

Preface

These proceedings contain the papers of the Computer Games Workshop (CGW 2014) held in Prague, Czech Republic. The workshop took place August 18, 2014, in conjunction with the 21st European Conference on Artificial Intelligence (ECAI 2014). The workshop received 20 submissions. Each paper was sent to two reviewers. In the end, 12 papers were accepted for presentation at the workshop, of which 11 made it into these proceedings. The Computer and Games Workshop series is an international forum for researchers interested in all aspects of artificial intelligence and computer game playing. Earlier workshops took place in Montpellier, France (2012), and Beijing, China (2013).

The published papers cover a wide range of topics related to computer games. They collectively discuss 11 abstract games: 7 Wonders, Amazons, AtariGo, Ataxx, Breakthrough, Chinese Dark Chess, Connect6, NoGo, Pentath, Othello, and Catch the Lion. Moreover, two papers are on General Game Playing, and four on video game playing. Below we provide a brief outline of the contributions, in the order in which they appear in the proceedings.

“Minimizing Simple and Cumulative Regret in Monte-Carlo Tree Search,” a joint collaboration by Tom Pepels, Tristan Cazenave, Mark Winands, and Marc Lanctot. In the paper a new MCTS variant, called Hybrid MCTS (H-MCTS), is introduced that minimizes cumulative and simple regret in different parts of the tree. H-MCTS uses SHOT, a recursive version of Sequential Halving, to minimize simple regret near the root, and UCT to minimize cumulative regret when descending further down the tree. The results show genuine performance increase in Amazons, AtariGo, and Breakthrough.

“On Robustness of CMAB Algorithms: Experimental Approach,” authored by Antonín Komenda, Alexander Shleyfman, and Carmel Domshlak experimentally analyzes the robustness of two state-of-the-art algorithms, Naive Monte Carlo (NMC) and Linear Side-Information (LSI), for online planning with combinatorial actions of the turn-based variant of the strategy game μ RTS. The results show that LSI is stronger with smaller budgets and shorter look-ahead.

“Job-Level Algorithms for Connect6 Opening Position Analysis,” by Ting-Han Wei, I-Chen Wu, Chao-Chin Liang, Bing-Tsung Chiang, Wen-Jie Tseng, Shi-Jim Yen, and Chang-Shing Lee, investigates job-level (JL) algorithms to analyze opening positions for Connect6. The paper first proposes four heuristic metrics when using JL-PNS to estimate move quality. Next, it introduces a JL Upper Confidence Tree (JL-UCT) algorithm and heuristic metrics, one of which is the number of nodes in each candidate move’s subtree. In order to compare these metrics objectively, the paper proposes two kinds of measurement methods to analyze the suitability of these metrics when choosing best moves for a set of benchmark positions. The results show that for both metrics this node count

heuristic metric for JL-UCT outperforms all the others, including the four for JL-PNS.

“Monte-Carlo Tree Search and Minimax Hybrids with Heuristic Evaluation Functions,” written by Hendrik Baier and Mark Winands, discusses three different approaches to employ minimax search with static evaluation functions in MCTS: (1) to choose moves in the play-out phase of MCTS, (2) as a replacement for the play-out phase, and (3) as a node prior to bias move selection. The MCTS-minimax hybrids are tested and compared with their counterparts using evaluation functions without minimax in the domains of Othello, Breakthrough, and Catch the Lion. Results show that introducing minimax search is effective for heuristic node priors in Othello and Catch the Lion. The MCTS-minimax hybrids are also found to work well in combination with each other.

“Monte-Carlo Tree Search for the Game of ‘7 Wonders’,” written by Denis Robilliard, Cyril Fonlupt, and Fabien Teytaud studies MCTS in the game of 7 Wonders. This card game combines several known challenging properties, such as imperfect information, multi-player, and chance. It also includes an inter-player trading system that induces a combinatorial search to decide which decisions are legal. Moreover, it is difficult to build an efficient evaluation function because the card values are heavily dependent upon the stage of the game and upon the other player decisions. The paper discusses how to effectively apply MCTS to 7 Wonders.

“Small and Large MCTS Playouts Applied to Chinese Dark Chess Stochastic Game,” by Nicolas Jouandeau and Tristan Cazenave, presents MCTS modifications to deal with the stochastic game of Chinese Dark Chess. Experiments are conducted with group nodes and chance nodes using various configurations: with different play-out policies, with different play-out lengths, with true or estimated wins. Results show that extending the play-out length is useful for creating more informed play-outs, and the usage of an evaluation function can increase or decrease player’s effectiveness through modifying the number of draw possibilities.

“On the Complexity of General Game Playing,” authored by Édouard Bonnet and Abdallah Saffidine, discusses the computational complexity of reasoning in General Game Playing (GGP) using various combinations of multiple features of the Game Description Language (GDL). Their analysis offers a complexity landscape for GGP with fragments ranging from NP to EXPSPACE in the single-agent case, and from PSPACE to 2-EXPTIME in the multi-agent case.

“Efficient Grounding of Game Descriptions with Tabling,” by Jean-Noël Vitaut and Jean Méhat, presents a method to instantiate game descriptions used in GGP with the tabling engine of a Prolog interpreter. Instantiation is a crucial step for speeding up the interpretation of the game descriptions and increasing the playing strength of general game players. The method allows one to ground almost all of the game descriptions present on the GGP servers in a time that is compatible with the common time settings of the GGP competition. It instantiates descriptions more rapidly than previous published methods.

