

AAAI 2001 Spring Symposium Series Reports

Lorraine Fesq, Ella Atkins, Lina Khatib, Charles Pecheur, Paul R. Cohen, Lynn Andrea Stein, Michael van Lent, John Laird, A. Proveti, and S. Tran Cao

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2001 Spring Symposium Series on Monday through Wednesday, 26 to 28 March 2001, at Stanford University. The titles of the seven symposia were

- Answer Set Programming: Toward Efficient and Scalable Knowledge Representation and Reasoning
- Artificial Intelligence and Interactive Entertainment
- Game-Theoretic and Decision-Theoretic Agents
- Learning Grounded Representations
- Model-Based Validation of Intelligence
- Robotics and Education
- Robust Autonomy

Answer Set Programming: Toward Efficient and Scalable Knowledge Representation and Reasoning

The goal of this symposium was to provide a forum for researchers and practitioners to discuss the recent developments, both in theory and implementation, in answer set programming (ASP).

ASP can be seen as bringing together the ability of default logic to capture commonsense reasoning and the technology underlying logic programming. It is based on the view of program statements as constraints on the solution of a given problem. Subse-

quently, each model of the program encodes a solution to the problem itself. For example, an ASP program encoding a planning scenario has as many models as valid plans. This schema is similar to that underlying the application of propositional satisfiability (SAT) algorithms. In fact, the ranges of applicability of these two techniques are similar.

Several ASP systems are now available, among them *DeRes*, *DLV*, *SMODELS*, and *XSB*; they support provably correct inferences and are about as fast and scalable as SAT checkers. These exciting results are attracting the attention of researchers from fields such as plan-

The American Association for Artificial Intelligence, in cooperation with Stanford University's Department of Computer Science, presented the 2001 Spring Symposium Series on Monday through Wednesday, 26 to 28 March 2001, at Stanford University. The titles of the seven symposia were

- Answer Set Programming: Toward Efficient and Scalable Knowledge Representation and Reasoning
- Artificial Intelligence and Interactive Entertainment
- Game-Theoretic and Decision-Theoretic Agents
- Learning Grounded Representations
- Model-Based Validation of Intelligence
- Robotics and Education
- Robust Autonomy

ning, cryptography, and verification.

The symposium began with Bart Selman's invited lecture on randomized methods for SAT. Article presentations were divided into six thematic

and two poster sessions. Overall, 21 regular, 8 posters, and 3 position papers were presented. The thematic sessions were entitled ASP from a Different Perspective, Computational and Representational Aspects, Parallelism, and Theory. We also had several system demonstrations and a rump session where participants presented their works in progress.

The attendance was strong, with 52 participants overall, 5 of whom were from Bay Area companies, 21 who were graduate students, and 17 who were from outside the United States and Canada.

Articles from the symposium can be retrieved from www.cs.nmsu.edu/~tson/ASP2001/.

— A. Proveti

University of Texas, El Paso

— S. Tran Cao

Stanford University

Artificial Intelligence and Interactive Entertainment

The 2001 AAAI Spring Symposium on AI and Interactive Entertainment brought together AI researchers and commercial game developers with the goal of building a community interested in applying AI research to commercial computer games. The approximately 70 attendees were about an even mix of AI researchers (academic and industry) and commercial game developers. From the perspective of the game developer, talking to the AI researchers can suggest new ideas, techniques, and algorithms that can make the agents in games appear to be smarter. However, experience has shown that more intelligent games aren't always more fun. The first session of the symposium consisted of a group exercise to discuss the issue of what AI behaviors make games fun and what behaviors decrease the fun? Among the conclusions was the feeling that fun AI agents should seek to lose (but just barely) to the human player rather than win and the belief that fun AI agents must exhibit personality.

From the perspective of the AI researcher, talking to the game developers can suggest new applications for research, constraints on solutions, and research topics. The last session on the

first day of the symposium sought to highlight these issues through a group exercise designed to identify what AI advance would be most useful to the game industry. The most frequently mentioned candidate was natural language interaction (interestingly, speech recognition and voice synthesis were viewed as less important). The remaining two sessions of the first day allowed a number of researchers to present their work in the areas of personality and humanlike characters and interactive narratives.

