

A Bayesian Poker Player

Kevin B Korb, Ann E Nicholson, et al.

Bayesian Poker Player (BPP) is a computer program built by a succession of students at Monash University that utilizes the Netica API and a Bayesian network to play poker. In its original form (in the honours work of Natalie Jitnah) it played five-card stud. As on-line poker has gained in popularity, it is a seven-card variation, Texas Hold'em, that has won out, especially in automated contests between "poker bots" (see <http://www.cs.ualberta.ca/pokert/>). BPP has been modified to play Texas Hold'em and currently plays at a weak amateur level.

1 Texas Hold'em

First I will explain roughly how poker is played, in particular the Texas Hold'em version. This will include basic ideas of bluffing and strategy.

2 BPP

I will introduce the existing Bayesian network, which models both the computer's and the opponent's hands, including unknown (hidden) cards. In order to simplify computations, similar possible hand types are grouped together. For example, instead of representing every possible busted hand combination (hands without pairs or anything better), busted hands are put into one of five categories, depending upon the high card (e.g., "Busted King High"). In any given situation, cards known to the computer player (its own cards plus those showing on the table) are entered as observations and the Bayesian net is updated, yielding a posterior distribution over outcomes in a showdown (win, lose or draw). This distribution is combined with a utility computation of potential winnings (or losses) to choose the optimal

action (check, bet, raise or fold). It also includes a primitive facility for bluffing.

This BPP can be played against at:

<http://www.csse.monash.edu.au/~stevenm/texasholdem/texasholdem.html>

3 Enhancements and Potential

There are numerous opportunities for enhancing the performance of the basic BPP. Some that have been attempted already include a simple representation of the opponent's psychological state, including conservative versus aggressive personalities and the opponent's tendency to bluff. We have also used "likelihood evidence" to merge the predictions of a combinatorial predictor with the Bayesian net's prediction of the final hands. Another thing studied is increasing the resolution with which hands are represented (i.e., the number of distinct hand types represented).

After reviewing these, I will discuss some possibilities for future exploration. The most important seems to be a much more refined and sophisticated opponent model, one that includes playing tendencies at different stages in the game and tendencies to bluff and to sandbag (bluff a weak hand). Others include improved estimation of "pot odds", i.e., the potential payoff when staying to a showdown (potsize minus the computer's contributions) and opportunities for incorporating the machine learning of opponents' behavior.