

Artificial Intelligence and Games

Georgios N. Yannakakis
Julian Togelius

Artificial Intelligence and Games

 Springer

Georgios N. Yannakakis
Institute of Digital Games
University of Malta
Msida
Malta

Julian Togelius
Department of Computer Science
and Engineering
New York University
Brooklyn, NY
USA

ISBN 978-3-319-63518-7 ISBN 978-3-319-63519-4 (eBook)
<https://doi.org/10.1007/978-3-319-63519-4>

Library of Congress Control Number: 2018932540

© Springer International Publishing AG, part of Springer Nature 2018

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

The use of general descriptive names, registered names, trademarks, service marks, 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.

The publisher, the authors and the editors are safe to assume that the advice and information in this book are believed to be true and accurate at the date of publication. Neither the publisher nor the authors or the editors give a warranty, express or implied, with respect to the material contained herein or for any errors or omissions that may have been made. The publisher remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Cover image by Rebecca Portelli

Printed on acid-free paper

This Springer imprint is published by the registered company Springer
International Publishing AG part of Springer Nature
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland



To our families

Foreword

It is my great pleasure to write the foreword for this excellent and timely book. Games have long been seen as the perfect test-bed for artificial intelligence (AI) methods, and are also becoming an increasingly important application area. Game AI is a broad field, covering everything from the challenge of making super-human AI for difficult games such as Go or *StarCraft*, to creative applications such as the automated generation of novel games.

Game AI is as old as AI itself, but over the last decade the field has seen massive expansion and enrichment with the inclusion of video games, which now comprise more than 50% of all published work in the area and enable us to address a broader range of challenges that have great commercial, social, economic and scientific interest. A great surge in research output occurred in 2005, coinciding with both the first IEEE Symposium (Conference) on Computational Intelligence and Games (CIG)—which I co-chaired with Graham Kendall—and the first AAAI AIIDE Conference (Artificial Intelligence in Digital Entertainment). Since then this rich area of research has been more explored and better understood. The Game AI community pioneered much of the research which is now becoming (or about to become) more mainstream AI, such as Monte Carlo Tree Search, procedural content generation, playing games based on screen capture, and automated game design.

Over the last decade, progress in deep learning has had a profound and transformational effect on many difficult problems, including speech recognition, machine translation, natural language understanding and computer vision. As a result, computers can now achieve human-competitive performance in a wide range of perception and recognition tasks. Many of these systems are now available to the programmer via a range of so-called cognitive services. More recently, deep reinforcement learning has achieved ground-breaking success in a number of difficult challenges, including Go and the amazing feat of learning to play games directly from screen capture (playing from pixels). It is fascinating to contemplate what this could mean for games as we stumble towards human-level intelligence in an increasing number of areas. The impacts will be significant for the intelligence of in-game characters, the way in which we interact with them and for the way games are designed and tested.

This book makes an enormous contribution to this captivating, vibrant area of study: an area that is developing rapidly both in breadth and depth as AI is able to cope with a wider range of tasks (and to perform those tasks to increasing levels of excellence). The service to the community will be felt for many years to come: the book provides an easier and more comprehensive entry point for newcomers to the field than previously available, whilst also providing an indispensable reference for existing AI and Games researchers wishing to learn about topics outside their direct field of interest.

Georgios Yannakakis and Julian Togelius have been involved with the field ever since its widespread expansion to video games, and they both presented papers at the first 2005 CIG. Over the years they have made an enormous contribution to the field with a great number of highly cited papers presenting both novel research and comprehensive surveys. It is my opinion that these authors are best qualified to write this book, and they do not disappoint. The book will serve the community very well for many years to come.

London, August 2017

Simon Lucas

Preface

Of all the things that wisdom provides for the complete happiness of one's entire life, by far the greatest is friendship.

Epicurus, *Principal Doctrines*, 27

Human beings, viewed as behaving systems, are quite simple. The apparent complexity of our behavior over time is largely a reflection of the complexity of the environment in which we find ourselves.

Herbert A. Simon

It would be an understatement to say that **Artificial Intelligence** (AI) is a popular topic at the moment, and it is unlikely to become any less important in the future. More researchers than ever work on AI in some form, and more non-researchers than ever are interested in the field. It would also be an understatement to say that **games** are a popular application area for AI research. While board games have been central to AI research since the inception of the field, video games have during the last decade increasingly become the domain of choice for testing and showcasing new algorithms. At the same time, video games themselves have become more diverse and sophisticated, and some of them incorporate advances in AI for controlling non-player characters, generating content or adapting to players. Game developers have increasingly realized the power of AI methods to analyze large volumes of player data and optimize game designs. And a small but growing community of researchers and designers experiment with ways of using AI to design and create complete games, automatically or in dialog with humans. It is indeed an exciting time to be working on AI and games!

