

Does Deep-Blue use AI?

Richard E. Korf

Computer Science Department
University of California, Los Angeles
Los Angeles, Ca. 90095
korf@cs.ucla.edu

Abstract

When Deep Blue played Garry Kasparov in February 1996 and May 1997, the extensive IBM web pages devoted to the site claimed that Deep Blue did not use artificial intelligence. We argue that this claim is inaccurate, is representative of a wide-spread phenomenon in the field, and is ultimately harmful to AI.

Does Deep-Blue use AI?

In February 1996 and May 1997, the reigning world chess champion, Garry Kasparov, played Deep Blue, a chess computer built by IBM and one of the strongest in the world (Keene, Jacobs, and Buzan, 1996). Kasparov won the first match 4-2. IBM sponsored extensive coverage of both matches including a web site with up-to-the-minute results, commentary, and background information. Surprisingly, there was almost no mention of artificial intelligence in any of the IBM web pages. The one exception was in a list of frequently asked questions, one of which was, "Does Deep Blue use artificial intelligence?" Even more surprisingly, IBM's answer to this question was "no"! (IBM, 1997)

To be fair, a careful reading of their answer to this question seems to equate artificial intelligence with the simulation of *human* intelligence. By that measure, Deep Blue doesn't use AI, since it plays chess very differently than a human does. For example, Deep Blue generates and evaluates about 200 million chess positions per second, something no human can do. On the other hand, it is well understood that AI is much broader than simply the simulation of human thought processes, and encompasses a great deal of techniques that are clearly not psychologically plausible.

By this broader definition, does Deep Blue use artificial intelligence? While there is relatively little published on Deep Blue, it is well-known that its main algorithm, at least for mid-game play, is alpha-beta minimax search with a heuristic static evaluation function. We argue that if any technique deserves to be called AI, this one does.

Chess was one of the original AI problems, and remains a canonical AI task. In fact, computer chess predates the term "artificial intelligence". The first paper

on the subject, written by Claude Shannon, and entitled "Programming a Computer To Play Chess", was published in 1950 (Shannon, 1950). In all the earliest books on AI, such as *Computers and Thought* (Feigenbaum and Feldman, 1963), chess and Samuel's checkers program (Samuel, 1963) were featured prominently.

One of the earliest AI techniques, heuristic search, was developed primarily to deal with problems such as computer chess. In particular, AI added to the minimax algorithm of (Von Neumann and Morgenstern, 1944) the ideas of heuristic static evaluation and alpha-beta pruning, all of which were developed in the 1950s. These ideas form the core of all chess programs even today, including Deep Blue.

Thus, clearly Deep Blue is a product of artificial intelligence, even if it isn't of the cognitive simulation variety. The fact that it won a game against Kasparov in 1996 should be viewed as an achievement for AI, and if it wins their rematch in May of 1997, it should be heralded as a glorious success for AI, tackling a problem that has been worked on hard and continuously for almost fifty years.

The Problem with AI

Why isn't Deep Blue viewed as an artificial intelligence program, at least by its sponsor, IBM? I believe that the term AI has become tainted in the field of computer science, and in the public mind as a whole. Partly this is a result of the field failing to deliver on its early promises. One of those promises was to program a computer that will defeat the human world champion in chess within ten years, a claim that was made almost continuously from about 1960 on. Another reason for disenchantment with the field is probably the name itself, which seems much more grandiose and boastful than other early alternatives, such as "complex information processing".

For whatever reason, AI is viewed with skepticism within computer science. Work in AI is often thought to be lacking in rigor, and without scientific basis. To some extent, this is a product of the nature of the problem being studied, but it also is the fault of those of us in the field who have not been sufficiently dili-

gent in applying rigorous scientific standards to our work. The Deep Blue project, and these highly visible matches, are viewed as marketing tools by IBM. It's no surprise that they would want to associate it with positive images, and avoid labels such as AI that have come to have a negative connotation in the public eye.

