

Preface

These joint proceedings contain the papers of the Computer Games Workshop (CGW 2016) and the General Intelligence in Game-Playing Agents (GIGA 2016) Workshop, which were both held in New York, USA. These workshops took place on July 9 and 10, 2016, respectively, in conjunction with the 25th International Conference on Artificial Intelligence (IJCAI 2016). These two workshops reflect the large interest in artificial intelligence (AI) research for games.

The Computer and Games Workshop series is an international forum for researchers interested in all aspects of AI and computer game playing. Earlier workshops took place in Montpellier, France (2012), Beijing, China (2013), Prague, Czech Republic (2014), and Buenos Aires, Argentina (2015). For the fifth edition of the Computer Games Workshop, 17 submissions were received in 2016. Each paper was sent to three reviewers. In the end, 11 papers were accepted for presentation at the workshop, of which nine made it into these proceedings. The published papers cover a wide range of topics related to computer games. They collectively discuss five abstract games: Breakthrough, Go, Hex, SameGame, and Werewolf. Additionally, one paper deals with optimization problems such as bus regulations and weak Schur numbers, and two papers are on video games.

The GIGA workshop series has been established to become the major forum for discussing, presenting, and promoting research on general game playing (GGP). It aims at building intelligent software agents that can, given the rules of any game, automatically learn a strategy for playing that game at an expert level without any human intervention. The workshop intends to bring together researchers from subfields of AI to discuss how best to address the challenges and further advance the state of the art of general game-playing systems and generic AI. Following the inaugural GIGA Workshop at IJCAI 2009 in Pasadena (USA), follow-up events took place at IJCAI 2011 in Barcelona (Spain), IJCAI 2013 in Beijing (China), and IJCAI 2015 in Buenos Aires (Argentina). This fifth workshop on General Intelligence in Game-Playing Agents received eight submissions. Each paper was sent to three reviewers. All papers were accepted for presentation at the workshop, but in the end three were accepted for these proceedings. The accepted papers focus on general techniques for automated reasoning about new games and cover the topics of propositional networks, ground instantiations of game rules, and decomposition of game descriptions.

In all, 48% of the submitted papers for both workshops were selected for these proceedings. Here we provide a brief outline of the 12 contributions, in the order in which they appear in the proceedings. They are divided into two parts: the first nine belong to the Computer Games Workshop and the last three to the GIGA Workshop.

Computer Games Workshop

“NeuroHex: A Deep Q-learning Hex Agent,” a joint effort by Kenny Young, Gautham Vasan, and Ryan Hayward, considers deep Q-learning for the game of Hex. After supervised initializing, self-play is used to train NeuroHex, an 11-layer convolutional neural network that plays Hex on the 13×13 board. Hex is the classic two-player alternate-turn stone placement game played on a rhombus of hexagonal cells in which the winner is the one who connects their two opposing sides. Despite the large action and state space, their system trains a Q-network capable of strong play with no search. After two weeks of Q-learning, NeuroHex achieves win rates of 20.4% as first player and 2.1% as second player against a one-second/move version of MoHex, the current ICGA Olympiad Hex champion. The data suggest further improvement might be possible with more training time.

“Deep or Wide? Learning Policy and Value Neural Networks for Combinatorial Games,” by Stefan Edelkamp, raises the question on the availability, the limits, and the possibilities of deep neural networks for other combinatorial games than Go. As a step toward this direction, a value network for Tic-Tac-Toe was trained, providing perfect winning information obtained by retrograde analysis. Next, a policy network was trained for the SameGame, a challenging combinatorial puzzle. Here, the interplay of deep learning with nested rollout policy adaptation (NRPA) is discussed, a randomized algorithm for optimizing the outcome of single-player games. In both cases the observation is that ordinary feed-forward neural networks can perform better than convolutional ones both in accuracy and efficiency.

“Integrating Factorization Ranked Features in MCTS: An Experimental Study” authored by Chenjun Xiao and Martin Müller investigates the problem of integrating feature knowledge learned by the factorization Bradley–Terry model in Monte Carlo tree search (MCTS). The open source Go program Fuego is used as the test platform. Experimental results show that the FBT knowledge is useful in improving the performance of Fuego.

“Nested Rollout Policy Adaptation with Selective Policies,” by Tristan Cazenave, discusses nested rollout policy adaptation (NRPA). It is a Monte Carlo search algorithm that has found record-breaking solutions for puzzles and optimization problems. It learns a playout policy online that dynamically adapts the playouts to the problem at hand using more selectivity in the playouts. The idea is applied to three different domains: Bus regulation, SameGame, and weak Schur numbers. For each of these problems, selective policies improve NRPA.