This is a book about **AI and games**. As far as we know, it is the first *comprehensive* textbook covering the field. With *comprehensive*, we mean that it features all the major application areas of AI methods within games: game-playing, content generation and player modeling. We also mean that it discusses AI problems in many different types of games, including board games and video games of many genres. The book is also *comprehensive* in that it takes multiple perspectives of AI and games: how games can be used to test and develop AI, how AI can be used

to make games better and easier to develop, and to understand players and design. While this is an academic book which is primarily aimed at students and researchers, we will frequently address problems and methods relevant for game designers and developers.

We wrote this book based on our long experience doing research on AI for games, each on our own and together, and helping lead and shape the research community. We both independently started researching AI methods in games in 2004, and we have been working together since 2009. Together, we played a role in introducing research topics such as procedural content generation and player modeling to the academic research community, and created several of the most widely used game-based AI benchmarks. This book is in a sense a natural outgrowth of the classes on AI and games we have taught at three universities, and the several survey papers of the field and of individual research topics within it that we have published over the years. But the book is also a response to the lack of a good introductory book for the research field. Early discussions on writing such a book date back at least a decade, but no-one actually wrote one, until now.

It could be useful to point out what this book is not. It is not a hands-on book with step-by-step instructions on how to build AI for your game. It does not feature discussions of any particular game engine or software framework, and it does not discuss software engineering aspects or many implementation aspects at all. It is not an introductory book, and it does not give a gentle introduction to basic AI or game design concepts. For all these roles, there are better books available.

Instead, this is a book for readers who already understand AI methods and concepts to the level of having taken an introductory AI course, and the introductory computer science or engineering courses that led up to that course. The book assumes that the reader is comfortable with reading a pseudocode description of an algorithm and implementing it. Chapter 2 is a summary of AI methods used in the book, but is intended more as a reference and refresher than as an introduction. The book also assumes a basic familiarity with games, if not designing them then at least playing them.

The use case for this textbook that we had in mind when writing it is for a one- or a two-semester graduate-level or advanced undergraduate level class. This can take several different shapes to support different pedagogical practices. One way of teaching such a class would be a traditional class, with lectures covering the chapters of the book in order, a conventional pen-and-paper exam at the end, and a small handful of programming exercises. For your convenience, each of the main chapters of the book include suggestions for such exercises. Another way of organizing a class around this book, more in line with how we personally prefer to teach such courses, is to teach the course material during the first half of the semester and spend the second half on a group project.

The material offered by this book can be used in various ways and, thus, support a number of different classes. In our experience, a traditional two-semester class on game artificial intelligence would normally cover Chapter 2 and Chapter 3 in the first semester and then focus on alternative uses of AI in games (Chapters 4 and 5) in the second semester. When teaching the material in compressed (one-semester) fash-

ion instead, it is advisable to skip Chapter 2 (using it as a reference when needed), and focus the majority of the lectures on Chapters 3, 4 and 5. Chapters 6 and 7 can be used as material for inspiring advanced graduate-level projects in the area. Beyond the strict limits of game AI, Chapter 4 (or sections of it) can complement classes with a focus on game design or computational creativity whereas Chapter 5 can complement classes with a focus on affective computing, user experience, and data mining. It is of course also possible to use this book for an introductory undergraduate class for students who have not taken an AI class before, but in that case we advise the instructor to select a small subset of topics to focus on, and to complement the book with online tutorials on specific methods (e.g., best-first search, evolutionary computation) that introduce these topics in a more gentle fashion than this book does.

Chania, Crete, Greece
New York, NY, USA

Georgios N. Yannakakis
Julian Togelius

September 2017

Acknowledgments

Writing a book of this size would be impossible without the support and the contributions of a large number of people. First, we would like to thank Ronan Nugent, our editor at Springer, who guided us and helped us from the book proposal phase all the way to its final production.

