

The Games Computers (and People) Play

Jonathan Schaeffer

Computing Science Dept.

University of Alberta

Edmonton, Alberta

Canada T6G 2H1

jonathan@cs.ualberta.ca

Abstract

The development of high-performance game-playing programs has been one of the major successes of artificial intelligence research. The results have been outstanding but, with one notable exception (Deep Blue), they have not been widely disseminated. This talk will discuss the past, present, and future of the development of games-playing programs. Case studies for backgammon, bridge, checkers, chess, go, hex, Othello, poker, and Scrabble will be used.

The research emphasis of the past has been on high performance (synonymous with brute-force search) for two-player perfect-information games. The research emphasis of the present encompasses multi-player imperfect/non-deterministic information games. And what of the future? There are some surprising changes of direction occurring that will result in games being more of an experimental testbed for mainstream AI research, with less emphasis on building world-championship-caliber programs.

One of the most profound contributions to mankind's knowledge has been made by the artificial intelligence (AI) research community: the realization that intelligence is not uniquely human.¹ Using computers, it is possible to achieve human-like behavior in nonhumans. In other words, the illusion of human intelligence can be created in a computer.

This idea has been vividly illustrated throughout the history of computer games research. Unlike most of the early work in AI, game researchers were interested in developing high-performance, real-time solutions to challenging problems. This led to an ends-justify-the-means attitude: the result—a strong chess program—was all that mattered, not the means by which it was achieved. In contrast, much of the mainstream AI work used simplified domains, while eschewing real-time performance objectives. This research typically used human intelligence as a model: one only had to emulate the human example to achieve intelligent behavior. The battle (and philosophical) lines were drawn.

The difference in philosophy can be easily illustrated. The human brain and the computer are different machines, each with its own sets of strengths and weaknesses. Humans are good at, for example, learning, reasoning by analogy, and

image processing. Computers are good at numeric calculations, repetitious computations, and memorizing large sets of data. These machine architectures are largely complimentary: the human's processing strengths are the computer's weaknesses and the computer's strengths are human weaknesses. Given a problem to be solved and a specified architecture (human brain or silicon computer), a good solution should cater to the strengths of the machine being used, not the weaknesses. When viewed in this light, it is not surprising that the unhuman-like approaches have won out.

Building high-performance game-playing programs has been one of AI's major triumphs. This is due, in part, to the success achieved in games such as backgammon, chess, checkers, Othello, and Scrabble, where computers are playing as well as or better than the best human players. However, it is also due to the examples it set to the research community. These include tackling challenging problems (rather than trivial subsets, as is still often seen in AI research) and the emphasis on the results of the system without regard for the methods used to achieve those results (the ends justify the means). More details can be found in (Schaeffer 1999).

The success of *Deep Blue* at chess was an enormous publicity win for the AI community. With the 50-year challenge of building a world-class chess program finally out of the way, no other games-related grand challenge problem has emerged that fires the imagination of AI researchers like chess did. A new generation of games are being actively researched (e.g. bridge, go, poker, shogi) which promise to produce research results that will likely have a wider impact on the AI community. The games cannot benefit significantly from the silver bullet of chess—deep search—and it will take a plethora of new ideas to achieve success. Research into building game programs continues, but the work will be less visible and more academic.

Acknowledgments

This work was funded by the Natural Sciences and Engineering Research Council of Canada. Thanks to Haym Hirsh for prompting me to put these ideas into words.

References

Schaeffer, J. 1999. The role of games in understanding computational intelligence. *IEEE Intelligent Systems* (November/December):10–11.