The second day of the symposium kicked off with two sessions discussing different approaches to agent architectures. The question of an “AI engine” (similar to the graphics engines used in many games) was discussed and seemed to meet with less opposition from the game developers than in previous years. The afternoon involved two presentations on AI for strategy games and reports from three ongoing efforts in applying AI research to gamelike environments.

The third day started with an interesting talk-demonstration of the soon-to-be-released game *Flogian Bros.* by Visual Concepts. This Sega Dreamcast game includes multiple “sidekick” characters that demonstrate many of the AI issues discussed earlier in the symposium, including personality, interactivity, and fun. The symposium wrapped up with a discussion of the current state of academic-industry cooperation, and it could be improved. As in previous years, even the subtlest inquiries about research funding from game companies were a great source of amusement for the game developers. However, discussions of internships and hiring graduates were taken seriously.

The symposium was a great success. Over the course of the symposium, a number of subgroups coalesced and made plans to start mailing lists. In addition, we’re working on a web site that will archive the three previous symposiums and hopefully provide a forum for discussion and cooperation until the 2002 spring symposium. Look for more information soon at www.aigames.org.

— Michael van Lent

— John Laird
University of Michigan

Game-Theoretic and Decision-Theoretic Agents

Over the last few years, decision and game theories have proved to be powerful tools with which to design autonomous agents and understand interactions in systems composed of many such agents. Decision theory has been adopted as a paradigm for designing agents that can handle the uncertainty of any moderately complex environments and act rationally to achieve their goals (or preferences). Decision theory defines *rationality* as behavior that maximizes the expectation of the degree to which the preferences are achieved (that is, behavior that maximizes the expected utility). Game theory, building on the assumption that agents are rational and self-interested, uses the notion of Nash equilibrium to design mechanisms and protocols for various forms of interaction and communication that result in the overall system behaving in a stable, efficient, and fair manner.

The symposium brought together researchers interested in theoretical aspects of these theories and those interested in applications. As an example of a challenge on the theoretical side, Kenneth Arrow’s invited talk “Some Thoughts on Decision Theory, Automation, and Discretionary Judgment” pointed to a difficulty of computing optimal behavior under computational constraints. Intuitively, deciding on a rational course of action can be difficult if the computation itself is costly. Attempts at handling this issue by including the cost of computation as one of the optimality criteria usually makes the problem even more difficult and costly to solve, thus indicating that a degree of arbitrariness somewhere in the solution procedure might be necessary. A related issue of optimality of actions versus optimality of plans was brought up by John Pollock. He pointed out that given that the length of possible plans of action is unrestricted, it might be computationally difficult to formulate and choose between all the plans. These issues brought on a general discussion about the means that can be used to automatically define decision problems in sound decision-theoretic

terms and in the role that qualitative decision theory could play in this context.

Other theoretical issues discussed were related to usefulness and applicability of Nash equilibrium as a solution concept for problems of designing intelligent distributed systems. Frequently, the assumption usually invoked to justify the use of Nash equilibria—that of agents having common knowledge about various aspects of the environment and the interactions—seems unrealistic. Further, in some cases, there can be multiple equilibria without a clear way to choose among them, and in still other cases, there can be no equilibrium solutions at all. In general, it might be necessary to explicitly consider the current states of knowledge of the interacting agents to study solutions, off-equilibrium properties, and transient phenomena in such systems. Promising approaches from this perspective include hypergames and the decision-theoretic approach to game theory; they can be used to show that in effect, “there is life after Nash.”

Still another fundamental issue brought up was related to the usefulness of the paradigm of rationality, as defined in decision theory, applied to distributed intelligent systems in the first place. Take the example of the game of 1000 pennies, during which agents alternate in choosing to take coins from a pile of, initially, 1000 pennies. If an agent takes one penny, then the game continues, but if any of the agents decides to take two pennies, then the game terminates after this agent’s turn. In this particular interaction, rational agents could use backward induction to conclude that it is better to end the game at the first round by choosing two pennies, thus leaving one agent with the total of two pennies and the other agent with nothing. Less perfectly rational agents, however, could play this game profitably until all pennies from the pile are exhausted and end up with close to 500 pennies each. This, and similar, examples show that decision-theoretic rationality exercised simultaneously by multiple agents is not without its pitfalls in some special cases, particularly when the interactions and the

states of agents' knowledge assume exotic and, hopefully, rare forms.