We also would like to thank the following persons who read all (or parts of a draft of) the book and provided useful feedback: Amy Hoover, Amin Babadi, Sander Bakkes, Vadim Bulitko, Phil Carlisle, Georgios Chalkiadakis, Dave Churchill, Mike Cook, Renato Cunha, Kevin Dill, Nathaniel Du Preez-Wilkinson, Chris Elion, Andy Elmsley, David Fogel, Bernardo Galvão, Kazu-ma Hashimoto, Aaron Isaksen, Emil Johansen, Mark Jones, Niels Justesen, Graham Kendall, Jakub Kowalski, Antonios Liapis, Nir Lipovetzky, Jhovan Mauricio López, Simon Lucas, Jacek Mańdziuk, Luciana Mariñelarena-Dondena, Chris Martens, Sean Mealin, Mark Nelson, Sven Neuhaus, Alexander Osherenko, Santiago Ontañón, Cale Plut, Mike Preuss, Hartmut Procha-ska, Christoffer Holmgård, Florian Richoux, Sebastian Risi, Christoph Salge, Andrea Schiel, Jacob Schrum, Magy Seif El-Nasr, Adam Smith, Gillian Smith, Dennis Soemers, Nathan Sturtevant, Gabriel Synnaeve, Nicolas Szilas, Varunyu Vorachart, James Wen, Marco Wiering, Mark Winands, Junkai Lu, Francesco Calimeri, Diego Pérez Liébana, Corine Jacobs, Junkai Lu, Hana Rudova, and Robert Zubek. Of these, we wish to especially thank Simon Lucas for writing the foreword of the book, Georgios Chalkiadakis for providing substantial input on parts of the book related to game theory, and Mark Nelson, Antonios Liapis, Mike Preuss and Adam Smith for reviewing large parts of the book and providing particularly detailed feedback. We also want to thank all those who granted us their permission to reproduce images and figures from their papers; they are all acknowledged in the figure captions of the book. Special thanks also go to Rebecca Portelli and Daniel Mercieca for the artwork featured on the cover page of the book.

Some chapters of this book build on papers or other book chapters that we have co-authored. In some cases the papers are co-authored by more than the two of us; for those papers our co-authors graciously gave us permission to reuse parts of the material and we wish to thank them for that. In particular,

- Chapter 1: [764, 700].
- Chapter 4: Chapter 2 and Chapter 3 from [616], and [381].
- Chapter 5: [778, 176, 782, 781].
- Chapter 6: [785].
- Chapter 7: [718, 458].

Writing this book has been a very long journey for both of us; a challenging and overwhelming journey at times. There are a number of people who have supported this effort that we would like to thank. Georgios wishes to thank the people at NYU Tandon School of Engineering for hosting him during the early days of book planning and the people at the Technical University of Crete for hosting him during the later stages of this project. Georgios would also like to thank the University of Malta for granting him a sabbatical leave, without which the book would not have been possible. Georgios and Julian also wish to thank each other for putting up with each other despite everything, and announce that they intend to buy each other drinks to celebrate when they meet next.

Last but not least, both of us wish to thank our families and all the people who showed us their support, care, encouragement and love at times when they were absolutely needed. You are too many to list here, but big thanks go to you all! Georgios would not have been able to write this book without the love and support of his family: Amyrsa and Myrto have been his core inspiration, Stavroula has been the main driving force behind the writing at all times; this book is dedicated to you!

Contents

Part I Background

1	Introduction	3
1.1	This Book	4
1.1.1	Why Did We Write This Book?	5
1.1.2	Who Should Read This Book?	6
1.1.3	A Short Note on Terminology	7
1.2	A Brief History of Artificial Intelligence and Games	8
1.2.1	Academia	10
1.2.2	Industry	11
1.2.3	The “Gap”	13
1.3	Why Games for Artificial Intelligence	15
1.3.1	Games Are Hard and Interesting Problems	15
1.3.2	Rich Human-Computer Interaction	17
1.3.3	Games Are Popular	18
1.3.4	There Are Challenges for All AI Areas	20
1.3.5	Games Best Realize Long-Term Goals of AI	21
1.4	Why Artificial Intelligence for Games	23
1.4.1	AI Plays and Improves Your Game	23
1.4.2	More Content, Better Content	24
1.4.3	Player Experience and Behavioral Data Analytics	25
1.5	Structure of This Book	25
1.5.1	What We (Don’t) Cover in This Book	26
1.6	Summary	28
2	AI Methods	29
2.1	General Notes	30
2.1.1	Representation	30
2.1.2	Utility	31
2.1.3	Learning = Maximize Utility (Representation)	32
2.2	Ad-Hoc Behavior Authoring	32

