Knowledge Discovery in
Deep Blue

*A vast database of human experience can be
used to direct a search.*

Murray Campbell

Deep Blue was the first chess computer to defeat a reigning human world chess champion in a regulation match. A number of factors contributed to the system's success, including its ability to extract useful knowledge from a database of 700,000 Grandmaster chess games—a process that has implications for any non-chess knowledge-discovery application involv-

ing large databases of expert decisions. A tremendous amount of information is contained in Deep Blue's database, which the system used in its matches against Garry Kasparov in 1996 and 1997, as well as in earlier games. The games in the database represent a rough consensus, developed over decades of play, of the best opening moves known.

It is a tribute to the richness of chess that this consensus, known to chess players as "opening theory," is surprisingly dynamic. Openings that may have been considered bad are frequently revived with a new idea in mind; blind faith in opening theory can be as hazardous as not knowing it at all. Grandmasters devote a great deal of time and study to opening theory, poring over the games of their chess-playing peers from around the world, subjecting a small subset of the resulting positions to very deep analysis.

Deep Blue takes a different approach to using the opening information in its database. Instead of subjecting a few selected positions to detailed analysis, it creates an additional database we call the "extended book." The extended book allows the system to quickly summarize previous Grandmaster experience in any of the several million opening positions in its game database. This summary is used to nudge Deep Blue in the consensus direction of opening theory. The system can combine its massive searching ability (200 million chess positions per second) with the summary information in the extended book to select opening moves. But Deep Blue does not blindly follow opening theory, using it instead to guide its search for the most promising move. This process gives Deep Blue the opportunity to not only find errors in opening theory, but to follow the branch of the theory that best fits its own chess-playing style.

Extended Book

Most chess-playing computers, including Deep Blue, have their own "opening book," a set of positions, along with associated recommended moves. While using the opening book to navigate the early stages of a chess game in which one of the included positions arises, the computer selects one of the recommended moves by way of some possibly random mechanism, then plays without further computation.

Creating a high-quality opening book is time- and labor-intensive, requiring a degree of chess skill. Because this task was not our primary focus, the opening book in Deep Blue and in its predecessor machines has contained only a few thousand positions. In order to compensate for this relatively small book, we introduced (in 1991) the idea of an extended book derived *automatically* from a Grand-

Figure 1. Position after White's move 19 in the second game of the 1997 match between Deep Blue (White) and Kasparov (Black), ultimately won by Deep Blue in 45 moves.

□ Deep Blue
■ Kasparov, G.

IBM Deep Blue vs. Kasparov Rematch, New York, 1997

1.e4 e5 2.Nf3 Nc6 3.Bb5 a6 4.Ba4 Nf6 5.0-0 Be7 6.Re1 b5 7.Bb3 d6 8.c3 0-0 9.h3 h6 10.d4 Re8 11.Nbd2 Bf8 12.Nf1 Bd7 13.Ng3 Na5 14.Bc2 c5 15.b3 Nc6 16.d5 Ne7 17.Be3 Ng6 18.Qd2 Nh7 19.a4
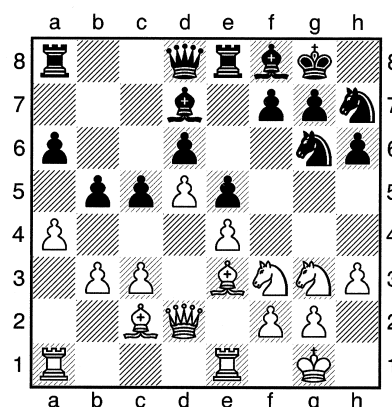
Nh4 20.Nxh4 Qxh4 21.Qe2 Qd8 22.b4 Qc7 23.Rec1 c4 24.Ra3 Rec8 25.Rca1 Qd8 26.f4 Nf6 27.fxe5 dxe5 28.Qf1 Ne8 29.Qf2 Nd6 30.Bb6 Qe8 31.R3a2 Be7 32.Bc5 Bf8 33.Nf5 Bxf5 34.exf5 f6 35.Bxd6 Bxd6 36.axb5 axb5 37.Be4 Rxa2 38.Qxa2 Qd7 39.Qa7 Rc7 40.Qb6 Rb7 41.Ra8+ Kf7 42.Qa6 Qc7 43.Qc6 Qb6+ 44.Kf1 Rb8 45.Ra6 1-0

master game database. For each of the (several million) positions arising in the first 30 or so moves in these Grandmaster games, the system computes an evaluation for each move that has been played. Moves that appear to be good are given positive values; moves that look weak are assigned negative values. Our aim was to bias the search, preferring the strong-looking moves while avoiding the weak ones.