In spite of the outstanding theoretical issues, applications of the decision- and game-theoretic paradigms have been promising. In his invited talk "Dimensions of Representations for Acting under Uncertainty: What We Want and Why We Can't Have It," David Poole gave an excellent unifying overview of the representation and solution concepts used for automated decision making. He also outlined a number of methods that can be used to make solutions practically feasible in the context of Bayesian networks and influence diagrams, which are the tools of choice for applied work in automated decision making and game theory. One specific approach to constructing approximate solutions, presented by Daphne Koller and collaborators, uses max-norm approximation to the value function. In another paper, also by Koller and colleagues, influence diagram formalism was extended to include the presence of multiple agents. The authors showed how game-theoretic equilibria can be defined within this representation and how the notion of strategic relevance can be used to compute the equilibria solutions efficiently.

Other applications of decision- and game-theoretic approaches presented included agent-based systems for decision support of humans and organizational units, where adjusting the autonomy of agents is an important issue, and a planetary rover control algorithm based on Markov decision processes. Electronic-commerce applications presented included a risk-averse auction mechanism integrated with a supply-chain management system, probabilistic pricebots used to automatically set prices on goods offered by online sellers, and a design for flexible trading agents operating in general auction and user settings.

Finally, a number of logical and other alternative approaches to decision making in single-agent and multiagent scenarios were presented. These approaches included comparisons of belief, desire, and intention (BDI) frameworks to qualitative decision theory; conflicts between desires and ways to resolve them; logical formula-

tion of games for agent dialogues; logical approach to coalitions in multiagent systems; reconsideration of agent intentions formalized as deliberation scheduling; and an interesting inverse game-theoretic approach to distributed system design.

— Simon Parsons
University of Liverpool

— Piotr Gmytrasiewicz
*University of Texas at
Arlington*

Learning Grounded Representations

When we speak of grounded representations, we're saying more about the meanings of the representations than the representations themselves: We're saying these meanings have a pretty direct connection to aspects of the environments in which agents operate. When we speak of learning grounded representations, we're saying these meanings are acquired with little or no intervention from us: intrinsic meanings, acquired by and for agents, which sets us at odds with a basic paradigm in AI. Much AI is technology for solving problems that seem to require understanding—natural language, vision, planning—without going to the trouble of saying what tokens mean or with the least possible effort. Programs consume and produce tokens that mean something to us only; to the machine, meaning is extrinsic, or derivative. This approach has been successful for decades, so why challenge it? Why do we care whether tokens mean anything to machines? First, we might see an economic imperative. Most syntactic operations don't produce meaningful results, and the cost of building programs that do is prohibitive. Semantic babysitting—checking whether a program has given us a meaningful answer and debugging it if not—is expensive. Second, we might think the vision of AI—of intelligent machines—remains unachieved as long as we must play semantic babysitter to our programs. Third, we might want to build more autonomous robots and other programs. Autonomy implies a lack of babysitting, semantic and otherwise. Fourth, we might be interested in a

range of questions from the philosophy of mind about the nature of intentionality. For example, is meaning causally epiphenomenal? Why do we share meanings? Is there really such a thing as intrinsic meaning? One can develop and seek empirical support for philosophical positions by trying to show how tokens might mean something to machines. Fifth, we might want to ground the semantics of language in the physical interaction between agents and their environments, or we might want to explain the emergence of denoting representations in the first year of life, or we might want a theory of perceptual development. All these interests were voiced at the symposium, which featured talks on language evolution, acquisition, and semantics; the development of physical knowledge and ontologies; symbol grounding and anchoring; the logic of grounded representations; and representations, whether they must be grounded in the physical world (for example, robots, much favored by attendees) or whether information agents in cyberspace are also grounded.

