This is a clear example of reasoning by the program.

The KPK domain is treated in Chapter 6. Previously, results were published in *Advances in Computer Chess 3*. Although very simple, the KPK endgame contains a surprising degree of complexity. According to the author:

> "Although almost trivial to a master, [the KPK endgame] is too complex for any but the most scrappy and vague rules to have been articulated in existing chess manuals"

Various successful attempts of catching this endgame into attributes and rules have been made in the past, for example by Bramer. Therefore, this domain serves as a test to check whether the developed method will yield equally good and brain-compatible results. The problem is decomposed into 5 subproblems, all independent, and 31 unique attributes are defined. An example of each subproblem is given. At the end of the chapter, a cost-analysis is presented in terms of programming time and run-time. Finally, in Chapter 7, Shapiro faces the challenge of solving the very complex KP(a7)KR endgame by induction of rules. For this particular endgame, no human codification seems to have been achieved at all. Again, the problem is broken up into (8) subproblems and (36 unique) attributes are defined. In this case however, the subproblems make up a tree-like structure of variable depth. As in the KPK endgame, each subproblem is outlined and illustrated by examples. The program was developed in four stages, each taking about 3-4 man-weeks of effort. Using self-commenting as a diagnostic aid, the system was tuned to yield a complete and correct rule of solution which is both machine-executable and brain-compatible.

*Structured Induction in Expert Systems* is an excellent introduction into the exciting field of artificial intelligence. No prior background knowledge on the subject is needed to enjoy this crisp and clean-cut book. By starting off with a seemingly trivial example (umbrella problem), the reader is introduced to basic techniques and methods used for tackling seemingly simple (KPK) and very complex (KP(a7)KR) domains. The text is adorned with abundant examples and figures. I enjoyed reading this book and strongly recommend it to anyone interested in machine learning.

## LITERATURE RECEIVED

## MASSIVELY PARALLEL RETROGRADE ENDGAME ANALYSIS

*Lewis Stiller*

Boston University
October 12, 1988

Boston University Computer Science Tech. Report #88-014

We quote the abstract:

"A massively parallel algorithm for the solution of certain chess endgames is described. The implementation on a 32K-processor Connection Machine Model CM-2 is currently about 4000 times as fast as previous coarse-grained implementations. Much of this speedup is due to greater RAM usage, rather than to increased processing power of the CM-2. The algorithm used is somewhat different from the most efficient serial or coarse-grained MIMD algorithm. The search space is encoded in the network topology: the motion of a point through the topology corresponds to the computation of the generator operator on that point. Related domains and possible extensions of the algorithm are discussed."