Machine Learning of Bridge Bidding

By Dan Emmons
Computer Systems Laboratory
2008-2009

The Goal

The goal of this project is to create an effective machine bidder in the card game of bridge. Factors like partial information and the multiplicity of the meanings of bids make this task difficult. This research proposes to overcome these problems with the use of a Monte Carlo simulation method for overcoming the limitation of partial information and a tree structure of constraints paired with sets of actions to store the bidding system used by a partnership. With this tree structure a machine partnership trains by continually swapping their new bidding inclinations to learn new decision networks. The performance of bidders is evaluated by having them play against a control pair in both directions for each hand and converting the results to an average IMP gain per hand. The results of this project will not only demonstrate the feasibility of having a machine learn to bid in this manner, but also may develop new bidding conventions useful to human bridge players.

Development

Monte Carlo simulation is used to allow bidders to envision other hands. Approximations of other hands become successively better throughout the auction because more bids give mo

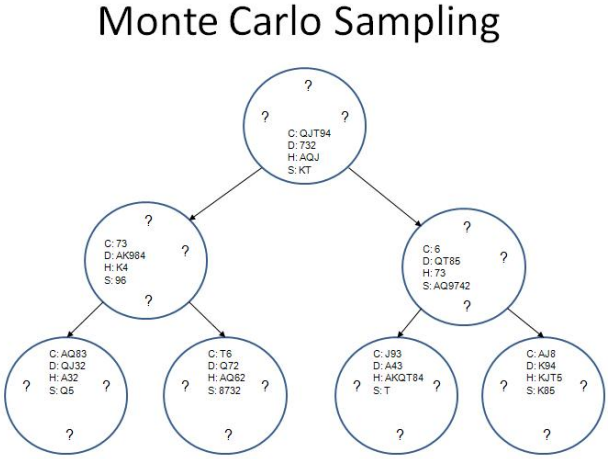

The bidding hierarchy used by the agents to determine what a certain conventional bid means or implies about their hand is represented as a tree, where each node contains hand constraints and actions. There are many priorities of pointers, and feasible actions are determined by taking the union of the sets of actions of the bottom-most nodes reached.

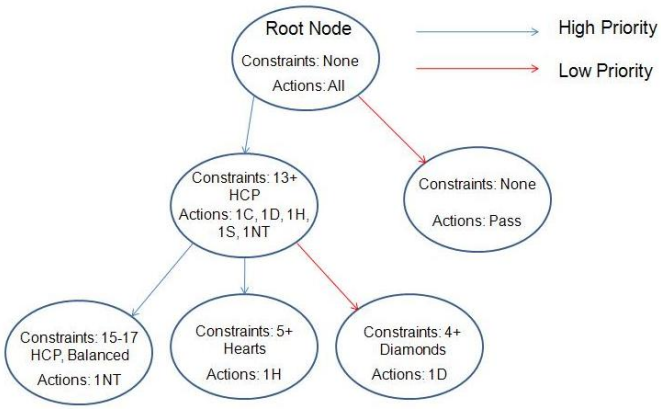

Performance

Here are sample hands that the program bid. In the first one, the program uses nothing other than its expected value algorithm with Monte Carlo simulation. The program earns a good score, but this is in part due to luck. In the second one the program combined these techniques with the use of bidding conventions to more accurately express its hand and earns a fantastic score that is very much grounded in skill.

Dealer: West
Vulnerable: None

North
Clubs: A K 7 6
Diamonds: J T 8 4
Hearts: Q T 8 3
Spades: 2

West
Clubs: 9 8 5 4
Diamonds: 9 7 6
Hearts: J 2
Spades: 8 7 6 3

East
Clubs: J 2
Diamonds: A Q 2
Hearts: A K 9 7 6 4
Spades: K 9

South
Clubs: Q T 3
Diamonds: K 5 3
Hearts: 5
Spades: A Q J T 5 4

South	West	North	East
		Pass	Pass
Pass			
2S	Pass	3H	Pass
3S	X	4C	Pass
4S	Pass	4NT	Pass
5C	Pass	5H	Pass
5S	X	Pass	Pass
Pass			

5SX Nonvul - South
Making Exact
Score: 650

Dealer: West
Vulnerable: E-W

North
Clubs: Q J 8 4 3
Diamonds: K T 4
Hearts: K J 9
Spades: K 7

West
Clubs: T 6 2
Diamonds: A 5 3 2
Hearts: 8 6 2
Spades: Q 6 2

East
Clubs: 5
Diamonds: 9 8 7 6
Hearts: Q 5 3
Spades: A J 9 4 3

South
Clubs: A K 9 7
Diamonds: Q J
Hearts: A T 7 4
Spades: T 8 5

West	North	East	South
Pass	1D	Pass	1H
Pass	2C	Pass	3NT
Pass	4C	Pass	4H
Pass	Pass	Pass	

4H Nonvul - South
Making Exact
420

Results

The implemented bidding agents have performed spectacularly. Without the use of bidding conventions, their bid sequences are awkward and unpredictable, but the agents frequently score well regardless. Using these conventions the agents bid quite skillfully, both against other computers and human opponents. This has been quantitatively verified with IMP scoring matches against players of both types.