Artificial and Computational Intelligence in Games

A Follow-up to Dagstuhl Seminar 12191

Edited by

Simon M. Lucas
Michael Mateas
Mike Preuss
Pieter Spronck
Julian Togelius

DAGSTUHL
FOLLOW-UPS

Editors

Simon M. Lucas
School of Computer Science and
Electronic Engineering
University of Essex
sml@essex.ac.uk

Michael Mateas
Center for Games and Playable Media
University of California, Santa Cruz
michaelm@cs.ucsc.edu

Mike Preuss
European Research Center for
Information Systems
University of Münster
mike.preuss@uni-muenster.de

Pieter Spronck
Tilburg Center for Cognition and Communication
Tilburg University
p.spronck@uvt.nl

Julian Togelius
Center for Computer Games Research
IT University of Copenhagen
julian@togelius.com

ACM Classification 1998
I.2.1 Applications and Expert Systems: Games

ISBN 978-3-939897-62-0

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at http://www.dagstuhl.de/dagpub/978-3-939897-62-0.

Publication date
November, 2013

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Digital Object Identifier: 10.4230/DFU.Vol6.12191.i

DFU – Dagstuhl Follow-Ups

The series *Dagstuhl Follow-Ups* is a publication format which offers a frame for the publication of peer-reviewed papers based on Dagstuhl Seminars. DFU volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

ISSN 1868-8977

www.dagstuhl.de/dfu

Contents

Artificial and Computational Intelligence in Games. *Dagstuhl Follow-Ups*, Volume 6, ISBN 978-3-939897-62-0.
Editors: Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius
DAGSTUHL FOLLOW-UPS Dagstuhl Publishing
Schloss Dagstuhl – Leibniz Center for Informatics, Germany

Preface

In May 2012, around 40 world-leading experts convened in Schloss Dagstuhl in Saarland, Southern Germany, to discuss future research directions and important research challenges for artificial and computational intelligence in games. The volume you are now reading is the follow-up volume to that seminar, which collects the distilled results of the discussions that went on during those May days. As organisers of the seminar and editors of the follow-up volume, it is our sincere hope that the chapters you are about to read will prove to be useful both as references for your existing research and as starting points for new research projects. In this introductory chapter, we give some background on the research field, describe how the seminar was held, and briefly present each of the following chapters.

The research field

Research into artificial intelligence (AI) and computational intelligence (CI)[1] started in the 1950s and has been growing ever since. The main focus of the research is to provide computers with the capacity to perform tasks that are believed to require human intelligence. The field is in constant need of good benchmark problems. All benchmark problems have their drawbacks – for instance, abstract mathematical problems might not be relevant to the real world, and complex robotics problems can be time- and resource-consuming. In the last decade, computer games have been considered a strong source of benchmark problems for human intelligence. Humans have played games for all of recorded history, and since 1980s video games have been a favourite pastime of people all over the world. Games provide a plethora of tough and interesting challenges for AI and CI, including developing artificial players, computationally creative systems that construct game content, agents that adapt to players, and systems that analyse and learn about players from their in-game behaviour. As games are developed to be challenging and interesting to humans, many game-related AI problems are relevant to understanding human cognition and creativity as well. Additionally, game design and development have a number of outstanding problems that could be solved by better and more appropriate AI, so there is an opportunity to make a contribution to real-world problems.

The study of AI and CI as applied to video games is rather new. A research field devoted to these topics has only started to coalesce within the last 8 years around the AAAI Artificial Intelligence and Interactive Digital Entertainment (AIIDE) and IEEE Computational Intelligence and Games (CIG) conferences[2]. The research field is not yet defined well, and the research community has been somewhat fractured along both the symbolic/non-symbolic axis and along the European/American axis. A set of common problems and a common terminology needed to be established, and ideas needed to be cross-fertilised. The Dagstuhl seminar which this volume is an outcome of was held in order to address these challenges.

[1] The terms AI and CI are here used more or less interchangeably, even though there is a historic divide in terms of both methods studied and membership of the research communities. In general, CI methods are more biologically inspired or statistical, whereas AI methods are more symbolical or logical, but there is a great deal of overlap.

[2] AI and CI research into classic board games has a longer history, and is concentrated around the ICGA Computer Games (CG) and Advances in Computer Games (ACG) conferences.

