

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Alfred Kobsa

University of California, Irvine, CA, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

TU Dortmund University, Germany

Madhu Sudan

Microsoft Research, Cambridge, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Gerhard Weikum

Max Planck Institute for Informatics, Saarbruecken, Germany

H. Jaap van den Herik Aske Plaat (Eds.)

Advances in Computer Games

13th International Conference, ACG 2011
Tilburg, The Netherlands, November 20-22, 2011
Revised Selected Papers

Volume Editors

H. Jaap van den Herik
Aske Plaat
Tilburg University
Tilburg Institute of Cognition and Communication
Warandelaan 2, 5037 AB Tilburg, The Netherlands
E-mail: {jaapvandenherik, aske.plaat}@gmail.com

ISSN 0302-9743 e-ISSN 1611-3349
ISBN 978-3-642-31865-8 e-ISBN 978-3-642-31866-5
DOI 10.1007/978-3-642-31866-5
Springer Heidelberg Dordrecht London New York

Library of Congress Control Number: 2012941836

CR Subject Classification (1998): F.2, F.1, I.2, G.2, I.4, C.2

LNCS Sublibrary: SL 1 – Theoretical Computer Science and General Issues

© Springer-Verlag Berlin Heidelberg 2012

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

This book contains the papers of the 13th Advances in Computer Games Conference (ACG 2011) held in Tilburg, The Netherlands. The conference took place during November 20–22, 2011, in conjunction with the 16th Computer Olympiad and the 19th 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. The Tilburg conference was definitively characterized by the progress of Monte Carlo Tree Search (MCTS) and the development of new games. 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), and Pamplona (2009).

The Program Committee (PC) was pleased to see that so much progress was made in MCTS and that on top of that new games and new techniques were added to the recorded achievements. Each paper was sent to at least three referees. If conflicting views on a paper were reported, the referees themselves arrived at an appropriate decision. With the help of many referees (see after the preface), the PC accepted 29 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 (1) to give authors the opportunity to include the results of the fruitful discussion after the lecture into their paper, and (2) to maintain the high-quality threshold of the ACG series. The authors enjoyed this procedure.

The above-mentioned set of 29 papers covers a wide range of computer games and many different research topics. We mention the topics in the order of publication: Monte Carlo Tree Search and its enhancements (10 papers), temporal difference learning (2 papers), optimization (4 papers), solving and searching (2 papers), analysis of a game characteristic (3 papers), new approaches (5 papers), and serious games (3 papers).

We hope that the readers will enjoy the research efforts made by the authors. Below we reproduce brief characterizations of the 29 contributions taken from the text as submitted by the authors. The authors of the first publication “Accelerated UCT and Its Application to Two-Player Games” received the Best Paper Award of ACG 2011.

“Accelerated UCT and Its Application to Two-Player Games” by Junichi Hashimoto, Akihiro Kishimoto, Kazuki Yoshizoe, and Kokoro Ikeda describes Monte-Carlo Tree Search (MCTS) as a successful approach for improving the performance of game-playing programs. A well-known weakness of MCTS is caused by the deceptive structures which often appear in game tree search. To overcome the weakness the authors present the Accelerated UCT algorithm

(Upper Confidence Bounds applied to Trees). Their finding consists in using a new back-up operator that assigns higher weights to recently visited actions, and lower weights to actions that have not been visited for a long time. Results in Othello, Havannah, and Go show that Accelerated UCT is not only more effective than previous approaches but also improves the strength of FUEGO, which is one of the best computer Go programs.

“Revisiting Move Groups in Monte Carlo Tree Search” is authored by Gabriel Van Eyck and Martin Müller. These authors also remark that the UCT (Upper Confidence Bounds applied to Trees) algorithm has been a stimulus for significant improvements in a number of games, most notably the game of Go. They investigate the use of move groups. Move groups is a modification that greatly reduces the branching factor at the cost of an increased search depth and as such it may be used to enhance the performance of UCT. From the results of the experiments, the authors arrive at a general structure of good move groups in which they determine which parameters to use for enhancing the playing strength.

In “PACHI: State-of-the-Art Open Source Go Program”, Petr Baudiš and Jean-loup Gailly start describing a state-of-the-art implementation of the Monte Carlo Tree Search algorithm for the game of Go. Their PACHI software is currently one of the strongest open source Go programs, competing at the top level with other programs and playing evenly against advanced human players. The paper describes their framework (implementation and chosen algorithms) together with three notable original improvements: (1) an adaptive time control algorithm, (2) dynamic komi, and (3) usage of the criticality statistic. Moreover, new methods to achieve efficient scaling both in terms of multiple threads and multiple machines in a cluster are presented.