— Paul R. Cohen
Tim Oates
University of Massachusetts

Model-Based Validation of Intelligence

The aim of this interdisciplinary symposium was to bring together researchers in software validation and intelligent systems to compare techniques, explore possible opportunities for cross-fertilization, and find answers to the fundamental questions related to the validation of intelligent software.

The symposium attracted over 40 participants. There were 19 technical presentations, which stimulated lively discussions. In addition, there were three invited talks and one panel discussion.

One of the highlights of this symposium was the three outstanding invited speakers. On Monday, Nicola Muscettola, from NASA Ames Research Center, gave a talk entitled "Deploying Robust Autonomous Systems: Lessons from the Remote

Agent.” Muscettola gave an overview of the evolution of the remote-agent autonomous controller, the first AI software to fly in space, and presented his vision of the design, development, and validation of future generations of this system. On Tuesday, Kim Larsen, from Aalborg University, gave a talk entitled “Real-Time Model Checking and Beyond.” Larsen gave an overview of UPPAAL, an advanced model-checking and -verification tool. Larsen also talked about ways his group has been extending model checking in directions that utilize AI technology and help to validate AI applications, in particular, in planning and scheduling. On Wednesday, Paolo Traverso, from IRST, Trento, gave a talk entitled “Model Checking and Planning” in which he gave a survey of the pioneering work of his group. Traverso presented NuSMV, a reimplementa-tion and extension of the SMV model checker, and MBP, a model-based planner that is based on model-checking techniques and implemented on top of NuSMV.

Another highlight of this symposium was the panel entitled “Silver Bullets for AI Softwolves,” chaired by Robert Goldman from Honeywell Laboratories. The panel took place over two sessions: (1) the warm-up session on Tuesday and (2) the main session on Wednesday. In the warm-up session, Goldman presented a review of the famed “No Silver Bullet” paper by Frederick Brooks. The panel witnessed experts in four validation methods (Cesar Munoz, David Dill, Rustan Leino, and Doron Drusinsky) defend their favorite silver bullets (respectively, theorem proving, model checking, static analysis, and run-time analysis), followed by a critical review by three “judges” (Michael Lowry, Nicola Muscettola, and Richard Waldinger). The general consensus that emerged from the panel discussion can best be summarized by using Munoz’s suggestion of a “Swiss-army knife” approach, in lieu of a silver bullet, with multiple blades (that is, complementary approaches) to be used, with a good dose of skill and tenacity, to wrestle softwolves into submission. Leino pointed out the need in the AI community for “sharpening” the blades, using an

example in static analysis to make his point.

From the feedback received from many participants, this symposium was successful. It appears that both the validation and AI communities are becoming aware of the similarities and differences in their techniques, and a fruitful cross-fertilization has been initiated, as exemplified by several technical presentations, as well as the lively discussion, that the symposium inspired. More information about the symposium can be found at ase.arc.nasa.gov/mvi.

The chairs would like to thank everyone who contributed to the success of this symposium. Special thanks go to Robert Morris and Klaus Havelund for their helpful feedback.

— Lina Khatib

Charles Pecheur

NASA Ames Research Center

Robotics and Education

This symposium brought together about 40 participants from a diverse set of institutions, including high school students, primary and secondary educators, and faculty from research universities and teaching colleges whose departmental affiliations included computer science, engineering, and education. It is clear that there is much interest in robotics in an unusually wide range of educational settings. Robots are a natural motivator for students: They’re fun and exciting. However, most of the participants in the Robotics and Education Symposium were drawn to robots, or continue with them, because of other pedagogic benefits seen by using an inherently multidisciplinary, hands-on, empirical tool in an educational setting.

