AI Magazine Volume 10 Number 3 (1989) (© AAAI)

AAAI News

Executive Council Meeting Minutes

January 12, 1989
(via telephone communications)

Attendees:
Douglas Lenat, Reid Smith, Bob Engelmore, Elaine Rich, Bill Clancey, Hector Levesque, Geoff Hinton, Danny Bobrow, Raj Reddy, Kathy McKeown, Howard Shrobe, and Richard Duda

Membership

1. Membership statistics indicate an increasing trend toward a reduction in the number of members.
2. Suggestion: Increase membership rate to $40 for regular members, $20 for students. The council unanimously agreed to the rate change.
3. Other suggestions: Make available to the members an automatic renewal of their membership by Visa or Master Card.
4. Discussed establishing a lifetime membership

   a. Issues included questions about the possibility of a net loss over time from establishing a lifetime status.

   b. Recommendation: Call American Association for the Advancement of Science and National Academy of Science to review the financial outcome of the type of membership category. Review the status again at the next council meeting.

National Technological University's Satellite Tutorials

1. After the collection of net discussions on AAAI sponsoring these tutorials, it was decided to proceed with an experimental series for 1989.

Carnegie-Mellon University's Project Mercury

1. CMU's staff has developed an Email interface that will soon be available to members. The CMU staff is working on the final version.

AAAI Press

1. To date, we have received proposals from Springer-Verlag, Morgan-Kaufmann, and The MIT Press. We are still waiting for proposals from Addison-Wesley and John Wiley. We expect to make a decision in February 1989.

Conference on Innovative Applications of Artificial Intelligence

1. We received about 100 submissions.

March 30, 1989 (Stanford University)

Attendees:
Bruce Buchanan, Daniel Bobrow, Bob Engelmore, Raj Reddy, Barbara Grosz, William J. Clancey, Hector Levesque, Elaine Rich, Howard Shrobe, Barbara Hayes-Roth, William Swartout, and Peter Patel-Schneider

Conference on Innovative Applications of Artificial Intelligence

1. There were 495 attendees for the first year's conference.
2. The conference was well received by the attendees, and it was generally felt that the applications represented a wide variety of domains and complicated decision-support functions.
3. Next year's conference will be held in Washington, D.C.

1989 Spring Symposium Series

1. The invitation-only policy will change to open enrollment.
2. It was observed that more discussion periods need to be introduced into each symposium.
3. The call for topics for 1990 will be distributed in late April.

The Grants and Scholarship Program

1. It was proposed we continue to support the PRIME and Mills College programs for disadvantaged students and women in 1989 for the same amount as 1988—$5,000 grant per institution.
2. It was recommended and approved that AAAI establish its own travel grant program for students attending IJCAI-89 rather than comingling funds with IJCAI.
3. It was recommended that AAAI support existing minority educational incentive programs rather than try to create its own.

CMU's Project Mercury

1. Project Mercury received additional funding from a number of other sources, including the Pugh Foundation; Digital Equipment Corporation; and, perhaps, Apple.
2. Daniel Bobrow, Bill Woods, Bruce Buchanan, and Robert Kahn, on behalf of AAAI, attended the project's status meeting at CMU on February 6. The group reviewed issues associated with the online distribution of information. These issues included mechanisms for retrieval, extent of the database, and liabilities and intellectual property rights. Project Mercury has hired a full-time librarian as well as significant programming staff.
3. Since January, the Email interface to bibliographic database's retrieval system has been undergoing beta testing. It will be available for use by members in May.
4. Everyone was reminded that one purpose of AAAI's funding of the project is to determine the amount of use the membership makes of this alternative form of information distribution. This experimental project is not a commercial development project. If use becomes overwhelming, then the database might be transferred to a commercial service.
5. Funding for 1990 was approved for $90,000.

Library Aid Program for Third World Countries

1. Many students in Third World countries cannot afford to buy textbooks. It was recommended that AAAI establish a Library Aid Program in which AAAI solicits book donations from publishers and acts as a clearinghouse for distribution. Libraries would submit applications to AAAI. A grant of $10,000 was approved for the first year.

AAAI-90

1. The issue of trying to allow articles more space in the proceedings for fuller disclosure of ideas was discussed. The conclusion by the group was to expand the number of allowable pages from 5 to 6.
2. Discussion about the proposed multisubmission paper policy by IJCAI was deferred until the August meeting.

Scientific Fraud

1. The group encouraged reciting the full set of data behind the experiment in dissertations or noting access to the data for further investigations.