“Time Management for Monte Carlo Tree Search in Go” is written by Hendrik Baier and Mark H.M. Winands. So far, little has been published on time management for MCTS programs under tournament conditions. The authors investigate the effects that various time-management strategies have on the playing strength in Go. They consider strategies taken from the literature as well as newly proposed and improved ones. Moreover, they investigate both *semi-dynamic* strategies that decide about time allocation for each search before it is started, and *dynamic* strategies that influence the duration of each move search while it is already running. In their experiments, two domain-independent enhanced strategies, EARLY-C and CLOSE-N, are tested; each of them provides a significant improvement over the state of the art.

“An MCTS Program to Play EinStein Würfelt Nicht!” by Richard Lorentz describes a game that has elements of strategy, tactics, and chance. The author remarks that reasonable evaluation functions for this game can be found. Nevertheless, he constructed an MCTS program to play this game. The paper describes the basic structure and its strengths and weaknesses. Then the MCTS program is successfully compared with existing mini-max-based programs and to a pure MC version.

“Monte Carlo Tree Search Enhancements for Havannah” is authored by Jan A. Stankiewicz, Mark H.M. Winands, and Jos W.H.M. Uiterwijk. The article shows how the performance of a Monte Carlo Tree Search (MCTS) player for Havannah can be improved by guiding the search in the *playout* and *selection* steps of MCTS. To enhance the *playout* step of the MCTS algorithm, the authors used two techniques to direct the simulations, Last-Good-Reply (LGR) and N-grams. Experiments reveal that LGR gives a significant improvement, although it depends on which LGR variant is used. Using N-grams to guide the playouts also achieves a significant increase in the winning percentage. Combining N-grams with LGR leads to a small additional improvement. To enhance the *selection* step of the MCTS algorithm, the authors initialize the visit and win counts of the new nodes based on pattern knowledge. Experiments show that the best overall performance is obtained when combining the visit-and-win-count initialization with LGR and N-grams. In the best case, a winning percentage of 77.5% can be achieved against the default MCTS program.

“Payout Search for Monte-Carlo Tree Search in Multi-Player Games” by J. (Pim) A.M. Nijssen and Mark H.M. Winands proposes a technique called Payout Search. This enhancement allows the use of small searches in the *playout* phase of MCTS in order to improve the reliability of the playouts. The authors investigate \max^n , Paranoid, and BRS for Payout Search and analyze their performance in two deterministic perfect-information multi-player games: Focus and Chinese Checkers. The experimental results show that Payout Search significantly increases the quality of the playouts in both games.

“Towards a Solution of 7x7 Go with Meta-MCTS” by Cheng-Wei Chou, Ping-Chiang Chou, Hassen Doghmen, Chang-Shing Lee, Tsan-Cheng Su, Fabien Teytaud, Olivier Teytaud, Hui-Ming Wang, Mei-Hui Wang, Li-Wen Wu, and Shi-Jim Yen is a challenging topic. So far, Go is not solved (in any sense of solving, even the weakest) beyond 6x6. The authors investigate the use of Meta-Monte-Carlo Tree Search, for building a huge 7x7 opening book. In particular, they report the 20 wins (out of 20 games) that were obtained recently in 7x7 Go against pros; they also show that in one of the games, with no human error, the pro might have won.

“MCTS Experiments on the Voronoi Game” written by Bruno Bouzy, Marc Métivier, and Damien Pellier discusses Monte Carlo Tree Search (MCTS) as a powerful tool in games with a finite branching factor. The use of MCTS on a discretization of the Voronoi game is described together with the effects of enhancements such as RAVE and Gaussian processes (GP). A set of experimental results shows that MCTS with UCB+RAVE or with UCB+GP are good first solutions for playing the Voronoi game without domain-dependent knowledge. Then the authors show how to improve the playing level by using geometrical knowledge about Voronoi diagrams, the balance of diagrams being the key concept. A new set of experimental results shows that a player using MCTS and geometrical knowledge outperforms the player without knowledge.

“4*4-Pattern and Bayesian Learning in Monte-Carlo Go” is a contribution by Jiao Wang, Shiyuan Li, Jitong Chen, Xin Wei, Huizhan Lv, and Xinhe Xu.

