AI Magazine Volume 15 Number 1 (1994) (© AAAI)

AAAI 1993 Fall Symposium Reports

■ The American Association for Artificial Intelligence held its 1993 Fall Symposium Series on October 22–24 in Raleigh, North Carolina. This article contains summaries of the five symposia that were conducted: Automated Deduction in Nonstandard Logics; Games: Planning and Learning; Human-Computer Collaboration: Reconciling Theory, Synthesizing Practice; Instantiating Intelligent Agents; and Machine Learning and Computer Vision: What, Why, and How?

The American Association for Artificial Intelligence held its 1993 Fall Symposium Series on October 22–24 in Raleigh, North Carolina. This article contains summaries of the five symposia that were conducted: Automated Deduction in Nonstandard Logics; Games: Planning and Learning; Human-Computer Collaboration: Reconciling Theory, Synthesizing Practice; Instantiating Intelligent Agents; and Machine Learning and Computer Vision: What, Why, and How?

Automated Deduction in Nonstandard Logics

Over the past decade, a wide variety of methods have been developed for automating deduction, with an even wider variety of nonstandard logics. The goals of the Automated Deduction in Nonstandard Logics Symposium were to bring together researchers working in this area with the aim of promoting comparisons of the various deduction methods that have been proposed, foster discussion of the different problems involved in automating the various logics, and obtain reports on the use of theorem provers for nonstandard logics in AI systems.

The logics considered included modal logics, many-valued logics, autoepistemic logics, conditional logics, argument-based defeasible reasoning, probabilistic logics, default logic, transaction logic, defeasible Prolog, and temporal logics. Some of the applications were reasoning about actions, commonsense reasoning involved in the use of natural language, and the specification and verification of the properties of a robot.

The symposium began with an invited talk by Melvin Fitting entitled "A Many-Valued Modal Logic, with a Curious Tableau System." The rest of the symposium consisted of paper presentations, a tutorial, and three panel sessions. On the afternoon of the first day, there was a session consisting of 10 short talks. The audience was asked to submit ballots indicating their top three choices for talks to be heard in a longer form on the last day.

The final day began with the presentations by the winners from the series of short talks. Siani Baker spoke on automated deduction in arithmetic with the omega rule. An epistemic logic with quantification over names was presented by Andrew Haas. Jana Koehler spoke on a planning system that uses temporal logic. A method for reasoning with indefinite knowledge was presented by Thorne McCarty. Jeff Pelletier spoke on semantic tableau methods for modal logics.

In many nonstandard logics (unlike classical first-order logic), the set of conclusions warranted from a set of premises is not, in general, recursively enumerable. Some of the different solutions were to limit consideration to the propositional case in which the conclusions are recursively enumerable; use heuristic methods closely related to those used in mathematical induction; or, as in John Pollock's system OSCAR, allow the reasoner to retract answers that it earlier proposed.

The symposium fostered some interesting discussion on the relationship between tableau methods, resolution, and logic programming. Another theme that came up in discussions of a variety of logics was the relative merits of direct methods of deduction, as opposed to those that work by translating the problem into another logic, and the relative advantages of different translation methods. Hans-Jürgen Ohlbach presented some exciting new work on this topic.

Much of the research on automated deduction for nonstandard logics has been going on in Europe, as seen by the fact that half of the entries in the working notes for this symposium were from people at European institutions. Thus, this symposium was an important opportunity to meet and discuss this material in North America. We were pleased that so many came such a long way to attend this symposium.

Peter Jackson
Clarkson University

Richard Scherl
University of Toronto

Games: Planning and Learning

Computer game playing addresses the fundamental issues in AI: knowledge representation, search, planning, and learning. The Games: Planning and Learning Symposium brought together researchers and practitioners from all over the world. The intent was to look again at game playing in the context of AI and cognitive science and strike a balance between theoretical issues and implemented systems. The symposium generated unflagging enthusiasm in its participants. Among us, there were strongly differing theses but active interest in and tolerance for

Copyright © 1994, AAAI. 0738-4602-1994 / $2.00

every point of view. In addition, we had fun, watching ant teams compete, learning new Scrabble words, and chasing ghosts through mazes.

