TABLE OF CONTENTS

POKER

Graham Kendall and Jonathan Schaeffer

For many years Chess (and perhaps more recently Go) has served as the *Drosophila* of AI research. Decades of research culminated in the defeat of Garry Kasparov by DEEP BLUE in May 1997. There is still an active research community that uses Chess as a test-bed for AI research (as seen in this journal), but the game is limited in the types of challenges that it can offer to the AI researcher. Being a game of *perfect information* (both players know the full state of the game at any given point) with a relatively small branching factor, researchers have reduced the challenge of building a strong AI for Chess to merely one of deep brute-force search. The research challenges are to create a good evaluation function, and to design an effective search algorithm. This "solution" to Chess is unappealing to many AI purists. Nevertheless, alternative AI approaches have been largely ineffective.

Poker, as an experimental test-bed for exploring AI, is a much richer domain than Chess (and Go).

1. *Imperfect information*. Parts of the game state (opponent hands) are not known.
2. *Multiple players*. Many popular poker variants can be played with up to 10 players.
3. *Stochastic*. The dealing of the cards adds a random element to the game.
4. *Deception*. Predictable play can be exploited by an opponent. Hence, deceptive play is an essential ingredient of strong play (e.g., bluffing).
5. *Opponent modelling*. Observing your opponent(s) and adjusting your play to exploit (perceived) opponent tendencies is necessary to maximize poker winnings.
6. *Information sparsity*. Many poker hands end in the players not revealing their cards. This limits the amount of data available to learn from.

In Poker, evaluation functions and effective search algorithms are also important issues. However, the six additional dimensions of Poker described above introduce new complexities to the problem-solving process.

Perfect-information domains are exceptions rather than the rule in the real world. The richness of Poker means that it is a much superior domain than Chess (and Go) for exploring the issues needed to achieve the ultimate goal of human-level AI.

The past three years have seen an enormous increase in the popularity of Poker worldwide. Along with this has come a strong interest in building strong computer poker-playing programs. As evidence of this, in July 2006 the American Association for Artificial Intelligence hosted its first computer poker tournament; the 2007 edition of this event has already been confirmed. Given the popularity of the game, the interest in building strong programs, and the relevance of the research to the broader AI community, it seems timely to have a special issue of the International Computer Games Association Journal devoted to this challenging game.

In this issue, two research papers are presented. Both papers use the game of Texas Hold'em to illustrate their research. Texas Hold'em is the most popular poker variant played worldwide, in part because of the high strategic complexity of the play. In limit Hold'em, the size of a bet is fixed. In no-limit Hold'em, a player can bet any amount, adding yet another dimension to the problem-solving complexity of the game. The rules of Texas Hold'em are simple, and they can be found at numerous web sites.

A Tool for the Direct Assessment of Poker Decisions (Darse Billings and Morgan Kan, University of Alberta, Canada) presents a new evaluation tool for assessing the performance of two-player Poker. Given the stochastic element in Poker, a player may win a match (win the most money) due to luck rather than skill. Hence many tens of thousands of poker hands may have to be played to properly identify the better player. The DIVAT method proposed by Billings and Kan is able to remove most of the "luck" element from the assessment process, thereby allowing one to obtain a meaningful result by playing significantly fewer hands. The ideas are presented for limit Texas Hold'em, but they are general over all variants of Poker.

The second paper, *Pseudo-Optimal Strategies in No-Limit Poker* (Rickard Andersson, University of Umea, Sweden), considers the challenge of building a strong no-limit Texas Hold'em player. Two-player limit Hold'em has a search space of $O(10^{18})$, with no-limit being even larger (depending on the size of the player's stack of chips). Andersson extends the University of Alberta's game-theory approach to limit Poker to the no-limit domain. The paper demonstrates that the strength of the resulting computer poker player is a function of the amount of chips (money) that it has to play with.

We hope you enjoy the papers, and find them stimulating.

Finally, we would like to thank the referees of this special issue. They were:
- Luigi Barone
- Darse Billings
- Kevin Burns
- Ann Nicholson
- Eric Saund
- Lyndon While
- Michael Littman
- Brian Sheppard

ICGA Journal readers who are interested in information on our publications are referred to our website. A complete list of all articles, notes, and literature reviews published in the *ICCA Journal* and the *ICGA Journal* is accessible on the Internet at http://www.icga.org

The credits of the photographs in this issue are to: Frank Berger, Michael Greenspan, Ryan Haward, I-Chen Wu, Shi-Jim Yen, Creative Services, University of Alberta.