The authors propose a new model of pattern, namely, 4*4-Pattern, to improve MCTS in computer Go. A 4*4-Pattern provides a larger coverage and more essential information than the original 3*3-Patterns, which are currently widely used. Due to the lack of a central symmetry, it takes greater challenges to apply a 4*4-Pattern compared to a 3*3-Pattern. Many details of a 4*4-Pattern implementation are presented, including classification, multiple matching, coding sequences, and fast lookup. Additionally, Bayesian 4*4-Pattern learning is introduced, and 4*4-Pattern libraries are automatically generated from a vast amount of professional game records according to the method. The results of the experiments show that the use of 4*4-Patterns can improve MCTS in 19*19 Go to some extent, in particular when supported by 4*4-Pattern libraries generated by Bayesian learning.

“Temporal Difference Learning for Connect6” is written by I-Chen Wu, Hsin-Ti Tsai, Hung-Hsuan Lin, Yi-Shan Lin, Chieh-Min Chang, and Ping-Hung Lin. In the paper, the authors apply temporal difference (TD) learning to Connect6, and successfully use TD(0) to improve the strength of their Connect6 program, NCTU6. That program won several computer Connect6 tournaments from 2006 to 2011. The best improved version of TD learning achieves about a 58% win rate against the original NCTU6 program. The paper discusses several implementation issues that improve the program. The program has a convincing performance removing winning/losing moves via threat-space search in TD learning.

“Improving Temporal Difference Learning Performance in Backgammon Variants” by Nikolaos Papahristou and Ioannis Refanidis describes the project. Palamedes which is an ongoing project for building expert playing bots that can play backgammon variants. The paper improves upon the training method used in their previous approach for the two backgammon variants popular in Greece and neighboring countries, Plakoto and Fevga. The authors show that the proposed methods result both in faster learning as well as better performance. They also present insights into the selection of the features.

“CLOP: Confident Local Optimization for Noisy Black-Box Parameter Tuning” is a contribution by Rémi Coulom. Artificial intelligence in games often leads to the problem of parameter tuning. Some heuristics may have coefficients, and they should be tuned to maximize the win rate of the program. A possible approach is to build local quadratic models of the win rate as a function of the program parameters. Many local regression algorithms have already been proposed for this task, but they are usually not sufficiently robust to deal automatically and efficiently with very noisy outputs and non-negative Hessians. The CLOP principle is a new approach to local regression that overcomes all these problems in a straightforward and efficient way. CLOP discards samples of which the estimated value is confidently inferior to the mean of all samples. Experiments demonstrate that, when the function to be optimized is smooth, this method outperforms all other tested algorithms.

“Analysis of Evaluation-Function Learning by Comparison of Sibling Nodes” written by Tomoyuki Kaneko and Kunihiro Hoki, discusses gradients of search values with a parameter vector θ in an evaluation function. Recent learning

methods for evaluation functions in computer shogi are based on minimization of an objective function with search results. The gradients of the evaluation function at the leaf position of a principal variation (PV) are used to make an easy substitution of the gradients of the search result. By analyzing the variations of the min-max value, the authors show (1) when the min-max value is partially differentiable and (2) how the substitution may introduce errors. Experiments on a shogi program with about 1 million parameters show how frequently such errors occur, as well as how effective the substitutions for parameter tuning are in practice.

“Approximating Optimal Dudo Play with Fixed-Strategy Iteration Counterfactual Regret Minimization” is a contribution by Todd W. Neller and Steven Hnath. Using the bluffing dice game Dudo as a challenge domain, the authors abstract information sets by an imperfect recall of actions. Even with such abstraction, the standard Counterfactual Regret Minimization (CFR) algorithm proves impractical for Dudo, since the number of recursive visits to the same abstracted information sets increases exponentially with the depth of the game graph. By holding strategies fixed across each training iteration, the authors show how CFR training iterations may be transformed from an exponential-time recursive algorithm into a polynomial-time dynamic-programming algorithm, making computation of an approximate Nash equilibrium for the full two-player game of Dudo possible for the first time.