The most highly accomplished game-playing programs were all well represented along with reports on their most recent prowess. One approach relies primarily on large, perfectly accurate knowledge bases. Such programs aspire to perfect play, with evaluation functions that become oracles once they reach the stored results of retrograde analysis. Several small games have been solved this way, and checkers play is nearing the best human performance with this approach. A second approach relies primarily on extremely fast, deep search to apply a good evaluation function deep in the game tree. Such programs are supplemented by book openings and end-game procedures to reduce reliance on the evaluation function where possible. One chess program is now able to play as well as all but the top 250 humans in the world. A third approach relies on training a neural net with more playing experience than any program or person has ever endured. The resulting backgammon program is one of the best players—human or machine—in the world.

There is a substantial difference between the way humans play games and the classic computer paradigm. People rely on experience; intuition; and highly focused, limited calculation. The traditional game-playing program does deep, minimax, alphabeta search. The new search algorithms introduced at the symposium, such as best-first minimax, proof-number search, threat-space search, and Bayesian-based search, are intended to select nodes for expansion based on features of the game tree itself, without large amounts of domain-specific knowledge.

The learning advocates, however, believe search alone will never suffice for games with large trees. They believe that a program requires knowledge, knowledge learned from experience by a variety of methods. It appears that a game-learning program must devote some effort not only to playing well but also to playing so that its experiences support thorough and efficient learning. The substantial impact of perception on a human's ability to learn delivered an important reminder about the role of knowledge representation and bias in learning.

Many aspects of game playing were discussed. One of the most interesting periods came when we worked as a group to refine our definition and understanding of imperfect information games. Both the theory and implementations for imperfect information games are in their early stages and offer substantial challenges. The concept of a metagame tournament, where game-learning programs would compete at computer-generated games, received support, although the first such tournament must await the construction of appropriate competitors. Finally, the symposium began a mailing list for research in computer game playing. (To join, contact Michael Frank at games-request@medg.lcs.mit.edu.)

Throughout this meeting, the diversity of games and the general principles that affect them provided a remarkable breadth of reference. There are clearly two different approaches to computer game playing: a high-performance determination to play better than any human and a cognitively oriented exploration of learning and behavior. A third approach, based on developing a mathematical understanding of heuristics and game playing, was also proposed. The competition and cooperation among these approaches should continue to drive some exciting and significant research.

Susan Epstein
Hunter College
University of New York

Robert Levinson
University of California
at Santa Cruz

Human-Computer Collaboration: Reconciling Theory, Synthesizing Practice

Human-computer collaboration deals with the theory and practice of collaborative problem solving between people and computers. *Collaborative problem solving* is a process in which several agents work together to achieve shared goals. In the Human-Computer Collaboration: Reconciling Theory, Synthesizing Practice Symposium, we sought a deep understanding of collaborative problem solving between one human and one computational agent, an understanding that takes into account the unique characteristics of each type of agent.

Given that a human and a computer are to work together to achieve a task, three key issues must be addressed: (1) how to distribute responsibility for accomplishing the task between the human and the computer, (2) how to communicate effectively about the task and the state of the collaboration, and (3) how to provide the system with sufficient knowledge and processing capabilities to enable it to make significant contributions to the problem solving. The symposium was organized around these three topics.

A number of important issues surfaced during discussion and were addressed from varying perspectives. First, a number of methodological issues were raised, for example, what the designers of collaborative computer systems could learn from studies of human-human collaboration, on the one hand, and from computer simulations of collaborative behavior, on the other. The consensus was that findings from both these areas could be useful but that great care was necessary in transferring these findings to system design. How to evaluate formal models of collaboration and collaborative systems also was discussed.

Second, much discussion centered on the knowledge a system requires to collaborate. It was agreed that real collaboration requires significant amounts of domain knowledge and some reflexivity on the system's part; for example, a collaborative spreadsheet might need to understand the spreadsheet model of computation. Arguments also were made that various additional types of knowledge, for example, knowledge of the human's plans and goals or of com-

Now in its Third Printing!

