Machine Learning of Bridge Bidding

By Dan Emmons
Computer Systems Laboratory
2008-2009

Abstract

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. My plan to overcome these problems is to use a Monte Carlo sampling method for overcoming the limitation of partial information and to use a tree structure of constraints to store the bidding system used by a partnership. Then a machine partnership can train by continually swapping their new bidding inclinations in order to learn new decision networks. The results of this project will not only demonstrate the feasibility of having a machine learn to bid in this manner but also the machine partnerships may develop new bidding conventions useful to human bridge players.

Introduction

Games such as bridge are frequently used as test of and a medium for developing intelligent decision-making algorithms. So far the bidding phase of a bridge hand as proved to be very difficult for machines to perform well for several reasons. Unlike many other games, bridge involves interactions with both cooperative and opponent agents. Also, only part of the total information is available to each agent while bidding, making it difficult to evaluate the effectiveness of given bid until the entire bidding sequence is complete. Perhaps the greatest challenge of all is that bids made by human players usually have more than one meaning. Bids often times suggest contracts to play, but just as often are made to provide the bidder's partner with more information about the bidder's hand. Sometimes they even direct the bidder's partner about what to do while defending. Because of the limited amount of information available to each bidder it can be difficult to determine which of the meanings was intended, or if multiple where intended. The limited information even makes it hard to decide what information would be most useful to give to one's partner about their hand. All these issues would seem to indicate that machine bridge bidding requires a very different category of solution from solutions to problems in other games.

The goal of this project is to create a machine bidding agent as skilled as possible. This will be done using a flexible bidding system framework that can evolve through co-training over the course of many hands. The important developmental points in this project will be designing this framework or choosing from among those already designed, choosing an effective method to overcome partial information, and finding an effective way for partnerships to continually exchange information about their bidding tendencies so that they might improve through co-training. Progress will be measured using the average change in IMPs scored per hand in simulated team matches. This project will not deal with the playing phase of a hand beyond what is required to evaluate the performance of bidding agents.

Development

The development of a fast double-dummy solver was the first step taken in this project. This tool is necessary for evaluating the performance of bidding agents.

The algorithm used to count the number of tricks available in a contract is based on the standard minimax algorithm. The algorithm is given the remaining cards of the hand that must play next, the cards that have been played so far in the current trick, a goal trick count to reach, and the number of tricks taken so far. It will recursively call itself with many different cards from its hand having been played each time. The algorithm is initialized first with the goal number of tricks being 7, and subsequent calls are made that shift the goal to the midpoint of the range that is known to contain the goal.

If the algorithm is run on a hand position in which the opponents have taken too many tricks to allow the declarer to reach the goal, the algorithm returns to its caller that the goal number of tricks cannot be reached from the hand position. If the algorithm is run on a hand position in which the declarer has already reached his goal number of tricks the algorithm returns to its caller that the goal number of tricks can be reached from the hand position. If a defender sees just one position in which the declarer cannot make the goal number of tricks, he prematurely returns that the goal cannot be reached. If the declarer or his partners see a position in which they can reach the goal number of tricks, they prematurely return this result.

Sample Hand Analysis

North:
Clubs: T 7 5 3 2
Diamonds: J
Hearts: A Q J T
Spades: T 9 7

West:
Clubs: 6
Diamonds: A K T 7 5
Diamonds: Q 9 8
Hearts: 9 8 4
Spades: Q J 6 2

East:
Clubs: A J 8
Hearts: 5 3
Spades: A K 8 5 4

South:
Clubs: K Q 9 4
Diamonds: 6 4 3 2
Hearts: K 7 6 2
Spades: 3

Trick Counts for Each Declarer (North, South, East, West):
Clubs: 9 9 3 3
Diamonds: 2 2 11 11
Hearts: 7 7 3 3
Spades: 0 0 11 11
No Trump: 2 2 8 8

Results

Currently no bidding agents have been implemented and so there are no results to report yet. However, once machine partnerships are able to co-train the double-dummy solving tool that has been completed will be able to compute the progress of the partnership. It is expected that machine partnerships should be able to bid at least as well as an intermediate player when the project is complete.