Report on the AAAI Scholarship Committee

The AAAI Scholarship Committee has two objectives: (a) to encourage participation of computer science graduate students in AAAI sponsored conferences and workshops; and (b) to encourage female, minority, and underprivileged individuals to enter computer science and other science and engineering professions. Toward these objectives, the committee made the following awards during the past year:

- Student Travel Scholarships for AAAI 1989 Spring Symposia, $8500.
- Student Travel Scholarships for IJCAI89, $25,773.
- Grants to Outreach Organizations (discussed below), $20,000.

In its grants to outreach organizations, the committee made four $5000 grants. Two grants were renewals of grants initiated last year. The Math Science Network, in Berkeley, CA, aims to promote "the continuing development in mathematics and science of all people, with particular emphasis on the needs of women and girls." They used last year's AAAI grant to help support their 1988 Expanding Your Horizons conferences, which included 74 conferences in 21 states, serving over 20,000 girls. PRIME, Inc., in Philadelphia, PA, aims to "create opportunities for minorities and women in engineering, pharmacy, and other mathematics and science-based professions." They used last year's AAAI grant to expand their PRIME Universities Program, which allowed 543 students to spend the month of July, 1988 in skill-building programs at five university campuses in the Philadelphia area. Two new grants were initiated this year. Operation SMART of the Girls Clubs of America, Inc., in New York, NY, aims to "create equitable opportunities for girls, for youngsters from minority and low-income backgrounds, and for children with disabilities." The A.T. Anderson Memorial Scholarship Program of the American Indian Science and Engineering Society, in Boulder, CO, aims to "help talented American Indian science and engineering students meet the financial demands of going to college."

In addition to the grant programs listed above, the Scholarship Committee is interested in supporting a High School Mentor Program, under which female, minority, and underprivileged high school students could participate in activities at AAAI members' laboratories. These activities could occur during the summer months or after school during the school year. Students need not be involved in highly technical activities. The primary goal is to give students exposure to the field, encouragement to enter the field, and opportunities they might no otherwise encounter. AAAI members interested in participating in the Mentor Program can apply for high school student support by letter to the Scholarship Committee, AAAI, 445 Burgess Drive, Menlo Park, CA 94025.

Upcoming AAAI-Sponsored and Cosponsored Events

Conference on Machine Learning

The Seventh International Conference on Machine Learning will be held at the University of Texas in Austin, 21-23 June 1990. The conference will include presentations of refereed papers, invited talks and poster sessions.

The deadline for submitting papers is 1 February 1990. Papers are limited to 12 double-space pages (including figures and references), and should be formatted with 12 point font. Authors will be notified of acceptance by 20 March 1990 and camera-ready copy is due 23 April 1990. In addition to reporting advances in current areas of machine learning, authors are encouraged to report results on exploring novel learning tasks.

Please send papers (3 copies) to:

Machine Learning Conference,
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712-1188

For information, please contact:

Bruce Porter or Raymond Mooney,
m190@cs.utexas.edu
(512) 471-7316.

Nonmonotonic Workshop

The third international workshop on nonmonotonic reasoning will be held 31 May-3 June 1990 in South Lake Tahoe, California. Topics include general theories of defeasible inference, comparison of formal systems, relation to probability models, argument-based systems, applications to planning, commonsense reasoning, knowledge update and truth maintenance, theories of inheritance with exceptions, and proof theory complexity and automation.

Attendance will be limited to 30-40 people, by invitation only. Those wishing to attend should submit a detailed abstract of current research to:

Kurt Konolige
SRI International EJ272
333 Ravenswood Avenue
Menlo Park, California 94025
(415) 859-2788;
konolige@ai.sri.com.

Electronic mail submissions are encouraged. Abstracts should consist of no more than ten double-spaced pages when printed (4000 words).

Submission deadline is 17 December 1989. Notification of acceptance will be made by 26 February 1990. The preprint receipt deadline is 1 May 1990. Contact Kurt Konolige for further information.

Third International Conference on Industrial & Engineering Applications of Artificial Intelligence and Expert Systems.

IEA/AIE 90 will be held 15-18 July 1990 at Mills House Hotel, Charleston, South Carolina. This conference continues its tradition of emphasizing

applications of AI and expert/knowledge-based systems to engineering and industrial problems. Also of interest are the AI technology and research supporting such applications. Topics of interest include but are not limited to pattern recognition, vision, sensor fusion, computer-aided manufacturing, computer-aided design, robotics, planning/scheduling, diagnostics systems, intelligent interfaces, intelligent databases, autonomous systems, knowledge representation, knowledge acquisition, machine learning, natural language processing, neural networks, intelligent tutoring, reasoning under uncertainty, distributed and parallel architectures, qualitative models, blackboard systems, and industrial expert systems.

