COMPUTER GAMES WORKSHOP AT IJCAI 2013

Mark Winands[1] *and Tristan Cazenave*[2]

Maastricht, The Netherlands Paris, France

From August 6 to August 9, the Twenty-third International Conference on Artificial Intelligence was held in Beijing China. In the three days preceding, 36 workshops were additionally held. On the first day (Saturday, August 3) the Computer Games Workshop was held, chaired by Tristan Cazenave, Mark Winands and Hiroyuki Iida at Tsinghua University. A total of 15 papers were submitted: 10 were accepted and 5 were rejected. Moreover, 20 participants registered for the event. We will briefly report on the 9 presented papers below.

In the first talk, Bruno Bouzy presented his paper *Monte-Carlo Fork Search for Cooperative Path-Finding.* Monte-Carlo Fork Search (MCFS) is a new algorithm that solves Cooperative Path-Finding (CPF) problems with simultaneity. Its background is Monte-Carlo Tree Search (MCTS) and Nested Monte-Carlo Search (NMCS). The key idea of MCFS is to build a search tree balanced over the whole game tree. After a simulation, MCFS stores the whole sequence of actions in the tree, which enables MCFS to fork new sequences at any depth in the built tree. The algorithms is suited for CPF problems in which the branching factor is too large for MCTS or A*, and in which congestion may arise at any distance from the initial state. With sufficient time and memory, Nested MCFS (NMCFS) solves congestion problems in the literature finding better solutions than the state-of-the-art solutions. It also solves N-puzzles without hole near-optimally.

Subsequently, the paper *Building Large Compressed PDBs for the Sliding Tile Puzzle*, written by Robert Döbbelin, Thorsten Schütt, and Alexander Reinefeld, was presented. The authors computed 9-9-6, 9-8-7, and 8-8-8 Pattern Databases (PDB) for the 24-puzzle that are three orders of magnitude larger (up to 1.4 TB) than the 6-6-6-6 PDB. This was possible by performing a parallel breadth-first search in the compressed pattern space. Their experiments indicate an average 8-fold improvement of the 9-9-6 PDB over the 6-6-6-6 PDB for the 24-puzzle.

In the next talk, *Improving Performance in Imperfect-Information Games with Large State and Action Spaces by Solving Endgames*, written by Sam Ganzfried and Tuomas Sandholm, it was demonstrated that endgame solving can be successful in practice in imperfect-information games with large state and action spaces despite the fact that the strategy profile it computes is not guaranteed to be an equilibrium in the full game. Their experiments in two-player no-limit Texas Hold'em poker showed that their approach leads to a significant performance improvement.

After the coffee break Marc Lanctot presented joint research with Viliam Lisý and Mark Winands. Their paper *Monte Carlo Tree Search in Simultaneous Move Games with Applications to Goofspiel* discussed the adaptation of MCTS to simultaneous move games with and without chance events. They introduced a new algorithm, Online Outcome Sampling (OOS), which approaches a Nash equilibrium strategy over time. They compared both head-to-head performance and exploitability of several MCTS variants in Goofspiel. The result revealed that regret matching and OOS performed best and that all variants produced less exploitable strategies than UCT.

In the last talk of the session, *Decision Trees for Computer Go Features*, Francois van Niekerk presented his M.Sc. research, supervised by Steve Kroon. Their paper investigated the feasibility of using decision trees to generate features for guiding MCTS in Computer Go. The experiments showed that while this approach exhibits potential, their initial prototype is not as powerful as using traditional pattern features.

After lunch Nathan Sturtevant presented the paper *UCT Enhancements in Chinese Checkers Using an Endgame Database*, a joint collaboration with Max Roschke. They assessed the performance of MCTS-based AIs and the effectiveness of augmenting them with a lookup table containing evaluations of games states in the game of

[1] Games and AI Group, Faculty of Humanities and Sciences, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. Email: m.winands@maastrichtuniversity.nl.

[2] LAMSADE, Université Paris-Dauphine, Paris, France. Email: cazenave@lamsade.dauphine.fr.

Chinese Checkers. The lookup table is only guaranteed to be correct during the endgame, but serves as an accurate heuristic throughout the game. Experiments showed that using the lookup table only for its endgames is harmful, while using it for its heuristic values improves the quality of play. The research was performed on a board with 81 locations and 6 pieces, which is larger than previous work on lookup tables in Chinese Checkers. It is a precursor to using the 500GB full-game single-agent data on the full-size board with 81 locations and 10 pieces

Yuichiro Sato discussed *Automated Generation of New Concepts from General Game Playing*, joint research with Tristan Cazenave. He described how to extract explicit concepts from heuristic functions obtained using a simulation based approach. The proposed algorithm quickly learns new concepts without any supervision but from experience in the environment. Concepts to understand the semantics of Tic-tac-toe are generated by their approach. These concepts are also available to understand the semantics of Connect Four. The authors conclude that their approach is applicable to General Game Playing and is able to extract explicit concepts, which are able to be understood by humans.

The next presentation was entitled *WALTZ: A Strong Tzaar-Playing Program*, written by Tomáš Valla and Pavel Veselý. The authors introduced the game of Tzaar, part of the Project GIPF, to the AI community. It is an abstract strategy two-player game, which has recently gained popularity in the gaming community and has won several awards. The high branching factor makes Tzaar a difficult game for computers. The authors developed WALTZ a strong Tzaar-playing program, using enhanced variants of αβ and proof-number search. After many tests with computer opponents and a year of deployment on a popular board-gaming portal, the authors conclude that WALTZ can defeat all available computer programs and even strong human players.

In the last presentation *Perfectly Solving Domineering Boards*, Jos Uiterwijk presented his research in the game of Domineering. For this game he defined 12 knowledge rules, of increasing complexity. Of these rules, 6 can be used to show that the starting player (assumed to be Vertical) can win a game against any opposition, while 6 can be used to prove a definite loss (a win for the second player, Horizontal). Applying this knowledge-based method to all 81 rectangular boards up to 10×10 (omitting the trivial $1 \times n$ and $m \times 1$ boards), 67 could be solved perfectly. This is in sharp contrast with previous publications reporting the solution of Domineering boards, where only a few tiny boards were solved perfectly, the remainder requiring up to large amounts of search. Applying this method to larger boards with one or both sizes up to 30 solves 216 more boards, mainly with one dimension odd. All results fully agree with previously reported game-theoretic values. Moreover, among the larger boards solved 4 were not solved before, i.e., 10×11, 14×11, 18×6, and 18×11, all being a win for Vertical.

The workshop was enjoyed by all participants. The workshop proceedings will be published with Springer in their Communications in Computer and Information Science series.