2008

GS3 and Tartanian: Game Theory-based Heads-up Limit and No-limit Texas Hold'em Poker-playing Programs

Andrew Gilpin
Carnegie Mellon University

Tuomas W. Sandholm
Carnegie Mellon University

Troels Bjerre Sørensen
University of Aarhus

Published In
Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS), Demonstration Track..

GS3 and Tartanian: Game theory-based heads-up limit and no-limit Texas Hold'em poker-playing programs

(Demo Paper)

Andrew Gilpin
Computer Science Dept.
Carnegie Mellon University
Pittsburgh, PA, USA

Tuomas Sandholm
Computer Science Dept.
Carnegie Mellon University
Pittsburgh, PA, USA

Troels Bjerre Sørensen
Dept. of Computer Science
University of Aarhus
Åbogade 34, Århus, Denmark

ABSTRACT

We demonstrate two game theory-based programs for heads-up limit and no-limit Texas Hold'em poker. The first player, *GS3*, is designed for playing limit Texas Hold'em, in which all bets are a fixed amount. The second player, *Tartanian*, is designed for the no-limit variant of the game, in which the amount bet can be any amount up to the number of chips the player has. Both *GS3* and *Tartanian* are based on our potential-aware automated abstraction algorithm for identifying strategically similar situations in order to decrease the size of the game tree. *Tartanian*, in order to deal with the virtually infinite strategy space of no-limit poker, in addition uses a discretized betting model designed to capture the most important strategic choices in the game. The strategies for both players are computed using our improved version of Nesterov's excessive gap technique specialized for poker.

In this demonstration, participants will be invited to play against both of the players, and to experience first-hand the sophisticated strategies employed by our players.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Miscellaneous; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

General Terms

Algorithms, Economics

Keywords

Equilibrium finding, automated abstraction, Nash equilibrium, computational game theory, sequential games, imperfect information games, Texas Hold'em poker

1. INTRODUCTION

Poker is a complex game involving elements of uncertainty, randomness, strategic interaction, and game-theoretic reasoning. Playing poker well requires the use of complex, intricate strategies. Optimal play is far from straightforward, typically necessitating actions intended to misrepresent one's private information. For these reasons, and others, poker has been proposed as an AI challenge problem [1].

Cite as: GS3 and Tartanian: Game theory-based heads-up limit and no-limit Texas Hold'em poker-playing programs (Demo Paper), Andrew Gilpin, Tuomas Sandholm and Troels Bjerre Sørensen, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX.

Classic game tree search algorithms, such as minimax search with α-β-pruning, are not applicable to games like poker due to the presence of imperfect information; a situation cannot be solved based solely on information contained in the subtree emanating from the current situation. Rather, a game-theoretic equilibrium analysis of the entire game at once is needed. Although two-person zero-sum games with imperfect information are solvable in time polynomial in the size of the game tree, the standard solution algorithms scale neither in terms of time nor memory usage to massive real-world game trees, such as that of Texas Hold'em poker. The programs that we present in this demonstration are a product of the recent research we have conducted in making game-theoretic analysis tractable for real-world sized games.

In this demonstration, participants will be encouraged to play heads-up Texas Hold'em poker against our programs *GS3* [8] and *Tartanian* [9]. Section 2 of this short paper describes the rules of limit and no-limit Texas Hold'em poker. Section 3 overviews the component technologies we have developed that were used in the construction of our poker-playing programs. Section 4 briefly previews the actual demonstration.

2. HEADS-UP TEXAS HOLD'EM POKER

There are many variants of poker. In this demonstration, we will showcase our programs for playing heads-up (i.e., two-player) Texas Hold'em poker, in both the limit structure and no-limit structure. The rules for the no-limit variant are as follows.

Blinds Two players, the *small blind* and *big blind*, start every hand with 1000 chips. Before any cards are dealt, the small blind contributes one chip to the pot and the big blind contributes two chips.

Pre-flop Each player receives two *hole cards*, face down, from a deck of 52 cards. The small blind can then *fold* (ending the game and yielding all of the chips in the pot to the other player), *call* (contribute one more chip), or *raise* (call one more chip and add two or more chips to the pot). In the event of a call or a raise, the big blind has the option to take an action. The players alternate playing in this manner until either one of the players folds or calls. It is possible for a player to go *all in* at any point by raising all of their remaining chips.

Flop Three *community cards* are dealt face up. The players then participate in a second betting round, with the big blind going first. Bets must be at least two chips.

Turn One community card is dealt face up. The players again participate in a betting round as on the flop.

River A final community card is dealt face up. The players participate in a betting round as on the flop and turn.

