The Implementation of Artificial Intelligence and Machine Learning in a Computerized Chess Program

by James Mannion
Computer Systems, 2008-2009
Period 3

Abstract

Computers have developed to the point where searching through a large set of data to find an optimum value can be done in a matter of seconds. However, there are still many domains (primarily in the realm of game theory) that are too complex to search through with brute force in a reasonable amount of time, and so heuristic searches have been developed to reduce the run time of such searches. That being said, some domains require very complex heuristics in order to be effective. The purpose of this study was to see if a computer could improve its heuristic as it runs more searches. The domain used was the game of chess, which has a very high complexity. The heuristic, or evaluation function, of a chess program needs to be able to accurately quantify the strength of a player's position for any instance of the board. Creating such an evaluation function would be very difficult because there are so many factors that go into determining the strength of a position: the relative value of pieces, the importance of controlling the center, the ability to attack the enemy's stronger pieces – something that chess masters spend entire lives trying to figure out. This study looked to see if it was possible for a computer program to "learn" an effective evaluation function by playing many games, analyzing the results, and modifying its evaluation function accordingly.

Background

In 1950, Claude Shannon wrote a groundbreaking paper called "Programming a Computer for Playing Chess." It dealt with various issues, such as how one might go about writing an artificial intelligence program that could play chess well, and what the evaluation of such a program might look like. It discussed the problems involved with brute-force chess algorithms and suggested an AI algorithm that a computer chess program could use. It suggested using a 2-ply algorithm (an algorithm that looks two turns into the future before applying the evaluation function), however if the program finds that a move could lead to a check or checkmate, then it would investigate that move and subsequent moves out to as many as 10 turns. At the end of the paper, Shannon speculated about the possibility of computer programs that can "learn," but noted that such developments were probably man years down the road.

References

[1] Shannon, Claude E. "Programming a Computer for Playing Chess". 1950.
[2] Beal, D.F. and Smith, M.C. "Temporal Difference Learning for Heuristic Search and Game Playing". 1999.
[3] Moriarty, David E. and Miikkulainen, Risto. "Discovering Complex Othello Strategies Through Evolutionary Neural Networks".
[4] Huang, Shiu-li and Lin, Fu-ren. "Using Temporal-Difference Learning for Multi-Agent Bargaining". 2007.
[5] Russell, Stuart and Norvig, Peter. Artificial Intelligence: A Modern Approach. Second Edition. 2003.
[6] Asgharbeygi, Nima, Stracuzzi, David and Langley, Pat. "Relational Temporal Difference Learning".

Development

The first stage of the program simply involved a console-based chess game where two humans could input their moves into the command line, and the board would be re-printed with the new move, as shown
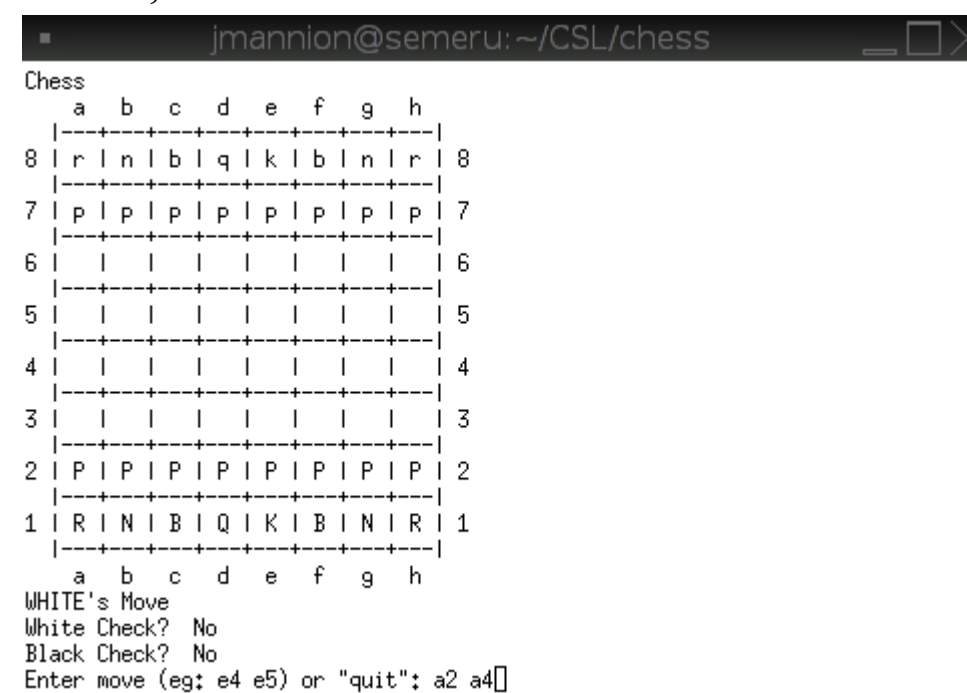
$$w\_i < -- w\_i + a((n\_ic - n\_ip)/(n\_ic))$$
$$h = c1(p1) + c2(p2) + c3(p3) + c4(p4) + c5(p5)$$
$$pi = sum(wi) - sum(bi)$$
$$p\_i = Sum(w\_i) - Sum(b\_i) \text{ over } i$$

Testing/Analysis

The program with TD learning was run against a static program. All weight values in the heuristic were initialized as 1. The win-loss differential of the learner was tracked over 2,000 games to see if the learner would start winning more and more as it underwent the learning process.

The altered piece-square tables were analyzed after the learning session was complete to see if the program developed generally accepted strategies such as maintaining control of the center. Higher numbers toward the center of the boards would indicate such a strategy.