Prior to conducting a search, Deep Blue first checks whether a move is available from the opening book. Finding a move, it plays it immediately. Otherwise, it consults the extended book; if it finds the position there, it uses the evaluation information to award bonuses and penalties to a subset of the available moves. Finally, Deep Blue carries out a search, with some preference for following successful Grandmaster moves. However, because there is always the possibility that opening theory has overlooked the strongest move in the position, Deep

Blue still has the opportunity to discover it. In some situations, where the bonus for a move is unusually large, Deep Blue can make a move without computation; such automatic extended-book moves allow the system to save time for use later in the game.

A number of factors goes into the evaluation function for the extended book, including:

- The number of times a move has been played. A move frequently played by Grandmasters is likely to be good.
- Relative number of times a move has been played. If move A has been played many more times than move B, then A is likely to be better.
- Strength of the players that play the moves. A move played by Kasparov is more likely to be good than a move played by a low-ranked master.
- Recentness of the move. A recently played move (usually played knowing prior games) is likely to be good, an effect that can in some cases dominate the second factor.
- Results of the move. Successful moves are likely to be good.
- Commentary on the move. Chess games are frequently annotated, with the annotator marking strong moves (with "!") and weak moves (with "?"). Moves marked as strong are likely to be good; moves marked as weak are likely to be bad.
- Game moves vs. commentary moves. Annotators of chess games frequently suggest alternative moves. In general, game moves are considered more reliable than commentary moves, and are thus likely to be better.

Deep Blue's designers combined these factors in a nonlinear function that produces bonuses as high as half the value of a pawn (50 points, where a pawn is worth 100 points) in situations in which a particular move has been played thousands of times with reasonable success.

The extended book idea has proved remarkably successful in the games Deep Blue and its predecessors have played. The system is often able to follow opening theory without any opening book at all, relying only on the extended book combined with its own search. An interesting example of the value of the extended book was demonstrated in the second game of Deep Blue's 1996 match with Kasparov. In that game, a configuration error caused Deep Blue to play the entire game without an opening book. Despite that limitation, the extended book allowed Deep Blue to follow opening theory up to move 13, after which Kasparov introduced a new move. Although this game, and the 1996 match, were ulti-mately won by Kasparov, the extended book allowed Deep Blue to obtain a respectable opening position.

More noteworthy was Deep Blue's victory over Kasparov in the second game of their 1997 match (see Figure 1). In that game, the first nine moves of a Ruy Lopez opening were taken directly from the opening book. The next eight moves (10–17) were all automatic extended-book moves. Two moves were given extended-book bonuses: 18.Qd2 was given a 17-point bonus, and 19.a4 was given a nine-point bonus. On move 19, Kasparov played a new move that led to a position not found in Deep Blue's extended book. From this point on, Deep Blue played without any extended-book bonuses, having established a very comfortable position from the opening.

Useful in Non-Chess Domains

Deep Blue uses knowledge extracted from a Grand-master game database to improve its performance in actual play. The extended book technique, involving summarized human decisions to bias a search, also appears to be of general value in non-chess domains with access to large databases of expert decisions, such as other games, medical diagnosis, and stock trading, especially when there is a useful similarity measure between differing situations for a given domain.

For playing chess, one could extract knowledge from a Grandmaster game database in ways other than the Deep Blue extended book. For example, we have experimented with methods for tuning Deep Blue's position-evaluation function by, say, trying to maximize the number of Grandmaster moves that can be predicted through a shallow search. Though this optimization problem is very difficult, due to the size of the parameter space that has to be explored, the automated turning method has yielded insight and guidance to the Deep Blue team in the design and programming of the Deep Blue system. **C**

REFERENCES
1. Hsu, F.-H. IBM's Deep Blue Chess Grandmaster Chips. *IEEE Micro 19*, 2 (Mar./Apr. 1999), 70–81.
2. Newborn, M. *Kasparov vs. Deep Blue: Computer Chess Comes of Age.* Springer-Verlag, New York, 1997.

MURRAY CAMPBELL (mcam@us.ibm.com) is a research scientist in the IBM T.J. Watson Research Center in Yorktown Heights, N.Y.