“The Global Landscape of Objective Functions for the Optimization of Shogi Piece Values with a Game-Tree Search” is written by Kunihiro Hoki and Tomoyuki Kaneko. The landscape of an objective function for supervised learning of evaluation functions is numerically investigated for a limited number of feature variables. Despite the importance of such learning methods, the properties of the objective function are still not well known because of its complicated dependence on millions of tree-search values. The paper shows that the objective function has multiple local minima and the global minimum point indicates reasonable feature values. It is shown that an existing iterative method is able to minimize the functions from random initial values with great stability, but it has the possibility to end up with a non-reasonable local minimum point if the initial random values are far from the desired values.

“Solving BREAKTHROUGH with Race Patterns and Job-Level Proof Number Search” is a contribution by Abdallah Saffidine, Nicolas Jouandeau, and Tristan Cazenave. BREAKTHROUGH is a recent race-based board game usually played on an 8×8 board. The authors describe a method to solve 6×5 boards based on race patterns and an extension of the Job-Level Proof Number Search JLPNS. Using race patterns is a new domain-specific technique that allows early endgame detection. The patterns they use enable them to prune positions safely and statically as far as seven moves from the end. The authors also present an extension of the parallel algorithm (JLPNS), viz., when a PN search is used as the underlying job.

“Infinite Connect-Four Is Solved: Draw” by Yoshiaki Yamaguchi, Kazunori Yamaguchi, Tetsuro Tanaka, and Tomoyuki Kaneko, describes the newly obtained solution for variants of Connect-Four played on an infinite board. The authors proved their result by introducing never-losing strategies for both players. The strategies consist of a combination of paving patterns, which are follow-up, follow-in-CUP, and a few others. By employing the strategies, both players can block their opponents to achieve the winning condition. This means that optimal play by both players leads to a draw in these games.

“Blunder Cost in Go and Hex” is a contribution by Henry Brausen, Ryan B. Hayward, Martin Müller, Abdul Qadir, and David Spies. In Go and Hex, they examine the effect of a blunder — here, a random move — at various stages of a game. For each fixed move number, they run a self-play tournament to determine the expected blunder cost at that point.

“Understanding Distributions of Chess Performances” by Kenneth W. Regan, Bartłomiej Macieja, and Guy M^cC. Haworth studies the population of chess players and the distribution of their performances measured by Elo ratings and by computer analysis of moves. Evidence that ratings have remained stable since the inception of the Elo system in the 1970s is given in three forms: (1) by showing that the population of strong players fits a straightforward logistic-curve model without inflation, (2) by plotting players’ average error against the FIDE category of tournaments over time, and (3) by skill parameters from a model that employs computer analysis keeping a nearly constant relation to Elo rating across that time. The distribution of the model’s *intrinsic performance ratings* can therefore be used to compare populations that have limited interaction, such as between players in a national chess federation and FIDE, and to ascertain relative drift in their respective rating systems.

“Position Criticality in Chess Endgames” is a contribution by Guy M^cC. Haworth and Á. Rusz. In some 50,000 Win Studies in Chess, White is challenged to find an effectively unique route to a win. Judging the impact of less than absolute uniqueness requires both technical analysis and artistic judgment. Here, for the first time, an algorithm is defined to help analyze uniqueness in endgame positions objectively. The key idea is to examine how critical certain positions are to White in achieving the win. The algorithm uses sub-n-man endgame tables (EGTs) for both Chess and relevant, adjacent variants of Chess. It challenges authors of EGT generators to generalize them to create EGTs for these chess variants. The algorithm has already proved to be efficient and effective in an implementation for Starchess, itself a variant of chess. The approach also addresses a number of similar questions arising in endgame theory, games, and compositions.

“On Board-Filling Games with Random-Turn Order and Monte Carlo Perfection” is a contribution by Ingo Althöfer. In a game, pure Monte Carlo search with parameter T means that for each feasible move T random games are generated. The move with the best average score is played. The author calls a game “Monte Carlo perfect” when this straightforward procedure converges to perfect play for each position, when T goes to infinity. Many popular games like Go, Hex,

and Amazons are NOT Monte Carlo perfect. In the paper, two-player zero-sum games are investigated where the turn-order is random: always a fair coin flip decides which player acts on the next move. A whole class of such random-turn games is proven to be Monte Carlo perfect. The result and generalizations are discussed, with example games ranging from very abstract to very concrete.

