Blondie24, Playing at the Edge of AI

By David B. Fogel, Morgan Kaufmann Publishers, 2002, ISBN 1-55860-783-8

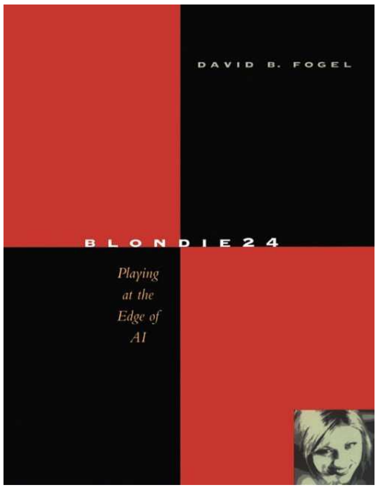

Reviewed by Eric Harley[1]

Computer science has failed abysmally at producing machines which display intelligence. According to Fogel, the last 50 years of effort in artificial intelligence have been on the wrong track, leaving us no closer to the goal than when we started. The wrong track has been the attempt to make the computer imitate our behavior. Scientists striving to build an artificial intelligence load the computer with knowledge on a chosen topic, along with an algorithm to do the associated task. The hope is that the computer will equal or surpass our intelligence for that subject. This approach dates back to the Turing Test, which Fogel points out has been misquoted and misinterpreted almost from day one. Misquoted or not, the Turing Test has a computer pretend to be human, and this paradigm became a signpost saying that the road to artificial intelligence is through mimicry of human behavior. According to Fogel, however, that path leads only to an illusion of intelligence — for example the kind of wooden intelligence exhibited by Deep Blue.

If imitation of human intelligence has not led to machine intelligence, then what will? What is intelligence? Fogel is critical in general of researchers in this field for skirting this last question. He believes that if the field of artificial intelligence had started with a proper definition of intelligence then there would have been a better chance of creating it. That is common sense — knowing what you are trying to build is crucial. Fogel therefore gives the following definition of intelligence.

> Intelligence is the capacity of a decision-making system to adapt its behavior to meet goals in a range of environments.

Fogel looks to nature for an example of intelligence. He describes the intricate and seemingly clever behavior of a certain species of wasp. However, he points out that an individual wasp is fixed in its behavior. An experiment by Jean Henri Fabré described in the Notes section of the book seems to make this clear. The wasp, Fogel says, is like the ‘proverbial robot’, an automaton with no adaptive behavior. Therefore the individual wasp is not intelligent. However, Fogel considers the species of wasp to be intelligent as a group or system, since it evolves to meet a changing environment. The species is what has learned the intricate behavior. The point is made that in general, intelligence requires a reservoir of knowledge and a mechanism to adjust that knowledge in the face of changes in the environment. The holder of the knowledge depends on the system — for a species it is the genetic pool, for a social group it is the culture, and for an individual it is the brain.

According to Fogel, a definition is neither right nor wrong, but rather useful or not useful. Perhaps that is true when one is making up a new definition, or defining a term for use in a local environment. The above definition of intelligence is useful, since it leads to an algorithm for creating an intelligent machine. The idea is to mimic evolution. First you create a population of machines, each with a reservoir of knowledge which could initially be basically garbage. Then you repeat a two-step process of variation and selection to cause the

[1] Ryerson University, Canada, eharley@scs.ryerson.ca

knowledge to evolve in some direction. The variation step duplicates the machines and applies some random changes to the knowledge in the daughter machines. The selection step tests the machines against some criterion defined in relation to a goal and removes the weaker half of the population. The expectation is that after some number of generations, the resulting machines will be excellent at achieving the desired goal.

To test his idea, Fogel decides to use the algorithm to create a program that can teach itself to play a game without being told ahead of time what are the characteristics of good moves or even what are the features to consider when assessing moves. This would be a learning machine, adapting to an environment that demands ever-increasing levels of skill. Fogel chooses the game of checkers, rather than chess, since checkers like chess involves very little if any luck, and because checkers is a simpler game, more appropriate for his modest computing resources (a 400 MHz Pentium II computer). He chooses a simulated neural network to be the reservoir of knowledge. The neural network is set up to be the evaluator of board positions. In essence, it is a mathematical function whose inputs represent the 32 squares on the checker board, each with a value denoting what piece if any is on the square. The knowledge is captured in the weights of the connections and the thresholds. Some of the history and theory of neural networks is described, so that the reader has a fairly clear idea of what these terms mean and how they relate to concepts such as 'hill-climbing' and 'local maxima'. Variation is carried out by a program which duplicates the neural networks and applies some constrained random variation to the weights. Selection is performed by a program which uses the neural networks as board evaluators in competitions within the population of neural networks. Each neural network plays at least five games, and the neural networks with the best scores live on for the next stage of variation, while the poorer performers are removed from the population. The hope is that after many generations of variation and selection the best neural network will be an expert at checkers (i.e., have a rating above 2000).