2.2.1	Finite State Machines	33
2.2.2	Behavior Trees	34
2.2.3	Utility-Based AI	37
2.2.4	Further Reading	39
2.3	Tree Search	39
2.3.1	Uninformed Search	40
2.3.2	Best-First Search	41
2.3.3	Minimax	42
2.3.4	Monte Carlo Tree Search	45
2.3.5	Further Reading	49
2.4	Evolutionary Computation	49
2.4.1	Local Search	50
2.4.2	Evolutionary Algorithms	52
2.4.3	Further Reading	57
2.5	Supervised Learning	57
2.5.1	Artificial Neural Networks	59
2.5.2	Support Vector Machines	66
2.5.3	Decision Tree Learning	68
2.5.4	Further Reading	71
2.6	Reinforcement Learning	71
2.6.1	Core Concepts and a High-Level Taxonomy	73
2.6.2	Q-Learning	74
2.6.3	Further Reading	76
2.7	Unsupervised Learning	77
2.7.1	Clustering	77
2.7.2	Frequent Pattern Mining	80
2.7.3	Further Reading	82
2.8	Notable Hybrid Algorithms	82
2.8.1	Neuroevolution	83
2.8.2	TD Learning with ANN Function Approximators	84
2.8.3	Further Reading	87
2.9	Summary	88

Part II Ways of Using AI in Games

3	Playing Games	91
3.1	Why Use AI to Play Games?	92
3.1.1	Playing to Win in the Player Role	93
3.1.2	Playing to Win in a Non-player Role	95
3.1.3	Playing for Experience in the Player Role	95
3.1.4	Playing for Experience in a Non-player Role	96
3.1.5	Summary of AI Game-Playing Goals and Roles	98
3.2	Game Design and AI Design Considerations	98
3.2.1	Characteristics of Games	98
3.2.2	Characteristics of AI Algorithm Design	105

3.3	How Can AI Play Games?	109
3.3.1	Planning-Based Approaches	109
3.3.2	Reinforcement Learning	115
3.3.3	Supervised Learning	118
3.3.4	Chimeric Game Players	118
3.4	Which Games Can AI Play?	119
3.4.1	Board Games	119
3.4.2	Card Games	121
3.4.3	Classic Arcade Games	123
3.4.4	Strategy Games	132
3.4.5	Racing Games	136
3.4.6	Shooters and Other First-Person Games	139
3.4.7	Serious Games	141
3.4.8	Interactive Fiction	144
3.4.9	Other Games	145
3.5	Further Reading	148
3.6	Exercises	149
3.6.1	Why Ms Pac-Man?	149
3.7	Summary	150
4	Generating Content	151
4.1	Why Generate Content?	152
4.2	Taxonomy	154
4.2.1	Taxonomy for Content	154
4.2.2	Taxonomy for Methods	155
4.2.3	Taxonomy of Roles	156
4.3	How Could We Generate Content?	157
4.3.1	Search-Based Methods	157
4.3.2	Solver-Based Methods	161
4.3.3	Grammar-Based Methods	162
4.3.4	Cellular Automata	165
4.3.5	Noise and Fractals	169
4.3.6	Machine Learning	171
4.4	Roles of PCG in Games	175
4.4.1	Mixed-Initiative	176
4.4.2	Autonomous	180
4.4.3	Experience-Driven	180
4.4.4	Experience-Agnostic	182
4.5	What Could Be Generated?	184
4.5.1	Levels and Maps	184
4.5.2	Visuals	186
4.5.3	Audio	187
4.5.4	Narrative	189
4.5.5	Rules and Mechanics	191
4.5.6	Games	193