“A Rollout-Based Search Algorithm Unifying MCTS and Alpha-Beta” by Hendrik Baier integrates MCTS and minimax tightly into one rollout-based hybrid search algorithm, MCTS- $\alpha\beta$. The hybrid is able to execute two types of rollouts: MCTS rollouts and alpha-beta rollouts. During the search, all nodes accumulate both MCTS value estimates as well as alpha-beta value bounds. The two types of information are combined in a given tree node whenever alpha-beta completes a deepening iteration rooted in that node by increasing the MCTS value estimates for the best move found by alpha-beta. A single parameter, the probability of executing MCTS rollouts vs. alpha-beta rollouts, makes it possible for the hybrid to subsume both MCTS as well as

alpha-beta search as extreme cases, while allowing for a spectrum of new search algorithms in between. Preliminary results in the game of Breakthrough show that the proposed hybrid outperforms its special cases of alpha-beta and MCTS.

“Learning from the Memory of Atari 2600,” written by Jakub Sygnowski and Henryk Michalewski, describes the training of neural networks to play the games Bowling, Breakout, and Seaquest using information stored in the memory of a video game console Atari 2600. Four models of neural networks are considered that differ in size and architecture: two networks that use only information contained in the RAM and two mixed networks that use both information in the RAM and information from the screen. In all games the RAM agents are on a par with the benchmark screen-only agent. In the case of Seaquest, the trained RAM-only agents behave even better than the benchmark agent. Mixing screen and RAM does not lead to an improved performance compared with screen-only and RAM-only agents.

“Clustering-Based Online Player Modeling,” a joint collaboration by Jason M. Bindewald, Gilbert L. Peterson, and Michael E. Miller, presents a clustering and locally weighted regression method for modeling and imitating individual players. The algorithm first learns a generic player cluster model that is updated online to capture an individual’s game-play tendencies. The models can then be used to play the game or for analysis to identify how different players react to separate aspects of game states. The method is demonstrated on a tablet-based trajectory generation game called Space Navigator.

“AI Wolf Contest — Development of Game AI using Collective Intelligence,” a joint effort by Fujio Toriumi, Hirotaka Osawa, Michimasa Inaba, Daisuke Katagami, Kosuke Shinoda, and Hitoshi Matsubara, introduces a platform to develop a game-playing AI for a Werewolf competition. First, the paper discusses the essential factors in Werewolf with reference to other studies. Next, a platform for an AI game competition is developed that uses simplified rules to support the development of AIs that can play Werewolf. The paper reports the process and analysis of the results of the competition.

“Semantic Classification of Utterances in a Language-Driven Game,” written by Kellen Gillespie, Michael W. Floyd, Matthew Molineaux, Swaroop S. Vattam, and David W. Aha, describes a goal reasoning agent architecture that allows an agent to classify natural language utterances, hypothesize about humans’ actions, and recognize their plans and goals. The paper focuses on one module of the architecture, the natural language classifier, and demonstrates its use in a multiplayer tabletop social deception game, One Night Ultimate Werewolf. The results indicate that the system can obtain reasonable performance even when the utterances are unstructured, deceptive, or ambiguous.

GIGA Workshop

“Optimizing Propositional Networks” authored by Chiara F. Sironi and Mark H.M. Winands analyzes the performance of a Propositional Network (PropNet)-based reasoner for interpreting the game rules, written in the Game Description Language (GDL). The paper evaluates four different optimizations for the PropNet structure that can help further increase its reasoning speed in terms of visited game states per second.

“Grounding GDL Game Descriptions” by Stephan Schiffel discusses grounding game descriptions using a state-of-the art answer set programming system as a viable alternative to the GDL specific approach implemented in the GGP-Base framework. The presented system is able to handle more games and is typically faster despite the overhead of transforming GDL into a different format and starting and communicating with a separate process. Furthermore, this grounding of a game description is well-founded theoretically by the transformation into answer set programs. It allows one to optimize the descriptions further without changing their semantics.

“A General Approach of Game Description Decomposition for General Game Playing,” a joint effort by Aline Hufschmitt, Jean-Noël Vittaut, and Jean Méhat, presents a general approach for the decomposition of games described in GDL. In the field of general game playing, the exploration of games described in GDL can be significantly sped up by the decomposition of the problem in subproblems analyzed separately. The discussed program can decompose game descriptions with any number of players while addressing the problem of joint moves. This approach is used to identify perfectly separable subgames but can also decompose serial games composed of two subgames and games with compound moves while avoiding reliance on syntactic elements that can be eliminated by simply rewriting the GDL rules. The program has been tested on 40 games, compound or not. It decomposes 32 of them successfully in less than five seconds.

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 2016 contributed substantially by bringing the researchers together.

January 2017

Tristan Cazenave
Mark H.M. Winands
Stefan Edelkamp
Stephan Schiffel
Michael Thielscher
Julian Togelius

Computer Games

5th Workshop on Computer Games, CGW 2016, and 5th
Workshop on General Intelligence in Game-Playing
Agents, GIGA 2016, Held in Conjunction with the 25th
International Conference on Artificial Intelligence, IJCAI
2016, New York, USA, July 9-10, 2016, Revised Selected
Papers

Cazenave, T.; Winands, M.; Edelkamp, S.; Schiffel, S.;

Thielscher, M.; Togelius, J. (Eds.)

2017, XII, 179 p. 59 illus., Softcover

ISBN: 978-3-319-57968-9