“Modeling Games with the Help of Quantified Integer Linear Programs” is written by Thorsten Ederer, Ulf Lorenz, Thomas Opfer, and Jan Wolf. Quantified linear programs (QLPs) are linear programs with mathematical variables being either existentially or universally quantified. The integer variant (Quantified linear integer program, QIP) is PSPACE-complete, and can be interpreted as a two-person zero-sum game. Additionally, it demonstrates a remarkable flexibility in polynomial reduction, such that many interesting practical problems can be elegantly modeled as QIPs. Indeed, the PSPACE-completeness guarantees that all PSPACE-complete problems, for example, games like Othello, Go-Moku, and Amazons, can be described with the help of QIPs, with only moderate overhead. The authors present the *dynamic graph reliability* (DGR) optimization problem and the game *Go-Moku* as examples.

“Computing Strong Game-Theoretic Strategies in Jotto” by Sam Ganzfried describes a new approach that computes approximate equilibrium strategies in Jotto. Jotto is quite a large two-player game of imperfect information; its game tree has many orders of magnitude more states than games previously studied, including no-limit Texas Hold’em. To address the fact that the game tree is so large, the authors propose a novel strategy representation called oracular form, in which they do not explicitly represent a strategy, but rather appeal to an oracle that quickly outputs a sample move from the strategy’s distribution. Their overall approach is based on an extension of the fictitious play algorithm to this oracular setting. The authors demonstrate the superiority of their computed strategies over the strategies computed by a benchmark algorithm, both in terms of head-to-head and worst-case performance.

“Online Sparse Bandit for Card Games” is written by David L. St-Pierre, Quentin Louveaux, and Olivier Teytaud. Finding an approximation of a Nash equilibrium in matrix games is an important topic. A bandit algorithm commonly used to approximate a Nash equilibrium is EXP3. Although the solution to many problems is often sparse, EXP3 inherently fails to exploit this property. To the authors’ best knowledge, there is only an offline truncation proposed to handle the sparseness issue. Therefore, the authors propose a variation of EXP3 to exploit the fact that the solution is sparse by dynamically removing arms; the resulting algorithm empirically performs better than previous versions. The authors apply the resulting algorithm to an MCTS program for the Urban Rivals card game.

“Game Tree Search with Adaptive Resolution” is authored by Hung-Jui Chang, Meng-Tsung Tsai, and Tsan-sheng Hsu. In the paper, the authors use an adaptive resolution R to enhance the min-max search with the alpha-beta pruning technique, and show that the value returned by the modified algorithm, called Negascout-with-resolution, differs from that of the original version by at

most *R*. Guidelines are given to explain how the resolution should be chosen to obtain the best possible outcome. The experimental results demonstrate that Negascout-with-resolution yields a significant performance improvement over the original algorithm on the domains of random trees and real game trees in Chinese chess.

“Designing Casanova: A Language for Games” is written by G. Maggiore, A. Spanó, R. Orsini, G. Costantini, M. Bugliesi, and M. Abbadi. The authors present the Casanova language, which allows the building of games with three important advantages when compared to traditional approaches: simplicity, safety, and performance. They show how to rewrite an official sample of the XNA framework, resulting in a smaller source and a higher performance.

“Affective Game Dialogues” by Michael Lankes and Thomas Mirlacher investigates natural game input devices, such as Microsoft’s Kinect or Sony’s Playstation Move. They have become increasingly popular and allow a direct mapping of player performance in regard to actions in the game world. Games have been developed that enable players to interact with their avatars and other game objects via gestures and/or voice input. However, current technologies and systems do not tap into the full potential of affective approaches. Affect in games can be harnessed as a supportive and easy to use input method. The paper proposes a design approach that utilizes facial expressions as an explicit input method in game dialogues. This concept allows players to interact with non-player characters (NPC) by portraying specific basic emotions.

“Generating Believable Virtual Characters Using Behavior Capture and Hidden Markov Models” by Richard Zhao and Duane Szafron proposes a method of generating natural-looking behaviors for virtual characters using a data-driven method called behavior capture. The authors describe the techniques (1) for capturing trainer-generated traces, (2) for generalizing these traces, and (3) for using the traces to generate behaviors during game-play. Hidden Markov models (HMMs) are used as one of the generalization techniques for behavior generation. The authors compared the proposed method with other existing methods by creating a scene with a set of six variations in a computer game, each using a different method for behavior generation, including their proposed method. They conducted a study in which participants watched the variations and ranked them according to a set of criteria for evaluating behaviors. The study showed that behavior capture is a viable alternative to existing manual scripting methods and that HMMs produced the most highly ranked variation with respect to overall believability.

