AI Magazine Volume 16 Number 3 (1995) (© AAAI)

The 1995 AAAI Spring Symposia Reports

■ The American Association for Artificial Intelligence held its 1995 Spring Symposium Series on March 27 to 29 at Stanford University. This article contains summaries of the nine symposia that were conducted: (1) Empirical Methods in Discourse Interpretation and Generation; (2) Extending Theories of Action: Formal Theory and Practical Applications; (3) Information Gathering from Heterogeneous, Distributed Environments; (4) Integrated Planning Applications; (5) Interactive Story Systems: Plot and Character; (6) Lessons Learned from Implemented Software Architectures for Physical Agents; (7) Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity; (8) Representing Mental States and Mechanisms; and (9) Systematic Methods of Scientific Discovery.

The American Association for Artificial Intelligence held its 1995 Spring Symposium Series on March 27 to 29 at Stanford University. This article contains summaries of the nine symposia that were conducted: (1) Empirical Methods in Discourse Interpretation and Generation; (2) Extending Theories of Action: Formal Theory and Practical Applications; (3) Information Gathering from Heterogeneous, Distributed Environments; (4) Integrated Planning Applications; (5) Interactive Story Systems: Plot and Character; (6) Lessons Learned from Implemented Software Architectures for Physical Agents; (7) Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity; (8) Representing Mental States and Mechanisms; and (9) Systematic Methods of Scientific Discovery. Proceedings of most of the symposia are available as technical reports from AAAI.

Empirical Methods in Discourse Interpretation and Generation

The symposium entitled Empirical Methods in Discourse Interpretation and Generation had two goals: (1) to investigate the empirical methods that can be used in the development and evaluation of computational theories of discourse and (2) to develop a set of shared resources for the computational discourse community. The workshop was successful at achieving the first goal because of the large number of papers submitted (48) surveying a range of different empirical methods. Methods discussed included Wizard of Oz studies, machine learning on corpora tagged for discourse features, human subjective evaluation of the system output in a natural language–generation system using a grading scheme, simulation test beds, the running of the implemented systems with different discourse modules on subjects and the evaluation of the quality of the interaction by various methods, and the coding of independent and dependent variables and the statistical calculation of their relationships. The discussions were lively and addressed issues such as determining what methods to use for various problems, suiting statistical tests to different types of tagging scheme, and determining the generalizability of different methods. The discussions also focused on the second goal, but progress on this goal consisted of airing the issues and forming a committee to get funding for a follow-on workshop to focus on developing shared resources.

Marilyn Walker
Mitsubishi Electric Research Labs

Extending Theories of Actions: Formal Theory and Practical Applications

Reasoning about action has been studied within a number of different subdisciplines of AI (and other fields), including knowledge representation, planning, reasoning under uncertainty, control theory, and decision theory. Work in these areas has often been motivated by different issues, problems, and underlying assumptions.

The aim of the symposium entitled Extending Theories of Action: Formal Theory and Practical Applications was to bring together researchers from these different areas to assess the state of the art, discuss common foundations of action representation and reasoning, and explore future directions for the extension and successful application of action theories. The symposium was well attended, attracting a diverse group of over 70 participants. There were five main panels: (1) Representation Languages, (2) Causality, (3) Observations, (4) Applications, and (5) Computational Issues.

The need for extended representations emerged as a main theme: For realistic applications, we must be able to handle concurrency, time, belief revision, complex actions, ability, utilities and preference, and so on. In addition, the need for new computational methods, especially approximation techniques, was discussed: The use of abstraction, hierarchical actions, relevance reasoning, and problem decomposition (especially in decision theoretic planning) was emphasized.

Because of the diverse approaches adopted by the participants, a tremendous amount of time was devoted to understanding the representations and issues addressed by different communities (for example, situation calculus, Bayes nets). Questions such as, What is the frame problem in a Bayes network? were common. Despite different models, it became clear that similar problems arise in all communities. Most participants went away with a better understanding and appreciation of the models and assumptions of other camps, and there is a definite sense that people are now eager to pursue these connections. The time seems ripe for substantive influence between disciplines.

Craig Boutilier
University of British Columbia

Information Gathering from Heterogeneous Distributed Environments

Given the rapidly increasing amount of information available online, locating, accessing, and integrating information from distributed and heterogeneous sources become critical problems. This symposium was aimed at bringing together researchers in AI who are working on the various aspects of this problem. The workshop attracted 75 people from a variety of areas, including knowledge representation, planning, learning, distributed AI, relevance reasoning, abstractions and approximations, case-based reasoning, and genetic algorithms as well as heterogeneous databases, operating systems, and networking.

Terry Winograd started off the workshop with the keynote speech, entitled "Information Gathering, for What?" He described how different types of information-gathering tasks will require different techniques. The discussion at the workshop focused on integrating multiple sources of related information, representing the contents of sources, efficiently processing queries to such sources, locating relevant sources, learning the structure of new information sources, and learning information to assist users in finding information.