4.6	Evaluating Content Generators	197
4.6.1	Why Is It Difficult?	197
4.6.2	Function vs. Aesthetics	197
4.6.3	How Can We Evaluate a Generator?	198
4.7	Further Reading	201
4.8	Exercises	201
4.8.1	Maze Generation	201
4.8.2	Platformer Level Generation	201
4.9	Summary	202
5	Modeling Players	203
5.1	What Player Modeling Is and What It Is Not	204
5.2	Why Model Players?	206
5.3	A High-Level Taxonomy of Approaches	207
5.3.1	Model-Based (Top-Down) Approaches	208
5.3.2	Model-Free (Bottom-Up) Approaches	212
5.3.3	Hybrids	213
5.4	What Is the Model's Input Like?	213
5.4.1	Gameplay	214
5.4.2	Objective	214
5.4.3	Game Context	217
5.4.4	Player Profile	218
5.4.5	Linked Data	219
5.5	What Is the Model's Output Like?	219
5.5.1	Modeling Behavior	220
5.5.2	Modeling Experience	220
5.5.3	No Output	229
5.6	How Can We Model Players?	230
5.6.1	Supervised Learning	230
5.6.2	Reinforcement Learning	235
5.6.3	Unsupervised Learning	237
5.7	What Can We Model?	238
5.7.1	Player Behavior	238
5.7.2	Player Experience	245
5.8	Further Reading	252
5.9	Exercises	252
5.9.1	Player Behavior	252
5.9.2	Player Experience	253
5.10	Summary	254

Part III The Road Ahead

6 Game AI Panorama 259

6.1 Panoramic Views of Game AI 260

6.1.1 Methods (Computer) Perspective 261

6.1.2 End User (Human) Perspective 262

6.1.3 Player-Game Interaction Perspective 264

6.2 How Game AI Areas Inform Each Other 264

6.2.1 Play Games 266

6.2.2 Generate Content 269

6.2.3 Model Players 273

6.3 The Road Ahead 275

6.4 Summary 277

7 Frontiers of Game AI Research 279

7.1 General General Game AI 279

7.1.1 General Play 281

7.1.2 General Game Generation and Orchestration 282

7.1.3 General Game Affective Loop 284

7.2 AI in Other Roles in Games 285

7.3 Ethical Considerations 287

7.4 Summary 290

References 293

Index 331

Acronyms

A3C	Asynchronous Advantage Actor-Critic
ABL	A Behavior Language
AFC	Alternative Forced Choice
AI	Artificial Intelligence
AIIDE	Artificial Intelligence and Interactive Digital Entertainment
ALE	Arcade Learning Environment
ANN	Artificial Neural Network
ASP	Answer Set Programming
BDI	Belief-Desire-Intention
BT	Behavior Tree
BWAPI	Brood War API
CA	Cellular Automata
CI	Computational Intelligence
CIG	Computational Intelligence and Games
CFR	Counterfactual Regret Minimization
CMA-ES	Covariance Matrix Adaptation Evolution Strategy
CNN	Convolutional Neural Network
CPPN	Compositional Pattern Producing Network
DQN	Deep Q Network
EA	Evolutionary Algorithm
ECG	Electrocardiography
EDPCG	Experience-Driven Procedural Content Generation
EEG	Electroencephalography
EMG	Electromyography
FPS	First-Person Shooter
FSM	Finite State Machine
FSMC	Functional Scaffolding for Musical Composition
GA	Genetic Algorithm
GDC	Game Developers Conference
GGP	General Game Playing
GSP	Generalized Sequential Patterns

GSR	Galvanic Skin Response
GVGAI	General Video Game Artificial Intelligence
HCI	Human-Computer Interaction
ID3	Iterative Dichotomiser 3
JPS	Jump Point Search
LSTM	Long Short-Term Memory
MCTS	Monte Carlo Tree Search
MDP	Markov Decision Process
MLP	Multi-Layer Perceptron
MOBA	Multiplayer Online Battle Arenas
NEAT	NeuroEvolution of Augmenting Topologies
NES	Natural Evolution Strategy
NLP	Natural Language Processing
NPC	Non-Player Character
PC	Player Character
PCG	Procedural Content Generation
PENS	Player Experience of Need Satisfaction
PLT	Preference Learning Toolbox
RBF	Radial Basis Function
ReLU	Rectified Linear Unit
RPG	Role-Playing Game
RTS	Real-Time Strategy
RL	Reinforcement learning
TCIAIG	Transactions on Computational Intelligence and AI in Games
TD	Temporal Difference
ToG	Transactions on Games
TORCS	The Open Racing Car Simulator
TRU	Tomb Raider: Underworld
TSP	Traveling Salesman Problem
SC:BW	StarCraft: Brood War
SOM	Self-Organizing Map
STRIPS	Stanford Research Institute Problem Solver
SVM	Support Vector Machine
UT2k4	Unreal Tournament 2004
VGDL	Video Game Description Language

Website

<http://gameaibook.org/>

This book is associated with the above website. The website complements the material covered in the book with up-to-date exercises, lecture slides and readings.