This book would not have been produced without the help of many persons. In particular, we would like to mention the authors and the referees for their help. Moreover, the organizers of the three events in Tilburg (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 2011 as essential for this publication. One exception is made for Joke Hellemons, who is gratefully thanked for all services to our games community.

Finally, the editors happily recognize the generous sponsors Tilburg University, Tilburg Center for Cognition and Communication, LIS, ICGA, Netherlands eScience Center, NWO NCF, SURFnet, CICK, The Red Brick, NBrIX, and Digital Games Technology.

April 2012

Jaap van den Herik
Aske Plaat

Organization

Executive Committee

Editors	H. Jaap van den Herik Aske Plaat
---------	-------------------------------------

Program Co-chairs	H. Jaap van den Herik Aske Plaat
-------------------	-------------------------------------

Organizing Committee

Johanna W. Hellemons (Chair)	Aske Plaat
H. Jaap van den Herik	Susanne van der Velden
Erik Kamps	Eva Verschoor

Sponsors

Tilburg University
Tilburg Center for Cognition and Communication (TiCC)
LIS
ICGA
Netherlands eScience Center
NWO NCF
SURFnet
CICK
The Red Brick
NBrIX
Digital Games Technology

Program Committee

Ingo Althöfer	Aviezri Fraenkel	Akihiro Kishimoto
Yngvi Björnsson	James Glenn	Yoshiyuki Kotani
Bruno Bouzy	Matej Guid	Clyde Kruskal
Ivan Bratko	Dap Hartmann	Richard Lorenz
Tristan Cazenave	Tsuyoshi Hashimoto	Ulf Lorenz
Jr-Chang Chen	Guy M ^c C. Haworth	Hitoshi Matsubara
Paolo Ciancarini	Ryan Hayward	John-Jules Meyer
Rémi Coulom	Tsan-Sheng Hsu	Martin Müller
Omid David Tabibi	Han-Shen Huang	Jacques Pitrat
Jeroen Donkers	Hiroyuki Iida	Christian Posthoff
David Fotland	Graham Kendall	Matthias Rauterberg

Alexander Sadikov	Gerald Tesauro	Hans Weigand
Jahn Saito	Yoshimasa Tsuruoka	Mark Winands
Jonathan Schaeffer	Jos Uiterwijk	I-Chen Wu
Yaron Shoham	Erik van der Werf	Georgios Yannakakis
Pieter Spronck	Peter van Emde Boas	Shi-Jim Yen
Nathan Sturtevant	Jan van Zanten	
Tetsuro Tanaka	Gert Vriend	

Additional Referees

Victor Allis	Robert Hyatt	David Silver
Darse Billings	Shaul Markovitch	Duane Szafron
Cameron Browne	Tod Neller	Olivier Teytaud
Johannes Fürnkranz	Pim Nijssen	John Tromp
Reijer Grimbergen	Maarten Schadd	Thomas Wolf
Shun-Chin Hsu	Richard Segal	

The Advances in Computers and Chess / Games Books

The series of Advances in Computer Chess (ACC) Conferences started in 1975 as a complement to the World Computer-Chess Championships, for the first time held in Stockholm in 1974. In 1999, the title of the conference changed from ACC to ACG (Advances in Computer Games). Since 1975, 13 ACC/ACG conferences have been held. Below we list the conference places and dates together with the publication; the Springer publication is supplied with an LNCS series number.

London, England (1975, March)
 Proceedings of the 1st Advances in Computer Chess Conference (ACC1)
 Ed. M.R.B. Clarke
 Edinburgh University Press, 118 pages.

Edinburgh, United Kingdom (1978, April)
 Proceedings of the 2nd Advances in Computer Chess Conference (ACC2)
 Ed. M.R.B. Clarke
 Edinburgh University Press, 142 pages.

London, England (1981, April)
 Proceedings of the 3rd Advances in Computer Chess Conference (ACC3)
 Ed. M.R.B. Clarke
 Pergamon Press, Oxford, UK, 182 pages.

London, England (1984, April)
 Proceedings of the 4th Advances in Computer Chess Conference (ACC4)
 Ed. D.F. Beal
 Pergamon Press, Oxford, UK, 197 pages.

Noordwijkerhout, The Netherlands (1987, April)
Proceedings of the 5th Advances in Computer Chess Conference (ACC5)
Ed. D.F. Beal
North Holland Publishing Comp., Amsterdam, The Netherlands, 321 pages.