Showdown Once the river betting round has concluded (and if neither player has folded), a *showdown* occurs. Both players form the best five-card poker hand using their two hole cards and the five community cards. The player with the best hand wins the chips in the pot. In the event of two equally ranked hands, the players split the pot.

The above description is for *no-limit* poker. In *limit* poker, there is a single legal amount that can be bet at any time, and the number of bets each player can make in a single round is bounded at 4. Otherwise, the rules are the same.

3. OVERVIEW OF OUR PLAYERS' COMPONENT TECHNOLOGIES

We developed two agents for playing heads-up Texas Hold'em. *GS3* [8] was designed for playing limit and *Tartanian* [9] was designed for playing no-limit. We constructed the players using three conceptually separate component technologies:

1. **Automated card abstraction.** Given how large the game tree is, it is necessary to abstract nature's moves (i.e., the dealing of the cards). Our recent research has introduced abstraction algorithms for automatically reducing the state-space of the game in such a way that strategically similar states are collapsed into a single state. This can result in a significant decrease in problem size with little loss in solution quality. Here we applied our *potential-aware* automated abstraction algorithm [8], which improves upon our previous abstraction algorithms [7, 4, 6].

2. **Equilibrium finding.** Two-person zero-sum games can be modeled and solved as linear programs using the simplex algorithm or interior-point methods. However, those algorithms do not scale to games of this size. Recently, we developed new gradient-based algorithms that scale to games many orders of magnitude larger than what was previously possible [10, 2]. We applied these new algorithms to our problem, and developed a system for automatically constructing the source code for computing the crucial part of the equilibrium computation directly from a description of the game [9].

3. **Discretized betting model.** In no-limit poker, a player may bet any quantity up to the amount of chips she has remaining. Therefore, in principle, the betting action space is infinite. Even if players are restricted to betting integral amounts of chips (as is the case in most brick-and-mortar casinos), the number of actions available is huge. (The small blind has nearly 1000 actions available at the time of the first action.) To deal with this huge strategy space, we use our recently developed *discretized betting model* [9]. This also entails a *reverse model* for mapping the opponent's actions—which might not abide to the discretization—into the game model.

4. DEMONSTRATION DESCRIPTION

In this demonstration, participants will be encouraged to play heads-up limit Texas Hold'em against *GS3*, and heads-up no-limit Texas Hold'em against *Tartanian*. The authors will be on hand to demonstrate particularly interesting aspects of the programs' play, including strategies such as bluffing, slow-playing, check-raising, semi-bluffing, and other sophisticated poker strategies that are often attributed to human psychology, but actually are simply consequences of the game-theoretic analysis of the game.

Based on our experience in previous public demonstrations of our game-playing technologies [3, 5], we expect the demonstration to be quite interactive, informative, and entertaining for the participants.

5. REFERENCES

[1] D. Billings, A. Davidson, J. Schaeffer, and D. Szafron. The challenge of poker. *Artificial Intelligence*, 134(1-2):201–240, 2002.

[2] A. Gilpin, S. Hoda, J. Peña, and T. Sandholm. Gradient-based algorithms for finding Nash equilibria in extensive form games. In *3rd International Workshop on Internet and Network Economics (WINE)*, San Diego, CA, 2007.

[3] A. Gilpin and T. Sandholm. Optimal Rhode Island Hold'em poker. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pages 1684–1685, 2005. Intelligent Systems Demonstration.

[4] A. Gilpin and T. Sandholm. A competitive Texas Hold'em poker player via automated abstraction and real-time equilibrium computation. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2006.

[5] A. Gilpin and T. Sandholm. A Texas Hold'em poker player based on automated abstraction and real-time equilibrium computation. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1453–1454, Hakodate, Japan, 2006. Demonstration Track.

[6] A. Gilpin and T. Sandholm. Better automated abstraction techniques for imperfect information games, with application to Texas Hold'em poker. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2007.

[7] A. Gilpin and T. Sandholm. Lossless abstraction of imperfect information games. *Journal of the ACM*, 54(5), 2007.

[8] A. Gilpin, T. Sandholm, and T. B. Sørensen. Potential-aware automated abstraction of sequential games, and holistic equilibrium analysis of Texas Hold'em poker. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2007.

[9] A. Gilpin, T. Sandholm, and T. B. Sørensen. A heads-up no-limit Texas Hold'em poker player: Discretized betting models and automatically generated equilibrium-finding programs. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2008.

[10] S. Hoda, A. Gilpin, and J. Peña. A gradient-based approach for computing Nash equilibria of large sequential games. Available at http://www.optimization-online.org/, 2007.