ARTIFICIAL INTELLIGENCE

Will a Computer Checkmate a Chess Champion at Last?

On one side of the table is artistry and human intelligence; on the other is sheer number-crunching power. Starting on 10 February, Gary Kasparov, recognized as the best chess player in the world, will go face-to-interface with Deep Blue, a computer chess program developed at IBM. And some experts think that for the first time, a computer has a chance of beating a human champion.

There's a tidy purse—$400,000 for the winner, $100,000 for the loser—at stake in the series of six games, sponsored by the Association of Computing Machinery (ACM). But the match, which will take place in Philadelphia, is also a test of technology and programming ingenuity that may pay off in other computer science problems. Deep Blue's developers, led by Chung-Jen Tan, have sharpened its chess insight with new strategies for making the best possible choice in the face of an astronomical number of possible outcomes. "If the computer lives up to its design," says Monty Newborn, a computer scientist at McGill University in Montreal and chair of the ACM's computer chess committee, "it's going to win."

Kasparov sees things differently. Machines may occasionally win at speed chess—Kasparov himself has lost a few fast games to Deep Blue's predecessor, Deep Thought. But "in serious, classical chess," Kasparov has declared in a public statement, "computers do not have a chance in this century."

If he's right, it won't be the first disappointment for computer chess. Nearly 40 years ago, pioneers in the field of artificial intelligence predicted it would only take about a decade for computers to master the game of kings—but then found it hard to program machines to play well even against beginners. In retrospect, say computer chess researchers, the mistake lay in trying to imitate the way humans think about chess rather than relying on computers' natural advantage in speed and precision. "One of the disappointments is that we still don't know what people do" when they play chess, says Newborn.

Computer chess made progress when programmers pursued the more straightforward approach of "outsearching" human opponents: looking many moves ahead and ferreting out what looks like the best line of play. That strategy has already raised some programs that look 10 or 11 moves ahead to the equivalent of International Grandmaster status—a ranking that only around 200 humans currently hold.

Tan and his IBM colleagues Feng-Hsiung Hsu, Murray Campbell, Joseph Hoane Jr., and Gershon Brody aren't taking chances. "They set out to be really sure that this thing would work," Newborn says. As a result, Deep Blue has the horses to look 15 moves ahead. For the Kasparov match, Campbell notes, Deep Blue will take a slightly different tack: It will look at all possibilities only up to 12 moves, but then follow certain high-stakes lines of play as much as 30 moves deep.

To get that kind of power, Deep Blue's designers have had to contend with one of the central problems in artificial intelligence: carrying out a thorough search in the face of what computer scientists call combinatorial explosion. Combinatorial explosion occurs whenever the number of things grows exponentially—a common problem in tasks such as automatic theorem proving, where a machine searches for a proof by examining all possible combinations of the facts it knows. In the case of chess, each player has roughly 30 different moves available each time his or her turn comes, which means that trying to look 15 moves ahead requires sorting through 30^{15}—nearly an Avogadro's number—of different lines of play, branching out from the starting point (see illustration).

AP/WIDE WORLD PHOTOS

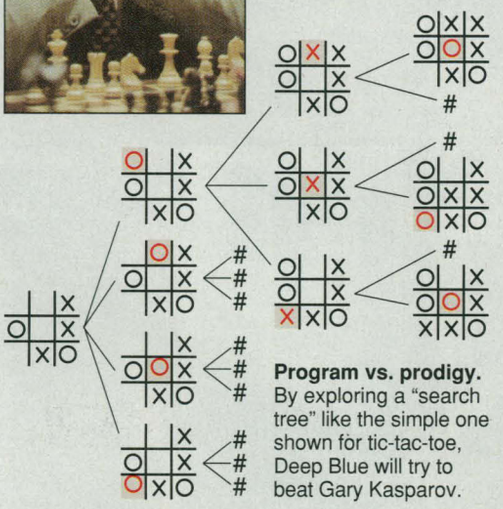

Program vs. prodigy. By exploring a "search tree" like the simple one shown for tic-tac-toe, Deep Blue will try to beat Gary Kasparov.

To choose the best move, Deep Blue has to explore this tree. The basic idea is to calculate all possible lines of play that begin with a particular move and assign each endpoint, 14 moves further down the line, a score representing the position's "value" to the computer—the stronger the position, the higher the score. The score is then propagated back along the tree to the starting position, using a simple criterion: Branch points corresponding to the computer's moves are assigned the highest of the scores they branch to, while branch points corresponding to the opponent's moves are assigned the lowest of the scores they branch to, on the assumption that the opponent will always steer the game toward the worst possible outcome for the computer. Deep Blue then makes the move that sends it along the branch with the highest score.

But because of the combinatorial explosion, even the fastest computer of today would need thousands of years to perform this full search for each move. Deep Blue defuses the explosion by exploring only a small fraction of the tree, relying on what is known as the "alpha-beta" algorithm to determine which lines of play it is safe to ignore. The basic idea is that once the program has identified a "good" move, it can cut short the consideration of an alternative move as soon as that alternative turns out to lead to a worse position.

In order to get the most pruning out of this strategy, the IBM team has also organized the search so that Deep Blue can identify good moves (if not the best one) relatively early. One technique is to make a "shallow" search first—say 10 moves deep—to get a rough indication of which moves are promising, and then redo the search at greater depth for those moves first. (Deep Blue actually uses snazzier, more complicated methods.) Deep Blue also curtails some searches by taking advantage of the fact that different lines of play frequently lead to the exact same position.

These clever search strategies cut the number of positions Deep Blue must look at from Avogadro's number down to something on the order of a hundred billion. "Those numbers are starting to look feasible," says Newborn. Even so, it's not something you can run on a PC. The Deep Blue system consists of several hundred special-purpose chess chips running in tandem, enabling it to search an average of 100 million positions per second.

Of course, Kasparov has a few billion processors (a.k.a. neurons) of his own devoted to the game. Campbell will only forecast "a very entertaining match" with "a number of interesting games." Kasparov, Newborn notes, is particularly strong in the opening phase of the game. But if Deep Blue can survive into the so-called middle game, Newborn thinks it stands a good chance of defeating the champ. And if not, he says, it's just a matter of time before Kasparov, or whoever comes next, plays second fiddle to an algorithm.

–Barry Cipra