London, England (1990, August)
Proceedings of the 6th Advances in Computer Chess Conference (ACC6)
Ed. D.F. Beal
Ellis Horwood, London, UK, 191 pages.

Maastricht, The Netherlands (1993, July)
Proceedings of the 7th Advances in Computer Chess Conference (ACC7)
Eds. H.J. van den Herik, I.S. Herschberg, and J.W.H.M. Uiterwijk
Drukkerij Van Spijk B.V. Venlo, The Netherlands, 316 pages.

Maastricht, The Netherlands (1996, June)
Proceedings of the 8th Advances in Computer Chess Conference (ACC8)
Eds. H.J. van den Herik and J.W.H.M. Uiterwijk
Drukkerij Van Spijk B.V. Venlo, The Netherlands, 332 pages.

Paderborn, Germany (1999, June)
Proceedings of the 9th Advances in Computer Games Conference (ACG9)
Eds. H.J. van den Herik and B. Monien
Van Spijk Grafisch Bedrijf Venlo, The Netherlands, 347 pages.

Graz, Austria (2003, November)
Proceedings of the 10th Advances in Computer Games Conference (ACG10)
Eds. H.J. van den Herik, H. Iida, and E.A. Heinz
Kluwer Academic Publishers, Boston/Dordrecht/London, 382 pages.

Taipei, Taiwan (2005, September)
Proceedings of the 11th Advances in Computer Games Conference (ACG11)
Eds. H.J. van den Herik, S-C. Hsu, T-s. Hsu, and H.H.L.M. Donkers
LNCS 4250, 372 pages.

Pamplona, Spain (2009, May)
Proceedings of the 12th Advances in Computer Games Conference (ACG12)
Eds. H.J. van den Herik and P. Spronck
LNCS 6048, 231 pages.

Tilburg, The Netherlands (2011, November)
Proceedings of the 13th Advances in Computer Games Conference (ACG13)
Eds. H.J. van den Herik and A. Plaat
LNCS 7168, 356 pages.

The Computers and Games Books

The series of Computers and Games (CG) Conferences started in 1998 as a complement to the well-known series of conferences in Advances in Computer Chess (ACC). Since 1998, seven CG conferences have been held. Below we list the conference places and dates together with the Springer publication (LNCS series number).

Tsukuba, Japan (1998, November)

Proceedings of the First Computers and Games Conference (CG98)

Eds. H.J. van den Herik and H. Iida

LNCS 1558, 335 pages.

Hamamatsu, Japan (2000, October)

Proceedings of the Second Computers and Games Conference (CG2000)

Eds. T.A. Marsland and I. Frank

LNCS 2063, 442 pages.

Edmonton, Canada (2002, July)

Proceedings of the Third Computers and Games Conference (CG2002)

Eds. J. Schaeffer, M. Müller, and Y. Björnsson

LNCS 2883, 431 pages.

Ramat-Gan, Israel (2004, July)

Proceedings of the 4th Computers and Games Conference (CG2004)

Eds. H.J. van den Herik, Y. Björnsson, and N.S. Netanyahu

LNCS 3846, 404 pages.

Turin, Italy (2006, May)

Proceedings of the 5th Computers and Games Conference (CG2006)

Eds. H.J. van den Herik, P. Ciancarini, and H.H.L.M. Donkers

LNCS 4630, 283 pages.

Beijing, China (2008, September)

Proceedings of the 6th Computers and Games Conference (CG2008)

Eds. H.J. van den Herik, X. Xu, Z. Ma, and M.H.M. Winands

LNCS 5131, 275 pages.

Kanazawa, Japan (2010, September)

Proceedings of the 7th Computers and Games Conference (CG2010)

Eds. H. J. van den Herik, H. Iida, and A. Plaat

LNCS 6515, 275 pages.