While IBM's mischaracterization of Deep Blue is a highly visible example of this phenomenon, it is by no means the only one. I believe that it has become commonplace for researchers and practitioners in areas that were once viewed as AI to distance themselves from the field, and particularly the term "artificial intelligence". Part of this is a result of simply the specialization of subfields of AI. Fields such as expert systems, computer vision, speech understanding, natural language processing, robotics, machine learning, planning, and knowledge representation, all of which started out as part of artificial intelligence, now have their own conferences and journals. Furthermore, many researchers in these subfields associate themselves with the subfields, and don't view themselves as AI researchers anymore.

To a large extent, the fact that artificial intelligence has spawned so many successful subfields that have reached critical mass is a success story for AI. The problem comes when the successful intellectual offspring of the field disown their parents. Every field has its successes and failures. If the successful subfields migrate away from the parent field, forming their own subfields, and the researchers in those fields don't consider themselves part of AI anymore, but associate with the subfield instead, the field of AI is left only with its failures, and newly emerging subfields, which have yet to demonstrate their promise. The result is a negative image of the AI field, both in the public eye and within computer science, an image under which it continues to suffer unfairly.

The best example of this phenomenon is the field of expert systems. Originally, expert systems was viewed as a subfield of artificial intelligence, but is now viewed as a separate field in itself, particularly by the people who build expert systems. Expert systems are by far the most important and widespread practical applications of artificial intelligence. See, for example, (Winston and Prendergast, 1984).

Interestingly, the IBM web pages give some credit to expert systems for the success of Deep Blue. What they cite as evidence of this is the program's opening book of moves. Clearly, this is an important knowledge base in the program, but it is simply a rote table of moves, rather than a set of rules. Furthermore, the opening book is only used for the first few moves of the game, after which the program must rely on its alpha-beta minimax search coupled with its heuristic static evaluation function. Modern high-performance game programs have little in common with expert systems. In spite of this, IBM gives credit to this technology, presumably because the term "expert systems" is

more palatable to the general public than the term "artificial intelligence", and better known than the term "heuristic search".

What to do About it?

Having identified this problem, the more difficult question is what to do about it. Changing the name of the field isn't a solution. The name is too well entrenched, having been around since the Dartmouth conference in 1956. In any case, the original connotations of the label eventually are lost and replaced by the connotations of the field itself. For better or worse, the image of the label "artificial intelligence" will be based primarily on the advances and ongoing research in the field that it refers to.

The only solution that comes to mind is a vigorous effort of education both within the field of computer science and among the public as a whole. This education should be aimed at making people aware of the successes of the field of AI and of the fields that have spun off from it. For example, we should reclaim computer chess as a classic AI problem. At some point, I believe that a chess machine will defeat the human world champion, and when that happens, it should be a public relations bonanza for the field of artificial intelligence, rather than some other field or organization.

Acknowledgements

This work was supported by NSF Grant IRI-9119825. Many thanks to Michael Dyer for his insightful comments on an earlier draft of this manuscript.

References

- E. Feigenbaum and J. Feldman (Eds.), *Computers and Thought*, McGraw-Hill, New York, 1963.
- International Business Machines, Deep Blue FAQ, <http://www.chess.ibm.com/meet/html/d.3.3a.html#ai>, 1997.
- Keene, R., B. Jacobs, and T. Buzan, *Man v Machine: The ACM Chess Challenge: Garry Kasparov v IBM's Deep Blue*, B.B. Enterprises, Sussex, 1996.
- Samuel, A.L., Some studies in machine learning using the game of checkers, in *Computers and Thought*, E. Feigenbaum and J. Feldman (Eds.), McGraw-Hill, New York, 1963, pp. 71-105.
- Shannon, C.E., Programming a computer for playing chess, *Philosophical Magazine*, Vol. 41, 1950, pp. 256-275.
- Von Neumann, J., and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, N.J., 1944
- Winston, P.H., and K.A. Prendergast, *The AI Business - Commercial Uses of Artificial Intelligence*, MIT Press, Cambridge, MA, 1984.