

Preface

This book contains the papers presented at the 15th Advances in Computer Games (ACG 2017) conference held in Leiden, the Netherlands. The conference took place during July 3–5, 2017, in conjunction with the 20th Computer Olympiad and the 23rd World Computer-Chess Championship.

The Advances in Computer Games conference series is a major international forum for researchers and developers interested in all aspects of artificial intelligence and computer game playing. Earlier conferences took place in London (1975), Edinburgh (1978), London (1981, 1984), Noordwijkerhout (1987), London (1990), Maastricht (1993, 1996), Paderborn (1999), Graz (2003), Taipei (2005), Pamplona (2009), Tilburg (2011), and Leiden (2015).

The Program Committee (PC) was pleased to see that so much progress was made in new games and that new techniques were added to the recorded achievements. In this conference, 23 papers were submitted. Each paper was sent to at least three reviewers. If conflicting views on a paper were reported, the reviewers themselves arrived at a final decision. The PC accepted 19 papers for presentation at the conference and publication in these proceedings. As usual, we informed the authors that they submitted their contribution to a post-conference editing process. The two-step process is meant (a) to give authors the opportunity to include the results of the fruitful discussion after the lecture in their paper, and (b) to maintain the high-quality threshold of the ACG series. The authors enjoyed this procedure.

The aforementioned set of 19 papers covers a wide range of computer games and many different research topics. We grouped the topics into the following four classes according to the order of publication: games and puzzles (seven papers), Go and chess (four papers), machine learning and MCTS (four papers), and (serious) gaming (four papers). The paper “Toward Solving “EinStein würfelt nicht!”” by François Bonnet and Simon Viennot received the Best Paper Award. In the proceedings, the award-winning paper is preceded by a paper from the same authors that can be read as a stepping stone for the second paper. For reference of self-containedness, the editors have allowed a similar introduction for both papers.

We hope that the readers will enjoy the research efforts presented by the authors. Here, we reproduce brief characterizations of the 19 contributions largely relying on the text as submitted by the authors. The idea is to show a connection between the contributions and insights into the research progress.

Games and Puzzles

The first paper, “Analytical Solution for “EinStein würfelt nicht!” with One Stone,” is written by François Bonnet and Simon Viennot. They discuss the board game “EinStein würfelt nicht!,” which is usually played on a 5×5 board with six stones per player and a die. In this contribution the authors study the game for the particular case when the players start with only one stone. In the case the random element from the use of a die disappears, an analytical analysis is possible. The authors describe and prove a winning strategy for the first (or second) player for all possible board sizes. In most cases, the first player can force a win, but depending on a precisely formulated condition on the board size, it is sometimes possible for the second player to win.

In their follow-up paper, “Toward Solving “EinStein würfelt nicht!,”” the same authors present an exact solution to some instances of the game, with fewer stones on smaller (or larger) boards. When the rules allow the players to choose their initial configuration, a solution consists in computing the exact optimal winning chances of the players for any initial configuration, and then computing the resulting Nash equilibrium between the two players. The most difficult result is the solution for a 4×4 board with six stones per player.

“Analysis of Fred Horn’s Gloop Puzzle,” written by Cameron Browne, presents the game of Gloop. It is a tile-based combinatorial puzzle game with a strong topological basis, in which the player is assigned a number of challenges to complete with a particular set of tiles. This paper describes the computer-based analysis of a number of representative Gloop challenges, including the computer-assisted solution of a difficult problem that had stood for over a decade.

“Set Matching: An Enhancement of the Hales–Jewett Pairing Strategy,” by Jos Uiterwijk, discusses the Hales–Jewett pairing strategy for solving k -in-a-Row games. It is a well-known strategy for proving that specific positions are (at most) a draw. It requires two empty squares per possible winning line (group) to be marked, i.e., with a coverage ratio of 2.0. In this paper a new strategy is presented, called *Set Matching*. A matching set consists of a set of nodes (the markers), a set of possible winning lines (the groups), and a coverage set indicating how all groups are covered after every first initial move. This strategy requires less than two markers per group. As such it is able to prove positions in k -in-a-Row games to be draws, which cannot be proven using the Hales–Jewett pairing strategy.

