

# Towards Approximately Optimal Poker

Jiefu Shi and Michael Littman

Department of Computer Science, Duke University

Durham, NC 27705

{jshi | mlittman}@cs.duke.edu

Creating strategies for different games forces us to grapple with different types of decision-making challenges. Poker is a stochastic game of imperfect information; unlike games of complete information, game-theoretic optimal strategies for poker can be randomized. Koller and Pfeffer [1] argue that two-player poker can be solved efficiently in the size of the game tree using a clever mapping to linear programming.

Texas Hold'em is the variant of poker used in championship tournaments. We are working on methods for generating approximately optimal strategies for Texas Hold'em. To give the flavor of our approach, we describe initial experiments in using abstraction.

For games with a single betting round, we use a grouping method to reduce the number of distinct hands considered. We enumerate all possible hands by their hand strength to obtain a ranking of each hand (a hand with higher value always beats the hand with a lower value). Next, we group hands into *bins*. Each bin contains hands with similar rankings, so, a hand belonging in a bin with a rank of three beats all hands in bins one and two. The game is then solved at the level of bins: we imagine that players are randomly assigned to bins, with the highest ranking bin the winner; betting strategies are computed for the resulting game. The number of bins used in the approximation controls the degree of abstraction and can be adjusted to accommodate space and time requirements.

For our test-bed, we used a game with 200 possible hands. We first generated the optimal strategy for the first player (dealer). We then ran experiments dividing hands into from 4 to 200 bins and produced strategies for the second player (gambler) based on each of these groupings. Figure 1 shows how well the gambler fared against the optimal strategy of the dealer based on 100,000 games (the gambler has an advantage in this and most poker games). The results are quite encouraging; using as few bins as 10% of the number of hands, the resulting play is almost as good as that of the optimal strategy.

In games with multiple betting rounds, the principle concern is "hand potential". The player is forced make decisions based on partial hands. The same ranking trick would not work as in the one round case. A hand with high potential of developing into a strong hand may not have high hand value currently (for example, four-card

flush). For these games, we rank the hands based on the average value of all possible complete hands that a partial hand can develop into. We then group them into bins and play each bin against another to get the expected payoff; this is used in defining the payoffs for the "abstracted" game. And, for cards yet to come, we also use the notion of abstraction. Instead of performing our calculation based on each possible card to come, we also group those cards yet to come into bins (for example, only knowing a low diamond is coming rather than knowing that the three of diamonds is coming).

We introduced a reduced version of Texas Hold'em that consists of a 52 card deck but with 3 cards in play. Figure 2 shows the result of the gambler versus the dealer (the dealer is using the optimal strategy). Just as in the game with the single betting round, we can do quite well just using a very small number of bins.

Figure 1. No Hand Potential





Figure 2. Hand Potential



## References

- [1] D. Koller and A. Pfeffer, 1997. "Representations and Solutions for Game-Theoretic Problems," *Artificial Intelligence*, vol. 94, no. 1-2, pp. 167-215