PREFACE

These proceedings contain the papers of the Computer Games Workshop (CGW 2017) held in Melbourne, Australia. The workshop took place on August 20, 2017, in conjunction with the 26th International Conference on Artificial Intelligence (IJCAI 2017). The Computer and Games Workshop series is an international forum for researchers interested in all aspects of artificial intelligence (AI) and computer game playing. Earlier workshops took place in Montpellier, France (2012), Beijing, China (2013), Prague, Czech Republic (2014), Buenos Aires, Argentina (2015) and New York, USA (2016). For the sixth edition of the Computer Games Workshop, 18 papers were submitted in 2017. Each paper was sent to three reviewers. In the end, 14 papers were accepted for presentation at the workshop, of which 12 made it into these proceedings.

The published papers cover a wide range of topics related to computer games. They collectively discuss six abstract games: Chinese Checkers, Chinese Dark Chess, Hex, Othello, Poker, and SameGame. Additionally, three papers deal with video games, two papers on General Game Playing and one discusses a web-based game. Here we provide a brief outline of the 12 contributions, in the order in which they appear in the proceedings.

"Atari Games and Intel Processors", a joint effort by Robert Adamski, Tomasz Grel, Maciej Klimek, and Henryk Michalewski, presents the results on learning strategies in Atari games using a Convolutional Neural Network (CNN), the Math Kernel Library, and TensorFlow framework. They also analyze effects of asynchronous computations on the convergence of reinforcement learning algorithms.

"Computer Hex using Move Evaluation Method based on Convolutional Neural Network", written by Kei Takada, Hiroyuki Iizuka, and Masahito Yamamoto, proposes a CNN model for the game of Hex. It evaluates all candidate moves by taking as input all sets consisting of 3 mutually adjacent cells. The proposed CNN model is tested against an existing CNN model called NeuroHex. The results show that their CNN model is superior to NeuroHex on the 13×13 board even though the CNN model is trained on the 11×11 board. The proposed model is also used as an ordering function and subsequently tested against the world-champion Hex program MoHex 2.0 on the 11×11 board. The results show that the proposed model can be used as a better ordering function than the ordering function created by minimax search. It obtains a win rate of 49.0% against MoHex 2.0, at 30 seconds per move.

"Deep Preference Neural Network for Move Prediction in Board Games", by Thomas Philip Runarsson, studies the training of deep neural networks for move prediction in board games for the game of Othello. The paper presents a general deep preference neural network. The problem of over-fitting becomes an immediate concern when training such a deep preference neural networks. The paper shows how dropout may combat this problem to a certain extent. It illustrates how classification test accuracy does not necessarily correspond to move accuracy and discussed the key difference between preference training versus single-label classification. The careful use of dropout coupled with richer game data produces an evaluation function that is a better move predictor but does not necessarily produce a stronger game player

"Deep Reinforcement Learning with Hidden Layers on Future States", written by Hirotaka Kameko, Jun Suzuki, Naoki Mizukami, and Yoshimasa Tsuruoka, presents a novel deep reinforcement learning architecture that can both effectively and efficiently use information on future states in video games. First, the authors demonstrate that such information is quite useful in deep reinforcement learning by using exact state transition information obtained from the emulator. They propose a method that predicts future states using Long Short Term Memory (LSTM), such that the agent can look ahead without the emulator. The paper applies their method to the asynchronous advantage actor-critic (A3C) architecture. The experimental results show that their proposed method with predicted future states substantially outperforms the vanilla A3C in several Atari games.

"Neural Fictitious Self-Play in Imperfect Information Games with Many Players" is authored by Keigo Kawamura, Naoki Mizukami, and Yoshimasa Tsuruoka. The paper proposes to use Neural Fictitious Self-Play (NFSP) to calculate approximate Nash equilibrium solutions for imperfect information games with more than two players. Although there are no theoretical guarantees of convergence for NFSP in such games, the authors empirically demonstrate that NFSP enables to calculate strategy profiles that are significantly less exploitable than random players in simple poker variants with three or more players.

"On-line Parameters Tuning for Monte-Carlo Tree Search in General Game Playing", by Chiara F. Sironi and Mark H. M. Winands, proposes a method to automatically tune search-control parameters on-line for General Game Playing (GGP). This method considers the tuning problem as a Combinatorial Multi-Armed Bandit (CMAB). Four strategies designed to deal with CMABs are evaluated for this particular problem. Experiments show that on-line tuning in GGP almost reaches the same performance as off-line tuning. It can be considered as a valid alternative for domains where off-line parameters tuning is costly or infeasible

“Memorizing the Playout Policy”, authored by Tristan Cazenave and Eustache Diemert, proposes a modification to Playout Policy Adaptation with move Features (PPAF). It consists of memorizing the learned policy from one move to the next. The paper tests PPAF with memorization (PPAFM) against PPAF and UCT for Atarigo, Breakthrough, Misere Breakthrough, Domineering, Misere Domineering, Knightthrough, Misere Knightthrough, and Nogo.