There was little doubt that the technology discussed at the workshop could be useful for a wide range of applications, but a fair amount of discussion centered on which applications of information gathering would be particularly viable. There was an interesting discussion about how the information revolution would change the world in 20 years and what role AI researchers could play in the revolution. There was also discussion about which AI technologies are relevant to the problems that arise in information gathering. Knowledge representation and machine learning appear to be particularly well suited to this area.

Craig A. Knoblock
USC Information Sciences Institute

Alon Y. Levy
AT&T Bell Laboratories

Integrated Planning Applications

The planning field has transcended blocks world and has begun to take on real applications. Once a planning application becomes realistic, issues arise regarding the integration of the planner with its environment (that is, its application domain, the other software in the system, and the users). Because real applications are relatively new, organizers expected that planning practitioners would be at a stage at which they could benefit from one another's navigation of the pitfalls of development.

The range of presented applications was impressive: image processing, manufacturing, robotics, agents, oil spill cleanup, and space shuttle mission analysis. Similarly, the systems in which the planners were embedded included a variety of components: image-processing systems, simulators, sensors, real-time controllers, knowledge-acquisition systems, and user interfaces.

Participants generally agreed that the two primary open issues were user acceptance and acquiring and maintaining domain knowledge. Approaches from developing tools for environment support (for example, knowledge acquisition and debugging) to involving users from the beginning of the design process were discussed.

What was most surprising about the discussion was that planning as such was not an issue. Increasingly, the planning method was not as important as the system in which it was embedded and the tools available to support it. Thus, the important issues in integrated planning systems seem to be the same as those for any large software system, with the planner viewed simply as a high-level programming language. Participants agreed that an excellent follow-up to this symposium would be one that focused on the development of planning environment tools.

… the important issues in integrated planning systems seem to be the same as those for any large software system, with the planner viewed simply as a high-level programming language.

Adele Howe
Colorado State University

Interactive Story Systems: Plot and Character

Research in interactive story systems seeks to create a new computer-based art form providing experiences that are both meaningfully interactive and good stories. The user should have a lot of control over what happens, but whatever happens, the story should be coherent and interesting and have the temporal structure of a story (for example, rising tension followed by resolution). It has proven hard to achieve these two goals simultaneously.

Current interactive story systems can be divided into two classes: (1) story graphs and (2) simulated worlds. In a *story graph,* the user follows links from one predefined scene to another. In a *simulated world,* the user interacts with computer-simulated characters in a virtual world. Unfortunately, story graphs are only minimally interactive, and interacting with a simulated world seldom leads to anything that could be called a story.

The primary focus of current interactive story research is on adding plot control to simulated worlds. The goal is for a central plot controller to shape a user's interaction with a simulated world into a story by controlling what the simulated characters do and influencing what the user does (for example, by controlling when various pieces of information are presented). Progress has been made toward methods for controlling characters and influencing the user; however, no one has yet demonstrated effective plot control. If interactive story systems are eventually successful, the way that plot control is achieved will be the defining feature of the resulting art form.

Richard C. Waters
Mitsubishi Electric Research Labs

Lessons Learned from Implemented Software Architectures for Physical Agents

A software architecture for physical agents reflects the organizing principles that its designers have learned from many prior experiences in building such agents. In this symposium, organizers gathered researchers who have a great deal of experience in building architectures for physical agents and asked them to identify the lessons they have learned. The goal was to focus on the design choices faced by researchers building physical agents and the ways that these design choices affect the agent.

In his keynote address, James Albus proposed a reference architecture for physical agents called RCS. RCS is a hierarchical architecture with many different representations and modules. Albus argued strongly for this kind of reference architecture, but others felt that the community would settle on a standard naturally over time. Indeed, by the end of the symposium, a consensus seemed to emerge that architectures have at least three layers: (1) a top layer for symbolic planning or modeling; (2) a bottom layer for real-time, reactive control; and a (3) middle layer to serve as a differential between long-range reasoning and short-range reaction.

There were strong calls from participants to stop simply talking about architectures and start building robots. In particular, David Miller argued for building *niche robots,* which are designed to perform particular tasks well. Another theme of the symposium was evaluating and comparing the performance of different architectures. Erann Gat proposed a scientific methodology using simulated robots and environments as theoretical models to make predictions. These predictions then need to be verified or refuted by experimentation with real robots in real environments.

A large part of the symposium consisted of break-out sessions for the subareas of coordination, interfaces, representation, structure, performance, psychology, simulation, and learning. These sessions allowed for wide-ranging discussions, and participants found there is still much work to be done despite an emerging consensus on an overall architecture.

