discovering the critical factors in environments that have direct application in biological research. This methodology also raises a number of serious issues for technological development, such as muscle-like actuators.

Third is testing the system by replicating the experimental conditions used for testing the animal, which exploits the many well-developed methodologies for behavior evaluation that already exist in biology, for example, directly replacing the fly in a flight simulator with an analog VLSI circuit modeling the fly's optomotor response.

Potential collaborations and identification of areas of biology with good potential for applying this approach were among the results to emerge from this symposium, pointing to a strong future for this field.

Barbara Webb
University of Nottingham

Tom Consi
Massachusetts Institute of Technology