“Distributed Nested Rollout Policy for SameGame”, written by Benjamin Negrevergne and Tristan Cazenave, discusses Nested Rollout Policy Adaptation (NRPA), which is a Monte Carlo search heuristic for puzzles and other optimization problems. It achieves state-of-the-art performance on several games including SameGame. In this paper, the authors design several parallel and distributed NRPA-based search techniques, and provide a number of experimental insights about their execution. Finally, they use the best implementation to discover 15 new best scores for well-known SameGame boards.

“A Study of Forward versus Backwards Endgame Solvers with Results in Chinese Checkers”, by Nathan R. Sturtevant and Abdallah Saffidine investigates different approaches to endgame solvers based on retrograde analysis and describes a couple natural optimizations. The paper introduces a formal model of the state space to quantify the impact of these optimizations on the worst-case complexity of the solving algorithms. It performs also an empirical study on a small-size variant of Chinese Checkers.

“Validating and Fine-tuning of Game Evaluation Functions using Endgame Databases”, a joint effort by Hung-Jui Chang, Gang-Yu Fan, Jr-Chang Chen, Chih-Wen Hsueh, and Tsan-sheng Hsu, constructs a scheme to use the information from endgame databases to validate and fine-tune a manually designed evaluation function for Chinese Dark Chess. The method abstracts critical information from a huge database and validates the evaluation function on positions that are included in the endgame database. Using this information, the authors then discover meta knowledge to fine-tune and revise the evaluation function such that it better evaluates a position, even for the ones having many pieces. Experimental results show that their approach is successful.

“Applying Anytime Heuristic Search to Cost-Optimal HTN Planning”, a contribution by Alexandre Menif, Christophe Guettier, Eric Jacopin, and Tristan Cazenave, presents a framework for cost-optimal Hierarchical Task Network (HTN) planning. The framework includes an optimal algorithm combining a branch-and-bound with a heuristic search, which can also be used as a near-optimal algorithm given a time limit. It also includes different heuristics based on weighted cost estimations and different decomposition strategies. The different elements from this framework are empirically evaluated on three planning domains, one of which is modeling a First-Person Shooter game. The empirical results establish the superiority on some domains of a decomposition strategy that prioritizes the most abstract tasks. They also highlight that the best heuristic formulation for the three domains is computed from linear combinations of optimistic and pessimistic cost estimations.

*“*A Game for Eliciting Trust between People and Devices under Diverse Performance Conditions”, authored by Ingrid Zukerman, Andisheh Partovi, Kai Zhan, Nora Hamacher, Julie Stout, and Masud Moshtaghi, introduces a web-based game designed to investigate how different conditions affect people's trust in devices. The game is set in a retirement village, where residents live in smart homes equipped with monitoring systems. Players, who “work”' in the village, need to trade-off the time spent on administrative tasks (which enable them to earn extra income) against the time spent ensuring the welfare of the residents. The scenario of the game is complex enough to support the investigation of the influence of various factors, such as system accuracy, type of error made by the system, and risk associated with events, on players' trust in the monitoring system. In the paper, the authors describe the game and its theoretical underpinnings, and present preliminary results from a trial where players interacted with two systems that have different levels of accuracy.

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 IJCAI 2017 contributed substantially by bringing the researchers together.

November 2017,

Tristan Cazenave
Mark Winands
Abdallah Saffidine

Organization

Program Chairs

Tristan Cazenave, Université Paris-Dauphine, France

Mark Winands, Maastricht University, the Netherlands

Abdallah Saffidine, University of New South Wales, Australia

Program Committee

Michael Albert, University of Otago, New Zealand

Christopher Archibald, Mississippi State University, USA

Yngvi Björnsson, Reykjavik University, Iceland

Edouard Bonnet, Middlesex University, UK

Bruno Bouzy, Université Paris-Descartes, France

Tristan Cazenave, Université Paris-Dauphine, France

Stefan Edelkamp, University of Bremen, Germany

Ryan Hayward, University of Alberta, Canada

Amy Hoover, Northeastern University, USA

Hiroyuki Iida, JAIST, Japan

Nicolas Jouandeau, Université Paris 8, France

Jialin Liu, Queen Mary University of London, UK

Richard Lorentz, California State University, USA

Simon Lucas, Queen Mary University of London, UK

Jean Mehat, University Paris 8, France

Henryk Michalewski, University of Warsaw, Poland

Martin Müller, University of Alberta, Canada

Thomas Runarsson, University of Iceland, Canada

Abdallah Saffidine, University of New South Wales, Australia

Nathan Sturtevant, University of Denver, USA

Julian Togelius, New York University, USA

Mark Winands, Maastricht University, the Netherlands

Shi-Jim Yen, National Dong Hwa University, Taiwan

Additional Reviewers

Pascal Bercher
Chao Gao
Pim Nijssen
Maarten Schadd
Dennis Soemers
David Thue
Kenny Young