Robots have recently—within the last 5 to 10 years—become both inexpensive and robust enough that they are increasingly used in classroom settings. Popular platforms include the handy board (handyboard.com) and the Lego Mindstorms kit, both priced in the hundreds of dollars and usable with a variety of software interfaces. We also heard about experiences with many other commercially available or custom-built hardware platforms and

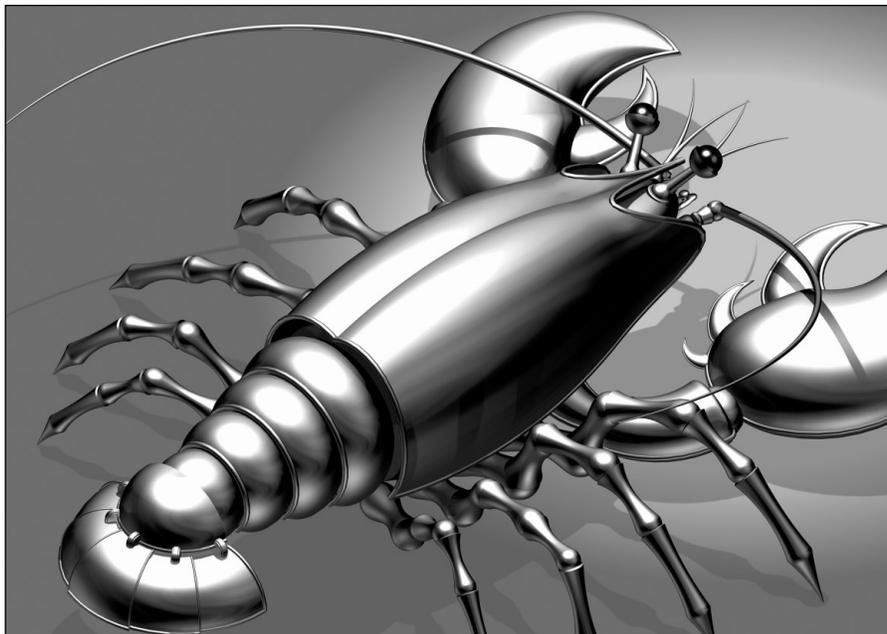
saw several new platforms not yet widely deployed.

Many of our participants used robotics in engineering settings, often as a multidisciplinary project lab. Although these courses covered material found in more traditional courses, the requirements of implementation and testing imposed by the physical robot also drew in a host of other skills—ranging from project management to pragmatic practice.

Introductory programming curricula are an increasingly common, if somewhat more unusual, place to find robots. There are many who argue that a robot gives a concrete, hands-on, pragmatic way of forcing novice programmers to confront software engineering challenges—unpredictability, changing requirements, need for maintenance (beyond initial development) of code—that substantially augment the introductory experience. These perspectives are not universally held, however, and there was discussion of the trade-offs between the potentially enhanced experience and the additional overhead imposed.

Not surprisingly, given the setting, several participants used robotics in the AI curriculum. These uses range from more traditional robotic applications to integration of robotics into a broad portion of the AI curriculum. However, not all the robotics applications were in intuitively technologically based areas. Other robotics-based courses included freshman writing seminars and introductions to engineering for humanities majors. We saw a program for Irish elementary schoolchildren who use robots to mount full-scale story-telling exhibitions and a soft and fuzzy robot intended to help elementary schoolchildren think about learning.

Perhaps the most exciting parts of the symposium concerned the many ways in which robots are being used as vehicles for outreach. Many of the college-level curricular programs also included involvement of, or outreach to, high school students. Other curricula were developed specifically for K to 12 education. Workshops and camps using robots are run to attract girls to science and technology. Other programs are targeted at underrepresent-



Biorobotics

Edited by Barbara Webb and Thomas R. Consi

Animal-like robots are playing an increasingly important role as a link between the worlds of biology and engineering. The new, multidisciplinary field of biorobotics provides tools for biologists studying animal behavior and testbeds for the study and evaluation of biological algorithms for potential engineering applications. This book focuses on the role of robots as tools for biologists.

An animal is profoundly affected by the many subtle and complex signals within its environment, and because the animal invariably disturbs its environment, it constantly creates a new set of stimuli. Biorobots are now enabling biologists to understand these complex animal-environment relationships. This book unites scientists from diverse disciplines who are using biorobots to probe animal behavior and brain function. The first section describes the sensory systems of biorobotic crickets, lobsters, and ants and the visual system of flies. The second section discusses robots with cockroach motor systems and the intriguing question of how the evolution of complex motor abilities could lead to the development of cognitive functions. The final section discusses higher brain function and neural modeling in mammalian and humanoid robots.