Artificial and Computational Intelligence in Games. *Dagstuhl Follow-Ups*, Volume 6, ISBN 978-3-939897-62-0.
Editors: Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius
DAGSTUHL FOLLOW-UPS Dagstuhl Publishing
Schloss Dagstuhl – Leibniz Center for Informatics, Germany

The seminar and the follow-up volume

In May 2012, we gathered around 40 researchers and practitioners of AI and CI in video games at Schloss Dagstuhl to discuss the challenges in the field and the approaches to make progress on important problems. The seminar was arranged so that participants formed work groups consisting of 3–10 people each, all experts on the topic of the group. The groups discussed their topic for one or two days, and then reported the consensus of their findings to the rest of the seminar. We explicitly instructed attendants to focus on the future of the field rather than the past, and to form groups around problems rather than particular methods. With so many world-class experts on these topics gathered in a single seminar to discuss the challenges of the future, it would have been a sin to not publish the outcome of the proceedings so it would be accessible to other researchers in the area.

Soon after the seminar, we published a Dagstuhl Report containing abstracts of the discussions in all groups. We then invited group members to write longer chapters on each topic, chapters which would include both surveys of the existing literature and discussions of future challenges, and thus serve as guides to the literature as well as starting points for new research projects. Unfortunately, not all groups wrote up a full chapter-length description of their conclusions for various reasons (mostly lack of time on part of the authors). The work groups that did not write full chapters were those on believable agents and social simulations, AI architectures for games, AI for modern board games, evaluating AI in games research, AIGameResearch.org and computational narrative. However, eight of the groups did write full chapters, and those are collected in this follow-up volume.

The chapters

This volume consists of eight chapters in addition to the Introduction you are currently reading. Each chapter is the outcome of a particular discussion group that met for two days, and wrote the chapter in the months after the symposium. To assure quality, single-blind peer review was carried out by other attendees of the seminar, and the final versions of the chapters have been edited to address the reviewers' concerns. In editing this volume, we have chosen to arrange the chapters so that they start with the more generic problems and methods and proceed to more specific applications. However, the web of interdependence between work on these topics is dense, with games for mobile platforms relying on pathfinding, general game playing on procedural content generation, procedural content generation on player modelling, etc.

Search in Real-Time Video Games

Almost every AI and CI technique can be seen as search in some way: search for solutions, paths, strategies, models, proofs, actions etc. Historically, the first applications of AI techniques to games – in particular, methods for playing board games – were all about searching the tree of possible actions and counter-actions, the "game tree". It therefore seems suitable to start this volume with a chapter about search. This chapter outlines the main types of search that are used in games today, and then proceeds to discuss the main challenges that search algorithms face when applied to real-time video games as opposed to board games. Arguably the greatest challenge is that real-time video games typically have near-continuous state and action space, so that the number of states and the number of actions that could be taken from each state (the branching factor) is enormous in comparison to the one for traditional board games. Also, the number of moves required to be taken before the game reaches an end point with a natural reward function is much larger. However, concepts as game tree and branching factor cannot easily be applied to real-time games as moves are neither strictly sequential nor necessarily alternating between players. We find that even measuring the hardness of these problems is difficult. Approaches to overcoming the challenges include clustering or partitioning states, and statistical approaches such as Monte Carlo methods.

Pathfinding in Games

Pathfinding is a particular kind of search, where the objective is to find the shortest path (according to some metric) between two points. The workgroup (and thus the chapter) on pathfinding was motivated by the central importance of pathfinding for most video games. Algorithms for pathfinding consume a considerable amount of processing power in modern games, and whereas pathfinding might be considered a "solved problem" in some areas of AI, it most certainly is not in computer games. Inferior pathfinding is a substantial problem in published commercial computer games, with their complex dynamic environments and real-time processing requirements. Recent advances in pathfinding include path computation methods based on hierarchical abstractions, informed memory-based heuristic functions, symmetry reduction and triangulation-based map representations. The chapter also outlines future research challenges, which mainly relate to the following three subjects: (1) the dynamic nature of game maps, which can change at any time with, for instance, a destructible environment; (2) the sheer size of game maps coupled with memory limitations of game consoles; and (3) collaborative pathfinding for multiple agents.