Please submit by 1 December 1989 four copies of an extended abstract (four to six double-spaced pages) to Manton M. Matthews, Department of Computer Science, University of South Carolina, Columbia, South Carolina 29208 (803) 777-3285; matthews@cs.scarolina,edu. Authors will be notified by 1 February 1990, and final copies of papers will be due 1 April 1990.

The 1989 AAAI Spring Symposium Reports

AI and Limited Rationality

This symposium brought together more than 50 scientists and philosophers interested in a core AI problem —how can an intelligent, computational agent act rationally when it has insufficient resources of time and information to determine the best course of action? It has long been clear that implementations of classical prescriptive theories of rational problem solving or decision making generally require intractable computations. However, until recently, few theorists have heeded Herb Simon's early warnings that a proper solution requires more than simple tuning of extant prescriptive theories. Practical AI has been forced to resort to ad hoc techniques, particularly for achieving real-time performance and managing uncertainty. The many, thoughtful papers submitted to this symposium made it clear that the consequences of taking resource limitations seriously would be far reaching, affecting the way one thinks about everything from simple search programs to the concepts of knowledge and reasoning.

Twenty-three papers were presented, including contributions from AI, decision science, philosophy, and economics. Proposed approaches to limited rationality can roughly be divided into two main camps, the conservative and the radical, although split personalities abounded. Many in both camps suggested that the goal of AI should be to achieve bounded optimality—the best possible configuration of a finite system to maximize a given performance measure within constraints imposed by a given class of environments. The camps divide over how to achieve this goal. We crudely characterize the two camps as follows:

The conservative agenda was to preserve the notion of deliberation on declarative representations and achieve bounded optimality through approximation methods, metalevel control, defeasible action selection, and design- and run-time compilation.

The radical agenda was to abandon the notion of general, rational deliberative architectures; study efficient, restricted architectures and analyze and extend their range of applicability.

One avenue for synergy lies in using the efficient architectures of the radicals as targets for compilation in the conservative systems. A particularly interesting common thread in the radical schemes was the use of locality in the structure of the agent to constrain computation. Concrete results and working application systems of both flavors were described. A tentative continuum between conservative and radical architectural configurations began to be perceived. Open problems identified and discussed included the potential infinite regress of deliberative control, the adequacy of decision theory as a conceptual basis for theories of limited rationality, the compilation of complex decisions for control of reasoning and action, learning in reactive architectures, and the integration of deliberative and reactive architectures. In general, this symposium made it clear that limited rationality is an important and timely topic for AI and related disciplines.

—Michael Fehling and Stuart Russell

Knowledge System Development Tools and Languages

This symposium was based on the premise that if a fundamental goal of AI is to build intelligent artifacts, then the production of effective tools and languages for building such artifacts is both a primary means of achieving this goal and a leading indicator of progress in the field. The symposium addressed current issues relating to the incorporation of emerging acquisition, representation, reasoning, explanation, and user-interface technologies into tools and languages for developing knowledge-based systems.

The symposium included the "Author's Self-Critique Panel," chaired by Richard Fikes, in which the panelists reflected on the strengths and weaknesses of the tool systems they had been involved in developing; a panel entitled "Use and Abuse of the Common Lisp Object System Metaobject Protocol (CLOS MOP)," chaired by Danny Bobrow, in which the panelists considered the role of CLOS and MOP in a knowledge representation facility; and sessions focused on large knowledge bases, use of commonsense knowledge in building knowledge-based systems, new developments in knowledge representation tools, task specific tools, and knowledge-acquisition tools.

Summary observations from the presentations and discussion include the following:

First, tool designers consistently have difficulty anticipating what is important to tool users. In particular, a pervasive tension exists between principled representation and hacker escapes in which, for example, tools are designed to support representation and reasoning but are used as object-oriented programming environments.

Second, tools for system engineers (for example, CLOS MOP, Joshua, Impulse) can be as important in facilitating system development as the standard representation and reasoning tools for knowledge engineers.

Third, many tool builders share the vision of a tool environment consisting of a collection of problem-solving architectures, each of which when used on a given task indicates the knowledge that needs to be acquired about the task and provides methods for solving the problems associated with the task. Developing such problem solving architectures or even determining what level of generality would be useful for them, however, remains a basically open problem.

—Richard Fikes

Representation and Compilation in High-Performance Theorem Proving

Progress in automated theorem proving primarily depends on finding better theorem-proving methods, but finding the best implementation techniques is also vital. This symposium focused on the latter. The techniques discussed can yield an orders-of-magnitude improvement in performance, although not uniformly across all problems and applications.

