AI Magazine Volume 20 Number 1 (1999) (© AAAI)

AI Game-Playing Techniques: Are They Useful for Anything Other Than Games?

A Synopsis of the Panel Discussion at IAAI-98

Dana S. Nau

■ In conjunction with the American Association for Artificial Intelligence's Hall of Champions exhibit, the Innovative Applications of Artificial Intelligence held a panel discussion entitled "AI Game-Playing Techniques: Are They Useful for Anything Other Than Games?" This article summarizes the panelists' comments about whether ideas and techniques from AI game playing are useful elsewhere and what kinds of game might be suitable as "challenge problems" for future research.

In conjunction with the American Association for Artificial Intelligence Conference (AAAI-98) Hall of Champions exhibit, the Innovative Applications of Artificial Intelligence held a panel discussion entitled "AI Game-Playing Techniques: Are They Useful for Anything Other Than Games?" We had the following panelists:

David Fotland is a hardware designer at Hewlett Packard and author of the Go program THE MANY FACES OF GO.

Dana Nau (the panel chair and cochair of the Hall of Champions exhibit) is an AI planning researcher at the University of Maryland and coauthor of the latest version of the BRIDGE BARON.

Jonathan Schaeffer (senior chair of AAAI-98's Hall of Champions exhibit) is an AI games researcher at the University of Alberta and author of the checkers program CHINOOK.

Gerald Tesauro is a neural network researcher at IBM Research and author of the backgammon program TD GAMMON.

David Wilkins is an AI planning researcher at SRI International and has also done research on computer chess.

The early research on the alpha-beta search algorithm was useful in establishing a foundation for AI theories of heuristic search, and these theories have been useful in many areas of AI.

Several of the panelists (particularly Schaeffer, Wilkins, and Fotland) pointed out that the minimax search algorithms traditionally associated with AI have only a limited range of applicability.

More specifically, Schaeffer expressed disappointment that game-playing program performance turned out to be so strongly correlated with search depth. This correlation led researchers to devote a disproportionate amount of work to achieving very small improvements in alpha-beta search efficiency, which has little applicability outside games.

Fotland pointed out that the search code is the smallest part of most AI game-playing programs. Most of the effort goes into domain-specific knowledge engineering; so, even if chess search algorithms are useful outside computer games, the vast bulk of effort in computer chess has to do with chess-specific knowledge that is not useful outside computer games.

Fotland also mentioned that the AI game algorithms are not very useful in general for computer games. The most popular computer strategy games (WARCRAFT, COMMAND AND CONQUER, CIVILIZATION, and so on) have abysmal AI. Despite these limitations, game playing has been useful as a test bed for developing several ideas and techniques that are useful elsewhere. The panelists gave several examples of how game-playing research has led to results useful elsewhere:

The early research on the alpha-beta search algorithm was useful in establishing a foundation for AI theories of heuristic search, and these theories have been useful in many areas of AI.

In his work on CHINOOK, Schaeffer had to develop new parallel algorithms for information storage and retrieval. Those algorithms are now finding use in DNA sequencing and other areas.

Tesauro used backgammon as a test bed for the development of learning algorithms that have inspired work by other researchers on several other problems. Examples include job-shop scheduling for the National Aeronautics and Space Administration space shuttle (Zhang and Dietterich at Oregon State University), elevator dispatch control (Crites and Barto at the University of Massachusetts), cellphone channel assignment (Singh at the University of Colorado), and assembly-line manufacturing (Moore at Carnegie Mellon University and Mahadevan at Michigan State University).

To adapt hierarchical task-network planning techniques for use in bridge, Stephen Smith (Hood College) and Nau had to develop ways for the planner to perform complex numeric calculations, plan for multiple agents, consult external information sources, and reason about uncertain information. These same techniques are now proving useful in tasks very different from bridge play, such as generating and evaluating manufacturing process plans.

One question from the audience was whether it might have been better to start out by developing these techniques directly for the other real-world applications rather than working first on computer games. Here are some of the panelists' responses:

The game-playing problems were what originally caught the researchers' imagination and inspired them to do the work in the first place. Also, games have a much higher visibility with the public (witness the enormous attention to the DEEP BLUE matches).

Real-world application domains are often so complex that one can spend huge amounts of time on details that have little to do with research. In the cases described earlier, game playing was a useful test bed because it provided a smaller and more tractable domain, which made it easier for the researchers to focus on the ideas they were trying to develop.

Much of the subsequent discussion focused on what games might be useful as challenge problems for future research, in particular, what kinds of game are likely to foster the development of techniques that will be useful and significant for other areas of AI. Several kinds of game were proposed:

First are the games of Go and Poker. Go is a perfect-information zero-sum game like Chess and Checkers but one in which traditional search techniques do not appear to be sufficient. Poker is an imperfect-information game in which it is important to be able to model the opponent's betting and bluffing strategies.

Second are classes of games in which the rules can vary from game to game (such as the metagames proposed a few years ago by Barney Pell). These would again be perfect-information zero-sum games, but the variability in rules would make it more important to develop general approaches rather than techniques that are hard wired to a specific game.

Third are some of the popular video war games that require coordination among multiple agents. Little is known about how to develop effective AI techniques for such games.

Fourth are computerized versions of athletic games, such as the software Robocup competitions. These games require the development of a very diverse set of skills, such as image recognition, real-time decision making and acting, motor activities, and coordination among multiple agents.

In summary, a majority of the panelists felt that game playing has been useful as a test bed for developing ideas and techniques that are useful elsewhere—but that several of the more useful ideas to come out of AI game-playing research have little to do with the minimax search algorithms traditionally associated with AI game playing.

Dana Nau is a professor at the University of Maryland in the Department of Computer Science and the Institute for Systems Research (ISR) and is director of the Computer-Integrated Manufacturing Laboratory. He received his Ph.D. from Duke University in 1979, where he was a National Science Foundation (NSF) graduate fellow. He has more than 200 technical publications (for recent papers, see www.cs.umd.edu/users/nau). He has received an NSF Presidential Young Investigator award, the ISR Outstanding Systems Engineering Faculty award, and several "best papers" awards. He is coauthor of the computer program that won the World Bridge Computer Challenge in July 1997, and he is a fellow of the American Association for Artificial Intelligence. His research interests include AI planning and searching and computer-integrated design and manufacturing.