Learning and Game AI

Machine learning is a very active research field in its own right, with a large number of applications in various domains. It would seem natural for an academic researcher to think that there were ample applications for learning algorithms in computer games. However, it is rather rare to see machine learning used in any published games, and commercial game developers tend to view such methods with utmost suspicion. The chapter on Learning and Game AI goes through some of the many potential applications for machine learning in games, such as balancing games, balancing players and finding design loopholes. The chapter also discusses some of the considerable challenges that are impeding the adoption of learning algorithms by the game industry, including explaining the models induced by learning algorithms to designers and players, the problem with finding good reward functions that reflect the quality of the game, and the high computational cost of many learning algorithms.

Player Modelling

Player modelling is a specific application of machine learning to games, where the goal is to model the behaviour, preferences or experience of the player. A central question tackled in the chapter on player modelling is to what extent an accurate model of a player can be constructed based on observations of the player's behaviour in a game, potentially enriched by information from demographics, questionnaires and psychophysiological measurements. With the growing amount of networking that game players engage in, the potential to acquire data for building player models is increasing all the time. Many modern computer games now "phone home" and report detailed information to their developers' servers about their players.

This information might be used to make games more entertaining and captivating for their players, which in turn translates to revenue for their developers. Researchers in CI and AI have much to contribute here, given the plethora of methods that have been developed in academia for similar problems. In the chapter, approaches based on unsupervised learning are contrasted with "theory-based" approaches based on ideas and models from psychology. One conclusion is that while it is relatively easy to create models for populations of players which predict how that population will respond, it is quite hard to create a model for an individual player, that explains and can anticipate that player's behaviour. A potential solution is to create dynamic models, i.e., a model for a player that is constantly evaluated and dynamically adapted to observations of the player, and player responses.

Procedural Content Generation: Goals, Challenges and Actionable Steps

The chapter on procedural content generation discusses methods for automatically generating game content such as maps, levels, items, quests and game rules. This is a set of important problems in game development, both to alleviate production costs for game content and to make new kinds of games that feature infinite and perhaps adaptive content generation. In recognition of this, procedural content generation has recently become one of the most active research topics within the CI/AI and games field. The chapter proposes a vision of a system that can make a complete game world, including characters, buildings, quests, and items for a given game-engine at the press of a button. To reach this vision, a number of technical and conceptual challenges need to be overcome; the list includes items such as better understanding the process of searching for game content artifacts, and representing artistic style in a content generator. Further, the chapter lists a number of specific projects which could be undertaken right away, and would contribute to addressing the main challenges for PCG and ultimately to realising the grand vision.

General Video Game Playing

Much game AI is developed to work specifically with one game, meaning that the solutions developed might be narrow in scope and contributing little to the greater problem of general artificial intelligence. In the small field of *general game playing*, approaches to creating agents that can proficiently play a wide range of games is studied. However, in actual practice general game playing tends to focus on variations of board games. This is to a large extent due to the field's focus on the General Game Playing Competition, where AI agents are tested on unseen games which in practice all are rather simple board games. The chapter on general video game playing argues for the importance of testing agents on multiple unseen games, but that these need to be more complex and multi-faceted than the games which have hitherto been used for general game playing. In particular, video games provide plenty of challenges related to coordination, timing, and navigation, that board games do not provide. The chapter proposes that one necessary component of a general video game playing system would be a language in which it is possible to specify complete video games, so that they can be generated by a special game engine.

Towards a Video Game Description Language

This chapter continues where the chapter on general video game playing left off and puts forward a concrete suggestion for a language that can specify simple video games. The language was developed by analysing three simple 2D arcade games from the early 1980's –

Frogger, Lunar Lander and Space Invaders – and finding a vocabulary for describing their common parts. It was found that a language could be structured around individual game objects, and defining their movement logics and the effects of collisions. The chapter includes sketches of how the three aforementioned games would be implemented in the proposed language.

Artificial and Computational Intelligence for Games on Mobile Platforms

The final chapter addresses the particular challenges and opportunities arising when using AI and CI in games for mobile platforms, such as smartphones and tablets. Compared to desktop and laptop computers, these devices typically have limitations in terms of battery life, processing power and screen size. However, they also have several features that are not usually part of conventional systems, such as location awareness, personal ownership and relatively low demands on graphics. The chapter on AI and CI for games on mobile platforms argues that these features make such platforms very well suited for experimentation with new AI-based game designs, especially those based on procedural content generation, personalisation and ubiquitous gaming.