“Playing Hanabi Near-Optimally,” a contribution by Bruno Bouzy, describes a study on the multi-player cooperative card game Hanabi. In this game a player sees the cards of the other players but not his own cards. Previous work using the hat principle reached near-optimal results for five players and four cards per player: On average, the perfect score then was reached in 75% of the cases. In this paper the author has developed HANNIBAL, a set of players, aiming at obtaining near-optimal results as well. The best players use the hat principle and a depth-one search algorithm. For five players and four cards per player, the perfect score is reached in 92% of the cases on average. In addition, by relaxing a debatable rule of Hanabi, the paper generalizes the near-optimal results to other numbers of players and cards per player: The perfect score is reached in 90% of the cases on average. Furthermore, for two players, the hat

principle is useless, and a confidence player is used obtaining high-quality results as well. Overall, this paper shows that the game of Hanabi can be played near-optimally by the computer player.

“Optimal Play of the Farkle Dice Game,” written by Matthew Busche and Todd Neller, presents and solves optimality equations for the two-player, jeopardy dice game Farkle. For fair play, the authors recommend 200 compensation points at the beginning of the game for the second player. The authors then compute the strategy that maximizes the expected score and demonstrate a means for replicating such play with mental mathematics. This method is augmented so as to enable human Farkle play against which complex optimal play maintains only a small win advantage of $\sim 1.7754\%$.

“Deep df-pn and Its Efficient Implementations,” a joint contribution by Zhang Song, Hiroyuki Iida, and Jaap van den Herik, investigates depth-first proof-number search (df-pn). It is a powerful variant of proof-number search, widely used for AND/OR tree search or solving games. However, df-pn suffers from the seesaw effect, which strongly hampers the efficiency in some situations. This paper proposes a new proof-number algorithm called Deep depth-first proof-number search (Deep df-pn) to reduce the seesaw effect in df-pn. The difference between Deep df-pn and df-pn lies in the proof number or disproof number of unsolved nodes. This number is 1 in df-pn, while it is a function of depth with two parameters in Deep df-pn. By adjusting the value of the parameters, Deep df-pn changes its behavior from searching broadly to searching deeply. The paper shows that the adjustment is able to reduce the seesaw effect convincingly. For evaluating the performance of Deep df-pn in the domain of Connect6, the authors have implemented a relevance-zone-oriented Deep df-pn that worked quite efficiently. Experimental results indicate that improvement by the same adjustment technique is also possible in other domains.

Go and Chess

“Improved Policy Networks for Computer Go,” by Tristan Cazenave, utilizes residual policy networks in the Go engine GOLOIS. Two improvements to these residual policy networks are proposed and tested. The first one is to use three output planes. The second one is to add Spatial Batch Normalization.

“Exploring Positional Linear Go,” authored by Noah Weninger and Ryan Hayward, targets Linear Go, the Go variant played on the $1 \times n$ board. The paper investigates Positional Linear Go, which has a rule set that uses positional superko. The paper explores game-theoretic properties of Positional Linear Go, and incorporate them into a solver based on MTD(f) search, solving states on boards up to 1×9 .

“Influence of Search Depth on Position Evaluation,” written by Matej Guid and Ivan Bratko, demonstrates empirically for computer chess that with increasing search depth backed-up evaluations of won positions tend to increase, while backed-up evaluations of lost positions tend to decrease. The authors show three implications of this phenomenon in practice and in the theory of computer game playing. First, they show that heuristic evaluations obtained by searching to different search depths are not directly

comparable, and second that fewer decision changes with deeper search are a direct consequence of this property of heuristic evaluation functions. Third, they demonstrate that knowing this property may be used to develop a method for detecting fortresses in chess, which is currently an unsolved task in computer chess.