“SHPE: HTN Planning for Video Games,” written by Alexandre Menif, Éric Jacopin, and Tristan Cazenave, describes SHPE (Simple Hierarchical Planning

Engine). It is a hierarchical task network planning system designed to generate dynamic behaviors for real-time video games. SHPE is based on a combination of domain compilation and procedural task application/decomposition techniques in order to compute plans in a very short time-frame. The planner is able to return relevant plans in less than three milliseconds for several problem instances of the *SimpleFPS* planning domain.

“Predicting Player Disengagement in Online Games,” by Hanting Xie, Daniel Kudenko, Sam Devlin, and Peter Cowling, introduces a pure data-driven method to foresee whether players will quit the game given their previous activity within the game, by constructing decision trees from historical gameplay data of previous players. The method is assessed on two popular commercial online games: I Am Playr and Lyroke. The former is a football game while the latter is a music game. The results indicate that the decision tree built by their method is valuable for predicting the players’ disengagement and that its human-readable form allow us to search out further reasons about which in-game events made them quit.

“Coordinating Dialogue Systems and Stories Through Behavior Composition,” a joint effort by Stefano Cianciulli, Daniele Riccardelli, and Stavros Vassos, exploits behavior composition in AI as a formal tool for facilitating interactive storytelling in video games. This is motivated by (1) the familiarity of transition systems in video game development, and (2) the fact that behavior composition extends the spectrum of approaches for non-linear storylines by introducing a new paradigm based on planning for a target desired process instead of a goal state. Moreover, the approach provides support for the debugging of deadlocks in stories at design level. The paper describes the behavior composition framework, and shows the details for an interactive dialogue system scenario in order to illustrate how interactive storytelling can be phrased in terms of the framework. A simple architecture for implementing a demo game over the scenario using existing behavior composition tools is also reported.

These proceedings would not have been produced without the help of many persons. In particular, we would like to mention the authors and reviewers for their help. Moreover, the organizers of ECAI 2014 contributed substantially by bringing the researchers together.

November 2014

Tristan Cazenave
Mark Winands
Yngvi Björnsson

Organization

Program Chairs

Tristan Cazenave	Université Paris-Dauphine, France
Mark Winands	Maastricht University, The Netherlands
Yngvi Björnsson	Reykjavik University, Iceland

Program Committee

Yngvi Björnsson	Reykjavik University, Iceland
Bruno Bouzy	Université Paris-Descartes, France
Tristan Cazenave	Université Paris-Dauphine, France
Rémi Coulom	Université Lille 3, France
Stefan Edelkamp	University of Bremen, Germany
Nicolas Jouandeau	Université Paris 8, France
Peter Kissmann	University Bremen, Germany
Sylvain Lagrue	Université d'Artois, France
Marc Lanctot	Maastricht University, The Netherlands
Viliam Lisý	Czech Technical University in Prague, Czech Republic
Jean Méhat	Université Paris 8, France
Jochen Renz	The Australian National University, Australia
Abdallah Saffidine	University of New South Wales, Australia
Fabien Teytaud	Université du Littoral Côte d'Opale, France
Olivier Teytaud	Université Paris-Sud, France
Mark Winands	Maastricht University, The Netherlands

Additional Reviewers

Tom Pepels	Maastricht University, The Netherlands
Stephan Schiffel	Reykjavik University, Iceland
Tsan-sheng Hsu	Institute of Information Science, Academia Sinica, Taiwan

Table of Contents

Minimizing Simple and Cumulative Regret in Monte-Carlo Tree Search	1
<i>Tom Pepels, Tristan Cazenave, Mark H.M. Winands, and Marc Lanctot</i>	
On Robustness of CMAB Algorithms: Experimental Approach	16
<i>Antonín Komenda, Alexander Shleyfman, and Carmel Domshlak</i>	
Job-Level Algorithms for Connect6 Opening Position Analysis	29
<i>Ting-Han Wei, I-Chen Wu, Chao-Chin Liang, Bing-Tsung Chiang, Wen-Jie Tseng, Shi-Jim Yen, and Chang-Shing Lee</i>	
Monte-Carlo Tree Search and Minimax Hybrids with Heuristic Evaluation Functions	45
<i>Hendrik Baier and Mark H.M. Winands</i>	
Monte-Carlo Tree Search for the Game of “7 Wonders”	64
<i>Denis Robilliard, Cyril Fonlupt, and Fabien Teytaud</i>	
Small and Large MCTS Playouts Applied to Chinese Dark Chess Stochastic Game	78
<i>Nicolas Jouandeau and Tristan Cazenave</i>	
On the Complexity of General Game Playing	90
<i>Édouard Bonnet and Abdallah Saffidine</i>	
Efficient Grounding of Game Descriptions with Tabling	105
<i>Jean-Noël Vittaut and Jean Méhat</i>	
SHPE: HTN Planning for Video Games	119
<i>Alexandre Menif, Éric Jacopin, and Tristan Cazenave</i>	
Predicting Player Disengagement in Online Games	133
<i>Hanting Xie, Daniel Kudenko, Sam Devlin, and Peter Cowling</i>	
Coordinating Dialogue Systems and Stories through Behavior Composition	150
<i>Stefano Cianciulli, Daniele Riccardelli, and Stavros Vassos</i>	
Author Index	165