Table of Contents

Accelerated UCT and Its Application to Two-Player Games	1
<i>Junichi Hashimoto, Akihiro Kishimoto, Kazuki Yoshizoe, and Kokolo Ikeda</i>	
Revisiting Move Groups in Monte-Carlo Tree Search	13
<i>Gabriel Van Eyck and Martin Müller</i>	
PACHI: State of the Art Open Source Go Program	24
<i>Petr Baudiš and Jean-loup Gailly</i>	
Time Management for Monte-Carlo Tree Search in Go	39
<i>Hendrik Baier and Mark H.M. Winands</i>	
An MCTS Program to Play EinStein Würfelt Nicht!	52
<i>Richard J. Lorentz</i>	
Monte-Carlo Tree Search Enhancements for Havannah	60
<i>Jan A. Stankiewicz, Mark H.M. Winands, and Jos W.H.M. Uiterwijk</i>	
Payout Search for Monte-Carlo Tree Search in Multi-player Games	72
<i>J. (Pim) A.M. Nijssen and Mark H.M. Winands</i>	
Towards a Solution of 7x7 Go with Meta-MCTS	84
<i>Cheng-Wei Chou, Ping-Chiang Chou, Hassen Doghmen, Chang-Shing Lee, Tsan-Cheng Su, Fabien Teytaud, Olivier Teytaud, Hui-Ming Wang, Mei-Hui Wang, Li-Wen Wu, and Shi-Jim Yen</i>	
MCTS Experiments on the Voronoi Game	96
<i>Bruno Bouzy, Marc Métivier, and Damien Pellier</i>	
4*4-Pattern and Bayesian Learning in Monte-Carlo Go	108
<i>Jiao Wang, Shiyuan Li, Jitong Chen, Xin Wei, Huizhan Lv, and Xinhe Xu</i>	
Temporal Difference Learning for Connect6	121
<i>I-Chen Wu, Hsin-Ti Tsai, Hung-Hsuan Lin, Yi-Shan Lin, Chieh-Min Chang, and Ping-Hung Lin</i>	
Improving Temporal Difference Learning Performance in Backgammon Variants	134
<i>Nikolaos Papahristou and Ioannis Refanidis</i>	

CLOP: Confident Local Optimization for Noisy Black-Box Parameter Tuning	146
<i>Rémi Coulom</i>	
Analysis of Evaluation-Function Learning by Comparison of Sibling Nodes	158
<i>Tomoyuki Kaneko and Kunihito Hoki</i>	
Approximating Optimal Dudo Play with Fixed-Strategy Iteration Counterfactual Regret Minimization	170
<i>Todd W. Neller and Steven Hnath</i>	
The Global Landscape of Objective Functions for the Optimization of Shogi Piece Values with a Game-Tree Search	184
<i>Kunihito Hoki and Tomoyuki Kaneko</i>	
Solving BREAKTHROUGH with Race Patterns and Job-Level Proof Number Search	196
<i>Abdallah Saffidine, Nicolas Jouandeau, and Tristan Cazenave</i>	
Infinite Connect-Four Is Solved: Draw	208
<i>Yoshiaki Yamaguchi, Kazunori Yamaguchi, Tetsuro Tanaka, and Tomoyuki Kaneko</i>	
Blunder Cost in Go and Hex	220
<i>Henry Brausen, Ryan B. Hayward, Martin Müller, Abdul Qadir, and David Spies</i>	
Understanding Distributions of Chess Performances	230
<i>Kenneth W. Regan, Bartłomiej Maciejka, and Guy M^cC. Haworth</i>	
Position Criticality in Chess Endgames	244
<i>Guy M^cC. Haworth and Á. Rusz</i>	
On Board-Filling Games with Random-Turn Order and Monte Carlo Perfectness	258
<i>Ingo Althöfer</i>	
Modeling Games with the Help of Quantified Integer Linear Programs	270
<i>Thorsten Ederer, Ulf Lorenz, Thomas Opfer, and Jan Wolf</i>	
Computing Strong Game-Theoretic Strategies in Jotto	282
<i>Sam Ganzfried</i>	
Online Sparse Bandit for Card Games	295
<i>David L. St-Pierre, Quentin Louveaux, and Olivier Teytaud</i>	
Game Tree Search with Adaptive Resolution	306
<i>Hung-Jui Chang, Meng-Tsung Tsai, and Tsan-sheng Hsu</i>	

Designing Casanova: A Language for Games	320
<i>Giuseppe Maggiore, Alvisè Spanò, Renzo Orsini, Giulia Costantini, Michele Bugliesi, and Mohamed Abbadì</i>	
Affective Game Dialogues: Using Affect as an Explicit Input Method in Game Dialogue Systems	333
<i>Michael Lankes and Thomas Mirlacher</i>	
Generating Believable Virtual Characters Using Behavior Capture and Hidden Markov Models	342
<i>Richard Zhao and Duane Szafron</i>	
Author Index	355