Evolving Solutions that are Competitive with Humans

David B. Fogel
Natural Selection, Inc.
3333 N. Torrey Pines Ct., Suite 200
La Jolla, CA 92037
dfogel@natural-selection.com

Extended Abstract

Ever since the advent of the modern digital computer, we've tried to generate machines that are as intelligent as humans. Indeed, this was the primary goal of early artificial intelligence (AI) researchers. The "general problem solver" served as one example of efforts along this line, and it's lack of success was disappointing to many in the AI community. One of the primary deficiencies of early, and even modern, efforts in AI has been a lack of a definition of the intelligence that is sought. There is no accepted definition of intelligence, let alone "artificial intelligence." In place of careful definitions, the Turing Test evolved as a surrogate criterion for judging intelligence.

The famous Turing Test involves an interrogator asking questions of a woman and man via a teletype machine, with the goal of correctly identifying the woman. The woman presumably gives truthful answers to questions, while the man may lie. Turing proposed replacing the man with a computer and suggested that if a computer can fool the interrogator into believing that it is the woman as often as a man can fool the interrogator then the machine will be said to have passed the test. Interestingly, Turing never claimed that passing the test meant that such a machine would be "intelligent," an issue that he described as being "too meaningless to deserve discussion."

In retrospect, the Turing Test is no more a test for intelligence than it is a test for femininity. If a man can fool an interrogator into believing that he is the woman, that does not make the man a woman. Similarly, just because a computer might fool an interrogator into believing that it is intelligent does not make the computer intelligent. Nevertheless, the Turing Test has had a profound impact on efforts to simulate behaviors that we associate with intelligence, mainly as we observe them in ourselves. Unfortunately, over time, the impact was mainly to narrow the focus of AI to simply generate programs that could compete with humans in specialized areas, such as chess. The mechanism for generating the required behavior became irrelevant, as all that was of importance was the end result. The culmination of this process is Deep Blue, a very fast machine that can beat Garry Kasparov in chess, but is no more "intelligent" than a calculator, or a hammer.

Rather than begin from the perspective of the Turing Test, an alternative perspective begins with the concept of decision making. For an organism to be intelligent, it must make decisions. It is pointless to speak of the intelligence of something that does not make decisions. A decision can be defined to arise when available resources are allocated. Note that a range of possible decisions, possible allocations, must be available otherwise there really is no decision at all. Logically, decision making requires a goal, for decision making in the absence of a goal is pointless. Thus we must inquire as to where goals come from.

In natural systems, the primary goal instilled in all living organisms is survival. Those organisms that do not possess this goal may be "successful," but are uninteresting from an evolutionary perspective. Thus behaviors can be judged in how well they support this ultimate goal, and subgoals that underlie it. More generally then, intelligence can be defined as "the ability for a system to adapt its behavior to meet its goals in a range of environments." This capability of intelligent decision making can be observed strikingly in the evolving phyletic lines of organisms (the reader is free to choose which ones), such as frogs and insects in which cryptic coloration, poison, reliable signaling, and mimicking have all been invented to meet the primary goal of "avoiding being someone else's lunch." No single individual invented any of these "tricks," rather the intelligent organism in these cases is the evolving line of individuals.

Taking this cue from nature, it is reasonable to assess the intelligence capability of a machine that evolves solutions to problems in a manner similar to that of evolving phyletic lines in the natural environment. From the perspective of performance comparison, many efforts in evolutionary computation have been measured in light of human capabilities. For example, L. Fogel (Artificial Intelligence Through Simulated Evolution, Wiley, 1966) compared the ability of graduate students and an evolutionary program operating on finite state machines to predict sequences of symbols. The results showed that the evolutionary program was competitive or slightly more capable than its human competitors. In Germany, in the mid-1960s, H.-P. Schwefel used an evolutionary algorithm to create a new design for a flashing nozzle, which exceeded the capabilities of the previous human design. There are many other relevant results in the literature.

More recently, the author and a colleague (Kumar Chellapilla) investigated the ability for an evolutionary algorithm to learn to play checkers at a level that is commensurate with human experts, without relying on human expertise about checkers. Instead, neural networks were used to evaluate candidate board positions based only on the inputs found in the number, location, and types of pieces on the board. Furthermore, the neural networks were not told which games were won, lost, or drawn. Only an overall point value was assessed to each neural network, which signified the total value earned over a series of games.

Starting from randomly weighted connections, a population of neural networks evolved, using random variation of the weights of each neural network and a selective mechanism to eliminate poor-scoring networks, over 100 generations to be competitive with "Class B" human players on the Internet. With some modifications of the input design to the neural networks, which allowed a recognition that the game is played on a two-dimensional board, and 840 generations, the best-evolved neural network (called Blondie24) was able to compete with human experts and finished in the top 500 of over 120,000 people on the Internet site, www.zone.com. More details on this effort can be found in D. Fogel's book, Blondie24: Playing at the Edge of AI, Morgan Kaufmann, 2002.

The results indicate that evolution is a suitable mechanism for creating intelligent behavior in machines, and that it can learn to generate behavior that is competitive with human experts even without relying on human expertise. As computer hardware increases in speed according to Moore's Law, it is important to recall that this acceleration in speed is not sufficient to generate intelligent machines. The software that this hardware will execute is critically important. The results described here and presented in the plenary lecture indicate one step toward creating intelligent machines that may someday possess an ability to adapt their behavior, to meet their goals, in a range of environments that is commensurate with our own abilities.