Conclusions

The 2012 gathering at Schloss Dagstuhl was deemed a great success by all participants, and it drew a large part of this strength out of the agile and very adaptive style it was held in, with several unforeseen developments in themes and results. This follow-up volume exemplifies the high level of the scientific discussions and the strong focus on scientific progress of the seminar as a whole. We are pleased to announce that a follow-up seminar will be organized at Schloss Dagstuhl in 2014. Whereas the 2012 seminar treated the topics discussed as separate research areas, the 2014 seminar will focus on the integration of the various research fields. This is meant to achieve faster developments, improve visibility and acceptance of our algorithms and approaches in industry and open up new areas of research. We believe that integration is an exciting as well as necessary step in order to further shape and consolidate the research field of artificial and computational intelligence in games.

List of Authors

Elisabeth André
University of Augsburg
Germany
andre@informatik.uni-augsburg.de

Clare Bates Congdon
University of Southern Maine
USA
congdon@usm.maine.edu

Christian Bauckhage
Fraunhofer IAIS, Sankt Augustin
Germany
christian.bauckhage@iais.fraunhofer.de

Michael Buro
University of Alberta
Canada
mburo@cs.ualberta.ca

Michal Bida
Charles University of Prague
Czech Republic
Michal.Bida@mff.cuni.cz

Adi Botea
IBM Research, Dublin
Ireland
adibotea@ie.ibm.com

Bruno Bouzy
Université Paris Descartes
France
bruno.bouzy@parisdescartes.fr

Martin V. Butz
Eberhard Karls Universität Tübingen
Germany
butz@informatik.uni-tuebingen.de

Alex J. Champandard
AiGameDev.com, Vienna
Austria
alexjc@aigamedev.com

Peter I. Cowling
University of York
United Kingdom
peter.cowling@york.ac.uk

Marc Ebner
Ernst Moritz Arndt Universität, Greifswald
Germany
marc.ebner@uni-greifswald.de

Philip Hingston
Edith Cowan University
Australia
p.hingston@ecu.edu.au

Graham Kendall
University of Nottingham
United Kingdom and Malaysia
graham.kendall@nottingham.ac.uk

Pier Luca Lanzi
Politecnico di Milano
Italy
lanzi@elet.polimi.it

John Levine
University of Strathclyde
United Kingdom
john.levine@strath.ac.uk

Daniele Loiacono
Politecnico di Milano
Italy
daniele.loiacono@polimi.it

Michael Mateas
University of California, Santa Cruz
USA
michaelm@cs.ucsc.edu

Risto Miikkulainen
University of Texas, Austin
USA
risto@cs.utexas.edu

Héctor Muñoz-Avila
Lehigh University
USA
munoz@eecs.lehigh.edu

Dana Nau
University of Maryland
USA
nau@cs.umd.edu

Artificial and Computational Intelligence in Games. *Dagstuhl Follow-Ups*, Volume 6, ISBN 978-3-939897-62-0.
Editors: Simon M. Lucas, Michael Mateas, Mike Preuss, Pieter Spronck, and Julian Togelius
DAGSTUHL FOLLOW-UPS Dagstuhl Publishing
Schloss Dagstuhl – Leibniz Center for Informatics, Germany

Ana Paiva
INESC-ID, Lisboa
Portugal
ana.paive@inesc-id.pt

Mike Preuss
Technical University of Dortmund
Germany
mike.preuss@cs.tu-dortmund.de

Tom Schaul
New York University
USA
schaul@cims.nyu.edu

Moshe Sipper
Ben-Gurion University
Israel
sipper@cs.bgu.ac.il

Pieter Spronck
Tilburg University
The Netherlands
p.spronck@uvt.nl

Kenneth O. Stanley
University of Central Florida
USA
kstanley@eecs.ucf.edu

Tommy Thompson
University of Derby
United Kingdom
t.thompson@derby.ac.uk

Julian Togelius
IT University of Copenhagen
Denmark
julian@togelius.com

Georgios N. Yannakakis
University of Malta, Msida
Malta
georgios.yannakakis@um.edu.mt