Compilation has been used to implement Prolog-like theorem-proving systems that are logically sound and complete (unlike Prolog). Inference rates are high but have less control of search-space redundancy and diminishing effectiveness for deep proofs. Caching results can be feasible and allow some trade-off between inference rate and search-space size. Compilation has also been applied to other areas, such as forward chaining, truth maintenance, term rewriting, and connection graphs.

The indexing of terms is necessary in many systems to efficiently retrieve all terms that unify with, are instances of, or are generalizations of a term. Long present in the most capable theorem provers, indexing has recently received increased scrutiny, with methods such as FPA/path indexing, discrimination nets, and connection graphs being examined and enhanced.

Sorted logics are an effective approach for using taxonomic information in theorem provers. Efficient techniques for sort reasoning have been developed. How to divide the theorem-proving task between multiple reasoners or whether to simulate specialized reasoners by control strategies in a general reasoner are important issues.

Parallel-processing versions of conventional and compiled systems have demonstrated near-linear speedup on tens of shared- and nonshared-memory SIMD processors.

—Mark Stickel and Woody Bledsoe

Planning and Search Summary

The program committee for the planning and search symposium consisted of Matt Ginsberg of Stanford University, Rich Korf from the University of California at Los Angeles, Vipin Kumar of the University of Texas at Austin, and Dave Smith from Rockwell International Corporation. The goals for the workshop were to explore the similarities and differences between work in planning and search in the hope of achieving some cross fertilization of the two fields. Similarities include the simulation of sequences of actions, the interleaving of planning and execution, interaction with other agents, limited computation, and coping with uncertainty. Differences include dealing with open versus closed worlds, the ease of problem representation, the amount of search brought to bear, the types of uncertainty encountered, and the sources of knowledge used to reduce search. The presented papers could generally be characterized as falling into one of eight categories: representing planning and search problems, interleaving planning and execution, learning plans, debugging plans, chunking in chess, using subgoals to reduce search, learning heuristic evaluation functions, and developing parallel search algorithms.

—Richard E. Korf

Robot Navigation

Navigation is the process of getting from one place to another in the service of some goal. The problem is made difficult by the need to deal with incomplete information and the geometric reasoning that is required. Until recently, most work on robot navigation has principally approached the problem in terms of a combinatorial path-planning task. More recently, there has been a growing recognition that the integration of a variety of problem-solving techniques is required. Sensor fusion, uncertainty, representation of spatial, physical, and functional knowledge, situated activity, planning and dynamic replanning, as well as nonmonotonic reasoning, are all applicable to the navigation problem. Mobile robot navigation, thus, provides a near-ideal environment within which to conduct research in system-level AI.

Several hotly debated topics dominated much of the discussion. Two of the more interesting discussions had to do with issues involving architectures and spatial representations. The classical dependence on extensive planning was attacked as being impractical given the uncertainties involved and, in any case, leading to systems in which few resources are left to actually take useful action. Reactive control, proposed as an alternative to planning, suffers from difficulties in the explicit representation of goal and in learning. With respect to spatial information, conventional wisdom in mobile robotics is that perceptual systems are used to determine geometric properties of the environment. This geometry is then used to infer topological properties (for example, junctions and paths) relevant to navigation and control. A rather radical counterview argues that perception and control should be tightly coupled so as to directly allow the determination of topological properties. Geometric properties are then inferred from the topological representations only as necessary.

—William B. Thompson

Artificial Intelligence and Software Engineering

The presentations and discussions at this symposium underscored a shift that has recently occurred in the field. Traditional work emphasized algorithm design; however, the emphasis has changed in the last few years to other aspects of software engineering, including the use of domain knowledge in requirements analysis and specification, techniques for modeling design and implementation decisions, and explicit representations of software development processes. This shift has occurred in part because of an increased amount of research in the context of real-world software (for example, the work of Biggerstaff at MCC, Ginoux at Electricité de France, Soni at Siemens, and Barstow at Schlumberger) and because of attempts at commercialization (for example, Reasoning Systems Inc.).

Although the symposium primarily focused on these trends, reports were also given on significant progress that has been made in the more traditional areas. Especially noteworthy are the KIDS system, which integrates recent work in algorithm design and transformational implementation (Smith at Kestrel Institute) and advances in program analysis (for example, Biggerstaff at MCC, Letovsky at Yale, and Wills at MIT).

The discussions at the symposium were exciting and stimulating and showed progress in the application of AI techniques to software engineering. Clearly, however, there is a long road ahead of us, especially as we begin to address issues that are important to programming in the large, such as support for collaborative work and representations of design histories.

—David Barstow