Clearly, the above experiment requires a lot of programming, which is undertaken by Kumar Chellapilla, a graduate student with an interest in evolutionary algorithms. While Chellapilla sets up the program, Fogel studies the history of programs that play checkers. One chapter of this book is devoted to Arthur Samuel, a pioneer in artificial intelligence and in the world of checkers programs. Samuel, working at IBM in the 1950's created a checkers program which to some extent taught itself. Samuel constructed its evaluation function from a number of features of positions in checkers which human experts over the years have determined to be important. Samuel then allowed the program to play against itself and adjust the weights of these features to improve its play. The challenge that Samuel left open to future researchers was to design a program that could invent its own features, i.e., one that could learn the game from scratch without human advice. This is one of the goals of the experiments in this book, so it is named Samuel's challenge. Another famous pioneer of computer science, Allan Newell, expressed the view that it is very unlikely that win/lose information is sufficient to allow a game-playing program to improve its performance. This statement contrasts with Fogel's experimental design. The selection step in his evolutionary algorithm is applied after the neural networks have been used to play at least five games. The neural networks of course are given no feedback regarding which moves were good or even which games resulted in a win. Thus, Fogel sets his experiment historically in terms of the *Samuel-Newell challenge*. It is an attempt to do what Samuel wanted to do (make a program that learns features by itself) in a way that Newell thought would not work (using only win/lose feedback).

I do not want to go into the results of Fogel's experiments, since this comprises much of the excitement of the book. He and Chellapilla test the evolved programs on an internet checkers club, and these tests are described with humor and anticipation, as well as a scientific eye. I will just say that Blondie24, the most evolutionarily-advanced of his programs, attains a level which is truly amazing, considering that 'she' is given no knowledge beyond the basic rules and piece count.

In conclusion, let us consider what Fogel has achieved, besides writing a very entertaining and informative book on artificial intelligence. Did he meet the Samuel's challenge? I think the answer here is clearly yes, since the evolutionary program that he designed learned to play checkers at a very high level without input of human expert knowledge. The Notes state that some reviewers objected that his program may simply be using the piece count. He rules this out by showing that the program is much better than any program that bases its decisions simply on the differential and worth of pieces. What attributes or features of checkers did the program discover to be important? No one knows, since that information is buried in the array of weights, but the book presents some anecdotal evidence that mobility might be one of the discovered features. Did he meet Newell's challenge? Without a doubt, since there is no feedback to the program after moves. The program definitely learned to play expert checkers via selection based on the results of five or more games.

Did Fogel create artificial intelligence? I think the answer again is yes. The evolutionary program has a built-in goal to make its neural networks better checkers players. It has a self-created environment, which is the checkers tournament. It makes decisions regarding which neural networks to throw out and which to keep for future generations. It adapts, since the neural networks continually get better under this selective environment. An interesting point is that the end product which looks intelligent is Blondie, yet she is not in fact the intelligence. Like the individual wasp, Blondie is fixed in her responses. If she played a million games, she would not be an iota smarter. In this sense, she is like Deep Blue. In retrospect, if you consider not just Deep Blue but also the program which adjusted the weights of features in Deep Blue's evaluation function based on many master-level games, then that program is also an example of artificial intelligence.

Is Fogel on the right track for creating the kind of general artificial intelligence symbolized in the computer HAL of *2001, A Space Odyssey*? This we do not know. Fogel says he has no idea how to build such a creative computer. I wonder if he has good enough definition of intelligence, or if perhaps his interpretation is not quite right. His definition refers to a decision-making system, and Fogel seems to be quite liberal in his interpretation of what that can be. He says it can be a whole species of wasp, yet this does not fit well with our intuition. Changes in the behavior or anatomy of the wasp presumably arise because of variations in the genetic pool and selection by the environment for the individuals that are better suited for survival. Thus, the species evolves and adapts to the environment, but where are the decisions being made? Neither the species nor the environment seems to be involved in any decision-making. If the wind blows down all but one well-rooted, hardy tree, which goes on to make a forest, then we may say there was selection and adaptation, but not a decision. Therefore, I do not think a species of tree or a species of wasp is intelligent, based on Fogel's definition. Perhaps a better example of intelligence would be an ape that decides to use a stick to probe for termites, or a human, who can adapt her behavior to any number of new challenges.

Fogel believes that the future of artificial intelligence lies in machines that can learn, since only they will be able to solve the problems for which we do not have solutions. Furthermore, he thinks that the way to achieve this goal is by copying the process of evolution, which already has proved successful in generating intelligence. If Fogel is right, then probably what his evolutionary algorithm should be selecting for is the property we define as intelligence, or a concomitant (maybe a sense of humor =)). Time will tell if his definition of intelligence is useful enough to lead him or others to the proper experiment.