Henry Hexmoor
State University of New York at Buffalo

Representation and Acquisition of Lexical Knowledge: Polysemy, Ambiguity, and Generativity

Gather together a multidisciplinary group of impassioned and dedicated researchers who are at the forefront of lexical research and applications and determined to crack some of the core problems in the lexicon, and the result is predictable: a set of stimulating talks, question sessions, break-out sessions, and informal conversation. The lexicon is at the core of many natural language, information-retrieval, and knowledge representation systems and, thus, plays a central role in determining the success or failure of the endeavor. Nonetheless, points of sharp controversy have arisen concerning the most flexible and powerful way to represent the extensive variety of lexical information required. Despite the fact that the need for a common lexicon has been a recent research focus, there is no general agreement on a lexical representation that is (1) complete enough for specification of even basic syntactic environments and (2) flexible enough to handle the productivity and underspecification that are required for representing lexical semantics. The focus of the symposium was on ways to determine the optimal representation of lexical knowledge needed for flexible broad-coverage lexicons as well as on the acquisition of this knowledge. Papers were presented from the perspectives of philosophy,

If interactive story systems are eventually successful, the way that plot control is achieved will be the defining feature of the resulting art form.

psychology, human-language acquisition, machine-language acquisition, linguistics, system design, statistical modeling, and lexicology.

Two special events were planned for the workshop: One was an invited talk by Charles Fillmore on the notion of perspective in lexicography and lexicology, with comments on the future of lexical research. The other was a set of break-out sessions involving workshop participants, focusing on burning issues on lexicons for machine translation, information extraction, lexicons and linguistic theory, and incompleteness. Stimulating questions were discussed, such as, What is the most important work in this area that hasn't been done yet? Also, if you could design a three- to five-year project with no worry about budget, what would it look like? Break-out sessions reported on the questions driving participants' research and concluded that there are a generous number of hard problems to solve. The final presentation, by George Miller on WORDNET, concerned some cases where resolving sense ambiguity in corpora might not be desirable. This talk concluded the symposium with a promising note of relief.

Judith Klavans
Columbia University

Representing Mental States and Mechanisms

A system's ability to act and learn without supervision depends on how well it can reason about intelligent agents in its environment. In particular, programs intended to function autonomously in an agent-rich environment, such as an office or a classroom, require a great deal of general knowledge about mental states, mental capacities, and decision-making mechanisms. The symposium facilitated efforts to identify, organize, and represent such knowledge.

As suggested by the diverse backgrounds of symposium participants, the need to represent mental states and mechanisms transcends usual disciplinary boundaries. Many participants were interested in applications involving cooperation and communication between machine agents. Others were concerned with planning or learning tasks that require extensive self-knowledge and introspective capacity. A third group of researchers focused on tutoring, advising, and other tasks that involve reasoning about human mental states and capacities. Participants discovered considerable overlap in the knowledge requirements for otherwise different applications.

The symposium schedule balanced traditional paper sessions with an equal number of interactive sessions. John McCarthy opened the symposium with a talk and open discussion on machine introspection. Other interactive sessions included a discussion on goal-driven memory retrieval led by David Leake; a knowledge-sharing panel with Richard Fikes, Tom Gruber, and Pat Hayes; and a break-out session on the construction of naive psychological theories.

Participants learned about knowledge engineering efforts carried out by other researchers and about tools for sharing representations. Ultimately, it is hoped, the symposium will stimulate advances in our ability to construct programs capable of sophisticated reasoning about intelligent agents.

Michael Freed
NASA Ames Research Center

Michael Cox
Georgia Institute of Technology

Systematic Methods of Scientific Discovery

As suggested by the carefully chosen title, this symposium focused on systematic methods to carry out the creative reasoning in science that leads to discovery. The phrase *scientific discovery* and the call for papers made clear that organizers sought a strong connection to actual science: The abstraction of all science content into domain-free algorithms was not the norm, nor was discovery in business databases within the scope. The word *methods*, in the plural, indicated that we did not demand candidates for a single, comprehensive method for scientific discovery; specific techniques of limited generality within science were adequate. Finally, the word *systematic* might seem superfluous; what could be more systematic than a program running on a computer? However, the emphasis on science as a phenomenon of intrinsic interest (apart from being a mere outlet for AI or other techniques) led organizers to welcome methods that were more or less systematic but intended for execution on human processors rather than computer processors. Thus, scope was given to philosophical, psychological, and methodological work that did not contribute computerized methods but did address the central question of enhancing science through the use of more systematic methods. However, such work has tended to lead to computerized methods as well.

The symposium included a variety of technical presentations and posters on the themes of discovery systems at work, patterns of scientific discovery, lessons from history, and exploratory methods and systems. One panel session addressed the question, What does field *X* contribute to systematic methods of scientific discovery? Here, *X* included representatives from philosophy of science, psychology, machine learning, AI, and scientific practice. The second and closing panel addressed future scenarios for human-machine collaboration in science and the steps that are needed to reach these scenarios. Many nontechnical issues of a sociological nature arose here.

Two open-floor sessions were held. The first provided a forum for quick successive expositions of published discoveries in science, during which any discussion of the methods involved were suppressed. The second addressed promising new areas for machine discovery. An unexpected but welcome local participant made the case for striving to "blow the lid off" some problem in science with machine discovery, making an analogy with the early history of expert system work.

Raul Valdes-Perez
Carnegie Mellon University