“Evaluating Chess-Like Games Using Generated Natural Language Descriptions,” a contribution by Jakub Kowalski, Łukasz Żarczyński, and Andrzej Kisielewicz, continues their study of the chess-like games defined as the class of Simplified Board-games. The paper presents an algorithm generating natural language descriptions of piece movements that can be used as a tool not only for explaining them to the human player, but also for the task of procedural game generation using an evolutionary approach. The authors test their algorithm on some existing human-made and procedurally generated chess-like games.

Machine Learning and MCTS

“Machine Learning in the Game of Breakthrough,” written by Richard Lorentz and Teofilo Erin Zosa, is motivated by recent activity in using machine-learning techniques to game programming. The authors present a study of applying these techniques to the game of Breakthrough. Specifically, the paper shows that by using temporal difference learning in a Monte Carlo Tree Search (MCTS) setting results are achieved almost equal to those obtained by WANDERER, a strong program with a highly tuned evaluation function. The paper also shows that convolutional neural networks trained by using WANDERER as a provider of expert moves can produce a program much stronger than the original. Even in an environment with quite slow execution speeds, excellent results are achieved.

“A Curling Agent Based on the Monte-Carlo Tree Search Considering the Similarity of the Best Action Among Similar States” is authored by Katsuki Ohto and Tetsuro Tanaka. Curling is one of the most strategic winter sports. Recently, many computer scientists have studied curling strategies. The Digital Curling system is a framework used to compare curling strategies. Herein, the authors present a computer agent based on MCTS for the Digital Curling framework. The paper proposes a novel action decision method based on MCTS for Markov decision processes with continuous state space.

“Exploration Bonuses Based on Upper Confidence Bounds for Sparse Reward Games,” written by Naoki Mizukami, Jun Suzuki, Hirotaka Kameko, and Yoshimasa Tsuruoka, has a closer look at deep reinforcement learning algorithms that have achieved super-human-level performance in many Atari games. However, the performance of the algorithms falls short of humans in games where rewards are only obtained occasionally. One solution to this sparse reward problem is to incorporate an explicit and more sophisticated exploration strategy in the agent’s learning process. In this paper, the authors present an effective exploration strategy that explicitly considers the progress of training using exploration bonuses based on Upper Confidence Bounds (UCB). The method also includes a mechanism to separate exploration bonuses from rewards, thereby avoiding the problem of interfering with the original learning

objective. The method is evaluated on Atari 2600 games with sparse rewards, and achieves significant improvements over the vanilla asynchronous advantage actor-critic (A3C) algorithm.

“Developing a 2048 Player with Backward Temporal Coherence Learning and Restart,” by Kiminori Matsuzaki, investigates the popular puzzle game 2048. This is a single-player stochastic game played on a 4×4 grid. After the introduction of the game, several researchers have developed computer players for 2048 based on reinforcement learning methods with N -tuple networks. The paper shows that backward learning is quite useful for 2048, since the game has quite a long sequence of moves in a single play. It is also shown that a restart strategy improves the learning by focusing on the later stage of the game. The resulting player achieves better average scores than the existing players with the same set of N -tuple networks.

(Serious) Gaming

“A Little Bit of Frustration Can Go a Long Way,” written by Adam Boulton, Rachid Hourizi, David Jefferies, and Alice Guy, investigates the phenomenon of frustration in video games. Frustration is reported to impede player engagement but it is unlikely that a game that never frustrated at all would be enjoyable. In that context, further work is required to identify, understand, and model the character, timing, and context of frustrations that help rather than hinder a positive gaming experience. The paper investigates the relationship between frustration and engagement over time in a carefully selected video game. It reveals that engagement often falls as frustration rises (and vice versa) but also reports on situations in which a rise in frustration can give rise to an increase in engagement. Finally, the paper considers the implications of these results for both game developers and the wider community of HCI researchers interested in gamification and user engagement.

