FIRST SET OF NOTES

IMPLICATIONS OF DEEP BLUE FOR ARTIFICIAL INTELLIGENCE

(Notes by Herbert A. Simon, 5/18/97)

In all of its announcements about the recent chess match between Garry Kasparov and Deep Blue, IBM has carefully avoided the use of the term "artificial intelligence," and has focused its attention almost exclusively on the large parallel processor that examines 2x10^8 positions per second, or about 1.8x10^10 each three minutes (the average time allowed for a move). Thus: "The power behind Deep Blue is an IBM RS/6000* SP* system finely tuned with customized processor chips designed by IBM Research. This combination, IN ADDITION TO EXPERT KNOWLEDGE [italics ours], enables users to take on larger problems by analyzing a greater number of possible solutions."

One may be pardoned a certain skepticism about this characterization of Deep Blue's prowess, or its implications for computing in general and AI in particular. In particular, for the future, should we look to machine speed plus expert knowledge, or expert knowledge plus machine speed? It should be noted that both Joel Benjamin, the Grandmaster who served as consultant to the DB team, and Murray Campbell, the chess expert among its programmers, claimed that important progress had been made during the past year in advancing DB's chess knowledge, and not simply in allowing it to examine more positions.

Let us consider first, the matter of speed. During the year, the number of positions it could examine in a given amount of time was increased by a factor of two. Since the branching factor at each move (ply) averages around 20, this means that the search was increased by only a fraction of a ply; yet all Grandmasters who watched the match (and Kasparov as well) commented emphatically on the QUALITATIVE change in its play, and especially its ability to adapt its play to the character of the position it faced. So it is to chess knowledge that we must look for most of the improvement. What form could this knowledge take?

As far as we know, the knowledge takes three forms. First, DB has a large opening "book" and probably a large endgame "book," each containing tens of thousands (?) of game trees, both from actual games and from analysis. These books probably contain about the same order of magnitude of specific information about the opening and end games as is possessed by Kasparov or other Grandmasters, although the evaluations of the positions reached on these trees (except when they reach the game's end) are probably of poorer quality than Kasparov's evaluations. Second, DB has assigned a numerical value to each of the leaf positions, so that different positions can be compared; and of course it is able also to assign numerical values to new positions it reaches during the game.

Third, and perhaps most important, DB appears to be able to CHARACTERIZE positions; that is, to notice patterns of piece arrangements in a position that distinguish this kind of position from another. Ability to classify positions would allow it to adjust its evaluation function to the character of the position — to value different features differently in different kinds of positions. It would also allow it, if it did not search exhaustively at each branch, but selectively, to search in different directions, depending on the character of the position. The programmers of DB have explicitly stated that it adjusts its evaluation function to the type of position, and just as explicitly that, at least under certain circumstances, it conducts selective searches, sometimes very deep but narrow ones.

Thus, there are three directions, in addition to increased speed, in which DB's performance could have improved during the year, and every indication, from its play, that there was improvement along these lines. It is much easier to explain its strength in these terms than by supposing that a mere doubling in speed could have caused major changes in the character of its play.

What is especially interesting about these three kinds of chess knowledge is that they are exactly the kinds of knowledge that human chess players possess and that give them the capability to select good moves after only a modest amount of search (almost certainly never more than 10^3 branches on the game tree, except when the "book" is being followed). We do not know how DB's new speed is allocated, but most of it may be devoted, not to searching deeper, but to computing and recognizing more sophisticated positional patterns in order to permit more selective search. In fact it is possible (there are no known facts one way or the other) that 1997 DB carried out LESS search than 1996 DB. There are numerous precedents for selective search in chess programs, beginning with the NSS program of 1958.

Putting chess aside, what does this mean for the future of computers, applied to very complex, knowledge-rich tasks, and for the role of AI methods in performing these tasks? First, speed is not enough; it must be supplemented by knowledge. Moreover, if there is enough knowledge, the capability for recognizing cues, and thereby accessing knowledge associated with particular kinds of situations may replace speed as the principal tool for high performance.

This does not mean that computer power becomes unimportant. In the human, recognizing patterns in visual and auditory stimuli appears to be, perhaps by orders of magnitude, the most expensive computing task that the human brain performs, and the one whose real-time demands make parallel processing essential. We can think of the expert system as a very large indexed encyclopedia, where the index entries correspond to cues that can be recognized in the task environment, and where recognition of a cue (a pattern) provides access to the information stored in association with that pattern. The information will normally include recommendations for the actions to be taken (including additional information-gathering actions) whenever a given pattern is recognized.

Of course, pattern recognition is not the whole story. In addition to recognizing patterns, an intelligent system must be able to carry out a certain amount of inference, of the kinds we have become familiar with from means-ends analysis and other forms of selective search in problem solving systems. But if the problem-solving search is to be selective, as it must be in complex situations, then it must again be guided by the ability to recognize patterns at branch points that tell it which branches are most worth exploring.

But if pattern recognition for selecting and evaluating search and action alternatives is to be a central component of systems for dealing with complex situations, then methods are needed for acquiring and evaluating patterns. Programming cannot do the job, especially if the indexed encyclopedia must be updated continually as new knowledge becomes available. It becomes desirable, indeed necessary, to develop machine capabilities for learning patterns autonomously from information about the task environment. Of course, we have considerable experience, today, in building programs that learn — an experience that extends all the way back to Arthur Samuel's impressive program that not only played checkers very well but, as early as the mid-1950's, learned to improve its game by storing information about evaluated variations and to modify its evaluation function with experience.

Our experience with learning follows at least four main paths: learning in so-called neural networks, learning by adaptive production systems that build new productions and incorporate them in themselves, discrimination nets that learn (nets up to about 100,000 leaf nodes have been grown), and genetic algorithms. The experience with DB suggests that a high research priority should be given to incorporate learning capabilities of one or more of these four kinds to enable the system to learn from published games and from its own analyses.

In summary, the current emphasis on brute force — larger and faster computers — for advancing expert systems and AI does not seem to be supported by the evidence provided by the recent chess match between Deep Blue and Kasparov. On the contrary, improved chess knowledge seems to be the main key to the improved performance. We have good basic understanding of the kinds of knowledge that are required, and of techniques for building memories of the requisite size and organization. Handcrafting these memories, for chess or any other task, is not a pleasing prospect. Learning programs offer a much more promising direction.