

The 2006 AAAI Computer Poker Competition

Michael Littman¹

Martin Zinkevich²

¹ Rutgers Univ., Piscataway, NJ US; ² Univ. of Alberta, Edmonton, Alberta, Canada

1. SUMMARY

At AAAI, the University of Alberta hosted a computer poker competition in which participants were required to submit computer programs capable of playing Heads-Up 10/20 Limit Texas Hold'em. There were two tournaments: the bankroll tournament, and the series tournament. Five teams entered the competition.

- The team from the University of Alberta (Edmonton, Canada), consisting of Michael Bowling, Martin Zinkevich, Darse Billings, Nolan Bard, Morgan Kan, Michael Johanson, Robert Holte, Jonathan Schaeffer, Neil Burch, Carmelo Piccione, and Finnegan Southey, developed **Hyperborean** based on ideas from past work (Billings *et al.*, 2003). Two different variants of Hyperborean competed in the two tournaments.
- The team from Monash University (Victoria, Australia), consisting of Ann Nicholson, Kevin Korb, and Steven Mascaro, designed **Monash BPP** based on ideas from past work (Korb, Nicholson, and Jitnah, 1999). Monash BPP competed in both tournaments.
- Teppo Salonen of Irvine, CA, US designed **BluffBot** (Salonen, 2006). BluffBot competed in both tournaments.
- The team from Carnegie Mellon University (Pittsburgh, US), consisting of Andrew Gilpin and Tuomas Sandholm, developed **GS2** (Gilpin and Sandholm, 2006). GS2 competed in the Series Tournament.
- Morten Lynge of Ikast, Denmark developed **Teddy**. Teddy competed in the Bankroll Tournament.

In the bankroll tournament, four bots competed: Hyperborean, Monash-BPP, BluffBot, and Teddy. Each pair of bots played a series of forty-thousand hands of poker with one another. Their objective was to maximize the overall bankroll—the total winnings. Detailed results appear in Table 1.

Four bots also competed in the series tournament: Hyperborean, GS2, BluffBot, and Teddy. Each pair of bots played a series of twelve-thousand hands of poker with one another. Each series was decided in favor of the bot with the higher winnings. The tournament objective was to maximize the number of series won. Detailed results appear in Table 2.

Each series consisted of duplicate match pairs. A match is a sequence of one-thousand hands of poker. A duplicate match pair is two matches, where in each match the same sequence of cards is dealt with the seats of the players reversed for symmetry. The performance of a bot in a duplicate match pair can have lower variance than the bot in two independent matches. After every match, the bots were reset. Thus, by dividing a series of forty-thousand hands into twenty duplicate match pairs, we can measure both performance and receive an estimate of the variance. The final results were:

	Bankroll Tournament	Series Tournament
1st place	Hyperborean (Win rate: 0.3925 small bets/hand)	Hyperborean (3 wins, 0 losses)
2nd place	BluffBot (Win rate: 0.0954 small bets/hand)	BluffBot (2 wins, 1 loss)
3rd place	Monash-BPP (Win rate: -0.0273 small bets/hand)	GS2 (1 win, 2 losses)

	Hyperborean	BluffBot	Monash BPP	Teddy
Hyperborean		0.0514 ± 0.0171	0.7227 ± 0.0161	0.4067 ± 0.0247
BluffBot	-0.0514 ± 0.0171		0.5271 ± 0.0197	-0.1895 ± 0.0289
Monash BPP	-0.7227 ± 0.0161	-0.5271 ± 0.0197		1.1678 ± 0.0428
Teddy	-0.4067 ± 0.0247	0.1895 ± 0.0289	-1.1678 ± 0.0428	

Table 1: Details of the bankroll tournament. Positive indicates the row player won money, negative indicates the column player did. Results are in small bets/hand. Standard deviation is shown after the result.

	Hyperborean	BluffBot	GS2	Monash BPP
Hyperborean		0.1145 ± 0.0375	0.1843 ± 0.0079	0.7344 ± 0.0293
BluffBot	-0.1145 ± 0.0375		0.1200 ± 0.0408	0.5214 ± 0.0561
GS2	-0.1843 ± 0.0079	-0.1200 ± 0.0408		0.6513 ± 0.0398
Monash BPP	-0.7344 ± 0.0293	-0.5214 ± 0.0561	-0.6513 ± 0.0398	

Table 2: Details of the series tournament. Positive indicates the row player won money, negative indicates the column player did. Results are in small bets/hand. Standard deviation is shown after the result.

2. STATISTICAL SIGNIFICANCE

In the bankroll tournament, because of the fact that each duplicate match pair was repeated twenty times, we are very confident as to the accuracy of the estimates of the standard deviation and the means for the various series in this tournament (see Table 1).

In the series tournament, because of the fact that each duplicate match pair was only repeated six times, we have less confidence in the standard deviations and how well the central limit theorem applies.¹ Nonetheless, the standard deviations we did measure indicated that the results were significant, that is, that the estimate of the winning rate of each winning bot was several standard deviations above zero (see Table 2).

For more details, visit <http://www.cs.ualberta.ca/~pokert/>.

Acknowledgements

Christian Smith wrote a large portion of the code for the competition. Luke Duguid helped to debug the code once entries were submitted. Rod Johnson, Jon Martin, and John Bartoszewski set up and maintained the sixteen machines for the competition. Darse Billings helped with the original design of the competition rules.

3. REFERENCES

Billings, D., Burch, N., Davidson, A., Holte, R., Schaeffer, J., Schauenberg, T., and Szafron, D. (2003). Approximating Game-Theoretic Optimal Strategies for Full-Scale Poker. *Proceedings of the International Joint Conference on Artificial Intelligence*.

Gilpin, A. and Sandholm, T. (2006). A Competitive Texas Hold'em Poker Player Via Automated Abstraction and Real-Time Equilibrium Computation. *Proceedings of the National Conference on Artificial Intelligence*.

Korb, K. B., Nicholson, A. E., and Jitnah, N. (1999). Bayesian Poker. *Uncertainty in Artificial Intelligence*.

Salonen, T. (2006). The Official Website for BluffBot - Poker Robot for Heads-Up Limit Texas Hold'em. <http://www.bluffbot.com/>.

¹In this competition, very little learning was used by the agents: thus, the results of the duplicate match pairs were in and of themselves distributed in a Gaussian fashion.