KNOWLEDGE DISCOVERY IN DATABASES

Edited by Gregory Piatetsky-Shapiro and William J. Frawley

The explosive growth in the quantity of data stored in databases creates a need for techniques and tools for understanding and extracting knowledge from data. Addressed to computer scientists, MISS professionals, and those interested in machine learning and databases, *Knowledge Discovery in Databases* is the first book to bring together leading-edge research from around the world on discovery in databases. Encompassing many different approaches, its thirty chapters move from the general to the specific: from less reliance on domain knowledge to more; from less restricted forms of discovered knowledge to the more specialized; from general methods to domain-specific techniques.

542 pages, $39.50

ISBN 0-262-66070-9 PIAKP

AAAI Press

DISTRIBUTED by The MIT Press
Massachusetts Institute of Technology
Cambridge Massachusetts 02142

To order call toll free: (800) 356-0343 or (617) 625-8569.

MasterCard and VISA accepted.

municative resources available to the system, either were necessary for or improved system collaboration.

Third, another set of issues arose from considering the long-term use of collaborative systems. Suppose your goal is to master an application rather than simply to perform a particular task. In this case, system collaboration should be directed at helping you learn the application rather than solve the particular problem that you face, even at the price of short-term inefficiencies in problem solving. Further, if you will use an application over a period of time, its knowledge must evolve, allowing you to shape it into your personal knowledge system. This capability requires sophisticated tools to support end users in viewing and modifying system knowledge and might well require some sort of learning ability on the part of the system.

The symposium revealed an encouraging consensus concerning the fundamental issues in the field. Although there are a variety of approaches to these issues, in many cases, the approaches are complementary rather than competing. Highlighting these issues, examining assorted approaches, and exploring their relationships were positive outcomes of the symposium.

Loren Terveen
AT&T Bell Laboratories

Instantiating Intelligent Agents

The Instantiating Intelligent Agents Symposium was attended by approximately 40 researchers. Although the majority were robotics researchers, attendees represented subdisciplines of AI ranging from natural language processing to formal logic.

The purpose of the symposium was to discuss what it takes to build (and what is) an intelligent physical agent. To help direct the discussion, the symposium attendees were instructed to focus their attention on an intelligent agent for doing household vacuuming and related tasks. To help set the stage, the first speakers came from companies that actually make robotic vacuum cleaners (for office environments). They described some of their real-world experiences with interacting with robots, getting customers to interact with robots, designing safe robots with adequate sensing, establishing the minimal amounts of intelligence and autonomy necessary for a robot to be productive in an unengineered environment, and deciding what makes a good vacuum cleaner.

The next two days had talks, panels, and lively discussions on a variety of topics ranging from the dangers of reward-motivated vacuuming in houses with Christmas trees to the feeding habits of dust bunnies. Although the talks sometimes appeared to be heading off into the stratosphere, everyone at the symposium paid attention to the goal—an intelligent agent for household cleaning. Somehow, the discussions always managed to have a point.

Saturday night we broke into two groups. The first group's assignment was to come up with a set of design specifications for some AI vacuuming agents that could be created with current technology in one to three years. The other group's aim was to come up with some design criteria for ROBO-VAC 2000.

The groups reported their results Sunday morning. The first group came up with several designs for vacuum aids: agents with insect levels of intelligence that could clean areas or vacuum around baseboards. It is important to report that the participants generally agreed that general autonomous vacuuming in a home—without modifying the home, greatly lowering the performance expectations, or doing more research and development—was beyond the current state of the art in AI.

The second group laid out the capabilities and technologies for ROBO-VAC 2000. Although not quite Rosie the Robot, ROBO-VAC 2000 would be able to maintain multiple rooms in a house and take directions from its owner to do spot cleaning or modify its schedule. Everyone agreed that ROBO-VAC 2000 was well beyond the state of current technology, but they also agreed that all the necessary technology pieces did exist (or at least were claimed to exist) and had been demonstrated under laboratory conditions.