“Automated Adaptation and Assessment in Serious Games: A Portable Tool for Supporting Learning,” authored by Enkhbold Nyamsuren, Wim van der Wegt, and Wim Westera, introduces the Adaptation and Assessment (TwoA) component, an open-source tool for serious games, capable of adjusting game difficulty to player skill level. Technically, TwoA is compliant with the RAGE (Horizon 2020) game component architecture, which offers seamless portability to a variety of popular game development platforms. Conceptually, TwoA uses a modified version of the Computer Adaptive Practice algorithm. This version offers two improvements over the original algorithm. First, TwoA improves the balancing of a player’s motivation and game challenge. Second, TwoA reduces the selection bias that may arise for items of similar difficulty by adopting a fuzzy selection rule. The improvements are validated using multi-agent simulations.

“An Analysis of Majority Voting in Homogeneous Groups for Checkers: Understanding Group Performance Through Unbalance” is a contribution by Danilo Carvalho, Minh Le Nguyen, and Hiroyuki Iida. The paper argues that experimental evidence and theoretical advances over the years have created an academic consensus regarding majority voting systems that the group performs better than its components

under certain conditions. However, the underlying reason for such conditions, e.g., stochastic independence of agents, is not often explored and may help to improve performance in known setups by changing agent behavior, or find new ways of combining agents to take better advantage of their characteristics. In this paper, an investigation is conducted for homogeneous groups of independent agents playing the game of Checkers. The analysis aims to find the relationship between the change in performance caused by majority voting, the group size, and the underlying decision process of each agent, which is mapped to its source of non-determinism. A characteristic unbalance in Checkers, due to an apparent initiative disadvantage, serves as a pivot for the study, from which decisions can be separated into beneficial or detrimental biases. Experimental results indicate that performance changes caused by majority voting may be beneficial or not, and are linked to the game properties and player skill. Additionally, a way of improving agent performance by manipulating its non-determinism source is briefly explored.

“Yasol: An Open Source Solver for Quantified Mixed Integer Programs” is authored by Thorsten Ederer, Michael Hartisch, Ulf Lorenz, Thomas Opfer, and Jan Wolf. The paper discusses Quantified Mixed Integer Linear Programs (QMIPs), which are mixed integer linear programs (MIPs) with variables being either existentially or universally quantified. They can be interpreted as two-person zero-sum games between an existential and a universal player on the one side, or multistage optimization problems under uncertainty on the other side. Solutions of QMIPs are so-called winning strategies for the existential player that specify how to react on moves – certain fixations of universally quantified variables – of the universal player to certainly win the game. To solve the QMIP optimization problem, where the task is to find an especially attractive winning strategy, the paper examines the problem’s hybrid nature and presents the open source solver Yasol that combines linear programming techniques with solution techniques from game-tree search.

This book would not have been produced without the help of many persons. In particular, we would like to mention the authors and the reviewers for their help. Moreover, the organizers of the three events in Leiden (see the beginning of this preface) have contributed substantially by bringing the researchers together. Without much emphasis, we recognize the work by the committees of the ACG 2017 as essential for this publication. Moreover, we gratefully acknowledge the support by Monique Arntz, who helped us with the organization and the proceedings. Thank you, Monique. Finally, the editors happily acknowledge the generous sponsors Oracle, the Municipality of Leiden, SURFsara, ICT Shared Service Centre, the Leiden Institute of Advanced Computer Science, the Leiden Centre of Data Science, ICGA, and Digital Games Technology.

September 2017

Mark H. M. Winands
Jaap van den Herik
Walter A. Kusters

Advances in Computer Games

15th International Conferences, ACG 2017, Leiden, The Netherlands, July 3-5, 2017, Revised Selected Papers

Winands, M.; van den Herik, J.; Kusters, W. (Eds.)

2017, XX, 235 p. 117 illus., Softcover

ISBN: 978-3-319-71648-0