

## IJCAI Computer Games Workshop 2015

Mark Winands<sup>a,\*</sup> and Tristan Cazenave<sup>b</sup>

<sup>a</sup>*Games and AI Group, Department of Data Science and Knowledge Engineering, Maastricht University, Maastricht, The Netherlands*

<sup>b</sup>*LAMSADE, Université Paris-Dauphine, Paris, France*

From July 28 to July 31, the 24th International Conference on Artificial Intelligence was held in Buenos Aires, Argentina. In the three days preceding, 33 workshops were additionally organized. On the second day (Sunday, July 26) the Computer Games Workshop took place, chaired by Tristan Cazenave, Mark Winands, and Stefan Edelkamp at the Facultad de Ciencias Económicas. A total of 16 papers were submitted: 10 papers were accepted for oral presentation and 6 were rejected. Moreover, 35 participants attended this event. We briefly report on the accepted papers below.

In the first talk *Argumentative AI Director using Defeasible Logic Programming*, Ramiro Agis, Andrea Cohen, and Diego Martínez presented a novel implementation of an AI Director that uses argumentation techniques to decide dynamic adaptations in the level generation of a roguelike game called HermitArg. The architecture of the game introduces smart items with defeasible: information to be analysed in a dialectical process.

The next talk, *Path Planning with Inventory-Driven Jump-Point-Search*, written by Davide Aversa, Stavros Vassos, and Sebastian Sardina, discussed the issue that for many navigational domains it is common that the traversability of cells is conditioned on the path taken. This happens for instance in video games, where the agent has to acquire a certain object (i.e., a key or a flying suit) to be able to traverse specific locations (e.g., doors or high walls). To handle such scenarios the authors presented InvJPS, an “inventory-driven” pathfinding approach based on the highly successful grid-based Jump-Point-Search (JPS) algorithm. They showed, formally and experimentally, that the new algorithm preserves JPS’s optimality guarantees and its symmetry breaking advantages in inventory-based variants of game maps.

Subsequently, Bruno Bouzy presented his paper *An Experimental Investigation on the Pancake Problem*. This problem is NP-hard and linked to the genome rearrangement problem also called sorting by reversals (SBR). Until now, the best theoretical R-approximation was 2 with an algorithm which gives a 1.22 experimental R-approximation on stacks with a size smaller than 70. In this paper a Monte-Carlo Search (MCS) approach with nested levels and specific domain-dependent simulations is used. Bruno Bouzy showed that MCS is an alternative to Iterative Deepening Depth First Search (IDDFS) for sorting large stacks of pancakes. At a given level and with a given number of polynomial-time domain-dependent simulations, MCS is a polynomial-time algorithm as well. MCS at level 3 gives a 1.04 experimental R-approximation, which is a breakthrough. At level 1, MCS solves stacks of size 512 with an experimental R-approximation value of 1.20.

---

\*Corresponding author: Mark Winands, Games and AI Group, Department of Data Science and Knowledge Engineering, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. Tel.: +31 433882728; Fax: +31 43 38 84910; E-mail: m.winands@maastrichtuniversity.nl.

After the coffee break Tristan Cazenave proposed a predator-prey game in the paper *Multi-Agent Retrograde Analysis*. This domain is modelled as a board game where three predators are trying to capture a prey. Each agent has five possible moves: going up, down, left, right or staying at the same location. The game terminates if the prey is at the same location as a predator or if the prey cannot move to an empty location. Small boards up to  $9 \times 9$  have been solved using retrograde analysis. The outcome is that the predator-prey game is always lost for the prey when there are at least 3 predators.

Next, Marc Chee discussed *A Principled Approach to the Problem of Chunking in UCT*, joint research with Abdallah Saffidine and Michael Thielscher. He presented a solution to the problem of receiving results in chunks during the UCT search. Chunking can occur when using a transposition table or with some parallelisation techniques. It is important to be able to use this data without it adversely affecting the search. Marc Chee described three techniques that allow aggregating chunked data and propagating it in such a way that adheres to UCT without a loss of information. Additionally, it was shown how this accelerates the convergence of a game player.

In the last talk of the session, *Learning to Trade in Strategic Board Games*, written by Heriberto Cuayáhuitl, Simon Keizer, and Oliver Lemon, discussed a data-driven approach for automatic trading in the game of Settlers of Catan. Their experiments are based on data collected from human players trading in text-based natural language. The performance of Bayesian Networks, Conditional Random Fields, and Random Forests have been compared in the task of ranking trading offers, and are evaluated both in an offline setting and online while playing the game against a rule-based baseline. Experimental results show that agents trained from data from average human players can outperform rule-based trading behaviour, and that the Random Forest model achieves the best results.

The first presentation after lunch was entitled *485 – A New Upper Bound for Morpion Solitaire*, a joint collaboration of Henryk Michalewski, Andrzej Nagórko, and Jakub Pawlewicz. The authors showed a new upper bound of 485 moves for the 5T variant of the Morpion Solitaire game. This is achieved by encoding Morpion 5T rules as a linear program and solving 126,912 instances of this program on special octagonal boards. To show the correctness of this method the rules of the game have been analysed and the potential of a given position has been used. By solving continuous-valued relaxations of linear programs on these boards, an upper bound of 586 moves is obtained. Further analysis of original, not relaxed, mixed-integer programs leads to an improvement of this bound to 485 moves. However, this is achieved at a significantly higher computational cost.

Next, Mark Winands presented the paper *Sequential Halving for Partially Observable Games*, a joint collaboration with Tom Pepels and Tristan Cazenave. They investigated Sequential Halving as a selection policy in the following four partially observable games: Go Fish, Lost Cities, Phantom Domineering, and Phantom Go. Additionally, H-MCTS is studied, which uses Sequential Halving at the root of the search tree, and UCB elsewhere. Experimental results reveal that H-MCTS performs the best in Go Fish, whereas its performance is on a par in Lost Cities and Phantom Domineering. Sequential Halving as a flat Monte-Carlo Search appears to be the stronger technique in Phantom Go. Before the last coffee break, Nathan Sturtevant gave the talk *Challenges and Progress on Using Large Lossy Endgame Databases in Chinese Checkers*. He discussed using large endgame databases to improve the performance of minimax and MCTS-based agents in Chinese Checkers. Several challenges are faced in how to properly integrate the endgame databases and how to correct errors that occur because of the compression that is used when storing the endgame data. Preliminary experimental results suggest minimax-based approaches are able to do a better job of using the endgame data than MCTS approaches. However, these results could still change with further study.

In the last talk of the workshop Mark Winands presented the paper *The Surakarta Bot Revealed*. The ideas behind the agent SIA, which won the Surakarta tournament five times at the ICGA Computer Olympiads, are revealed. During his talk he described SIA's  $\alpha\beta$ -based variable-depth search mechanism. Enhancements such as Multi-Cut forward pruning and Realization Probability Search improve the agent considerably. Additionally, features of the static evaluation function were presented. Experimental results indicate that features, which reward distribution of the pieces and penalize pieces that clutter together, give a genuine improvement in the playing strength.

The workshop was enjoyed by all participants. The papers will be published with Springer in their Communications in Computer and Information Science series. The volume will also contain the best papers of the IJCAI workshop on General Intelligence in Game-Playing Agents (GIGA'15), which was held a day after.