300 pp, index. ISBN 0-262-73141-X

\$35.00 / £23.95 softcover

A Fall Book from 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.*

ed minorities or at students from less socioeconomically or technologically privileged backgrounds.

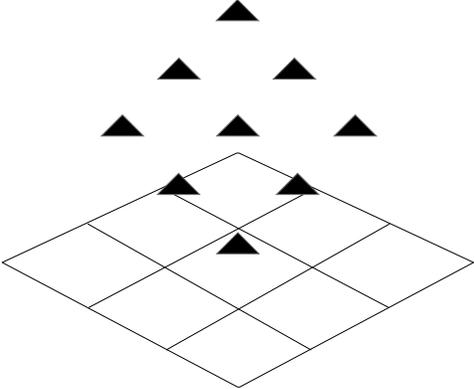
If you are inspired by the idea of using robots to help students learn, there are many opportunities to get involved. Any of the participants in this symposium will be happy to help you identify someone in your locale, or there are several national or international organizations that help start and support these programs. Many students encounter robotics through competitions, and teams are generally in need of mentors. Finally, if you are located in the greater San Jose, California, area, a Foothills High School teacher, who works with highly at risk students, has had remarkable success with a robotics-based curriculum and would love additional mentorship (www.thetech.org/robotics/activities/fhhs_intro.html).

— Lynn Andrea Stein
Olin College

Robust Autonomy

“Would you trust your autonomous system if your life depended on it?” When this question was posed to the Robust Autonomy Workshop participants, the response was a bit underwhelming—3 of the 40 participants raised their hands. When posed to the entire Spring Symposium community, only one additional person raised her hand. This lack of trust, even by the developers, provides the setting and the impetus to address the challenges in designing and deploying autonomous systems that are required to continue operating in the presence of failures and unanticipated events. Invited speaker Richard Doyle (Jet Propulsion Laboratory) kicked off the workshop with an inspiring discussion on future National Aeronautics and Space Administration (NASA) missions that will demand robustness—missions such as a Europa hydrobot that will be required to perform its mission under Europa’s “seas” without the luxury of communication with Earth.

The technical presentation topics ranged from lessons learned from deployed systems to tools and architectures for developing robust systems to managing environmental uncer-



*Submissions for the 2002 Spring
Symposium Series are due October 5th.
For more information, see the Symposia website at
www.aaai.org/Symposia/Spring/2002/*

tainties. These presentations sparked many lively discussions: Because a robust system must be able to handle unanticipated events, is there a trade-off between autonomy and predictability? As autonomy is introduced into a system, are we introducing more software (which usually is accompanied with more bugs), thereby reducing robustness? What are the advantages of introducing autonomy from the beginning (answer: efficiency, capability) versus retrofitting an additional capability (answer: minimal modification to a working system).

Experience indicates that autonomous systems find difficulty with user acceptance. We as a community must overcome this roadblock. Proposed steps include (1) involving users in the design process; (2) relying more on demonstrations of the technology versus technical papers to solicit understanding and buy-in; (3) gaining user-pilot trust through long-term use; (4) opening up the “black box” by providing concise explanations of the system’s behavior; and (5) allowing for gradual autonomy integration, the level of which can be controlled by the user.

No workshop would be complete without a discussion of terminology. As expected, definitions of robustness were plentiful: the ability to produce an executable plan despite run-time

variations in state, resources, and activity duration; graceful degradation; comprehensive event or fault detection, identification, and response; and identification of operational bounds or ranges for the system. There is also a level of robustness provided by software engineering approaches, such as understanding all requirements at the beginning of system development versus developing autonomous capabilities and then reengineering them to meet the domain requirements. In general, all agreed that a helpful step forward would be the definition of a clear set of metrics to determine the level of robustness. Some suggested metrics are to (1) determine how much error-safety margin is accommodated by the system, (2) borrow existing metrics such as the hardware-based approach of identifying the mean time between failure of the system, and (3) develop a generic simulation to benchmark the robustness of autonomous systems. Would metrics such as these provide sufficient evidence to trust a system if your life depended on it?

—Lorraine Fesq
Massachusetts Institute of
Technology

—Ella Atkins
University of Maryland