As a result of this design exercise, the symposium participants believed that ROBO-VAC 2000 could be the AI problem for the next decade. The AI community believes it has developed the pieces to make an intelligent agent. What we lack is the technology to integrate these pieces. Developing this integration technology and building a real-world agent with it should teach us a great deal about building agents, software integration, and the AI technologies that really work outside the laboratory. The participants left the symposium enthusiastic about the idea that a large integration AI project with a real goal that can be evaluated by users, rather than just technologists, is an important task for the community to take on. Weeks after the symposium, it still seems like a good idea.

R. Peter Bonasso
David Miller
MITRE Corp.

Machine Learning and Computer Vision: What, Why, and How?

Over 70 researchers attended the Machine Learning and Computer Vision: What, Why, and How Symposium, with strong attendance from both the machine learning and the computer vision communities. The symposium was divided into 90-minute sessions, some containing invited talks and panels but most devoted to moderator-author coverage of contributed papers. The moderator-author format for the contributed paper sessions proved interesting and valuable. Each moderator summarized and commented on five papers and then let the authors respond. Invited speakers included Tom Mitchell and Rich Sutton from machine learning and Chris Brown and Ramesh Jain from computer vision. Abe Waksman from the Air Force Office of Scientific Research organized the panel discussion.

The basic performance task of a vision system can always be viewed as mapping from sensory data (the input to the problem) to one or more of a set of possible decisions or actions (the output to the problem). Moreover, one can decompose the performance task of a traditional vision system into three subtasks that hold across the different categories. First, the system must transform its sensory input to a set of features, that is, early symbolic or qualitative abstractions of the input. In the second subtask, the vision system must go from inferred features to a set of partially instantiated entities that it can discern in the sensory input. This indexing or retrieval process can deal with models of specific objects, entire classes of objects, patterns of motion, contexts of scenes, or other phenomena. Finally, the vision system must go from these candidate models to decisions about the best models for the given sensory input, using some recognition process. Different systems can place different relative emphasis on these subtasks, but all systems must, in effect, deal with them in some manner.

Machine learning can help computer vision systems in a number of areas. A few examples are given here. Object models can be learned from examples or can be modified as examples are presented. Models can be constructed by observing objects in many different poses and abstracting the model. A system can learn which available features are most effective or learn composite features from the given lower-level features. Given a set of features, a system might learn to index from them to a label for an object.

The goal of improving the performance of computer vision systems presents a number of challenges to the field of machine learning:

Structured representations: Algorithms for machine learning are typically designed to operate with flat attribute-value formalisms. However, most research on computer vision assumes that knowledge about an image has inherent structure and, thus, represents information at multiple levels of aggregation.

Handling of uncertainty: Many learning algorithms represent acquired knowledge in logical terms that either match or mismatch a given instance, and even more assume that the features of instances are certain. However, many aspects of visual domains are inherently uncertain.

Partial information: Most work on machine learning assumes that all features are present during both training and testing. In contrast, images seldom contain all the information that would be useful in vision.

Focusing of attention: The performance components associated with most learning algorithms assume that information about instances falls outside the system's control and that no costs are involved in collecting such information. Recent work in computer vision assumes exactly the opposite, that focusing attention is central to the processes of visual inference and recognition.

Incremental learning: Typical machine-learning techniques process training instances in a nonincremental manner, using statistical regularities to direct search through the space of hypotheses. Although one can collect images for processing of this sort, a more natural approach attempts to learn incrementally from each image as one encounters it.

Learning with many classes: The majority of supervised induction techniques have been designed to handle only a few classes, and even unsupervised methods are seldom tested on domains with many different categories.

In summary, it is clear that computer vision and machine learning have much to contribute to each other. This symposium brought together a community of researchers that are excited about the great potential of the area, but it also revealed that the area has a long journey to travel before realizing this potential. Nevertheless, the symposium laid a good foundation for future work on this promising topic, and we hope future meetings will produce more significant results on the coupling of vision and learning.

Kevin Bowyer
Lawrence Hall
University of South Florida