A Vision on Analysing Approaches for Knowledge Representation and Reasoning Using Computer Games

Christian Eichhorn, Vanessa Volz, Richard Niland, and Tim Schendekehl

Faculty of Computer Science, TU Dortmund University, Dortmund, Germany
`<firstname>.<lastname>@tu-dortmund.de`

Abstract. Artificial intelligences (AIs) that interact with their environment are difficult to compare and evaluate as their formal properties easily become incomparable due to fundamentally different knowledge representations, reaction schemes, approaches and the general, very dendritic field of AI research. Nonetheless, AI approaches are regularly proposed as solutions to complex "real world" problems in areas such as self-driving cars or providing care for the elderly. Thus, the need for a safe and controllable proving ground for different AI approaches with scalable complexity emerges. Many researchers have argued that this need can be fulfilled by using computer games as a testbed [17,16]. In this paper, we propose a benchmark that specifically targets areas in AI research that still pose great challenges in AI and human-computer interaction research: the coordination of and cooperation among agents. We thus introduce and present the platform game ZooOperation as well as the corresponding competition involving this game at the KI 2017 conference, and illustrate how ZooOperation can serve as testbed for the coordination and cooperation skills of various AI approaches. On top of this, we discuss how this game, and computer games in general, can be used in comparative AI research, e.g. in testing for robustness, generalisability and human-computer interaction.

1 Introduction

Games have historically served as a testbed for artificial intelligence (cf. Alan Turing's chess-playing algorithm [17]). We, like many other researchers [16], argue that games continue to provide a great test environment for AI in general and both knowledge representation and reasoning approaches in particular, especially if adapted according to the (intended) "real world" AI applications.

Evidently, the increasing complexity of game benchmarks (Go [14], StarCraft [10]) has resulted in the advent of various non-classical reasoning approaches using Monte-Carlo simulations and deep learning. This is because for these games, classical AI approaches based on game theory such as alpha-beta-pruning (and the underlying minimax-search) are at a severe disadvantage due to the required (partial) enumeration of possible game states becoming computationally infeasible. The stochasticity of many games increases the state space even further and additionally requires statistical considerations. Similarly, AI approaches found in the area of knowledge representation and reasoning (KR) that use semantic methods are sidelined in these benchmarks, as they do not have the high reactiveness needed as the large number of possible states with random transitions leads to excessively time-consuming computations.

Moreover, games provide an abstract, controllable and nearly arbitrarily complex environment that can mimic the "real world" as closely as needed for a test while keeping interfering influences at bay. For instance, in computer games and other simulations, experiments can be conducted entirely without measurement noise so that the real effects can be investigated free from falsifications, or a specific, controlled amount of noise can be added deliberately to investigate how well the approaches cope with it. Additionally, games can often be sped up to enable repeated tests, as is needed, e.g., for evolutionary strategies. On top of that, unlike simulations, games have the desirable property of offering an easy approach to integrating humans into the loop by playing against or with AI players. This fact has already successfully been used in various studies on human cognition [7,2]. Games also have motivational and immersive aspects which facilitate both finding survey participants who make an honest effort, as well as measuring their genuine reactions.

In the following, we argue that platform games (or *platformers*), that is, games in the tradition of *Donkey Kong* (Nintendo, 1981), *Impossible Mission* (Epyx, 1984), *Prince of Persia* (Brøderbund, 1989) and the most influential[1] *Super Mario Bros.* (Nintendo, 1985), which require the player to overcome a multitude of obstacles (primarily by jumping between platforms, hence the name) on their way to the goal, have additional merits as proving grounds for AI. Advantages of using platformers as a testbed include, but are not limited to:

Scalability of Challenge Type The type of challenge can be varied easily, e.g. by limiting the information on the environment through a change in the agent's visual range, by changing the "physical" stretch of the level, or by restricting the time allowed for making decisions. As a result, both the long-term planning capabilities and reactiveness of an AI agent can be tested with a platformer.

Scalability of Challenge Difficulty Platform games provide a lot of different parameters (e.g., different types and counts of obstacles) which can be used to scale the difficulty of a level while keeping the core task unchanged.

Existence of Game Patterns It is not uncommon for levels in platformers that certain sets of different obstacle types can be overcome using the same techniques (e.g., jumping over a chasm, a body of water, or deadly spikes covering the same horizontal space), so it is possible for two levels to differ in their concrete obstacles while being identical in terms of the strategies needed to reach the goal.

[1] and, according to the Guinness World Records, best-selling video game of all-time `https://web.archive.org/web/20100224070604/http://gamers.guinnessworldrecords.com/records/nintendo.aspx`

Possibility of Multiple Solutions The layout of obstacles in a level usually allows for more than one way to reach the goal. This gives room for evaluating whether a strategy is successful without limiting its course of action more than necessary, thus providing room for uncommon or "creative" solutions.

Generalisation Test By using different types of levels, it is possible to test whether or not an approach is capable of generalising a valid solution (that is, the solution to one level) to a similar, but not identical task (i.e. another level following the same rules but differing in terms of design).

Following this introduction to the paper and the topic in general, we present the game ZooOperation in the subsequent Section 2 by consecutively introducing the game itself in Section 2.1, existing controllers for the avatars in Section 2.2 and finally the corresponding ZooOperation competition in the scope of the KI 2017 conference in Section 2.3. This is followed by a discussion of how computer games may support research in the various areas of artificial intelligence (Section 3), where we describe further questions and characteristics that are suitable for the analysis of KR approaches based on their empirical performance in the game. We afterwards conclude the article in Section 4 with a summary and our suggestions on future applications of ZooOperation.

2 ZooOperation

The game ZooOperation is a cooperative platform game inspired by the game *Geometry Friends* [12] and was created as a student project at TU Dortmund University [3]. Unlike this and other cooperative games such as *RoboCup* (Simulation League) [2], ZooOperation challenges planning, coordination and collaboration almost exclusively through removing the additional complexity of extensive physics simulations. Additionally, with ZooOperation, we challenge the AIs with avatars that each have different but closely defined abilities that have to be coordinated to reach the goal. This differs from other cooperative challenges as, for instance, the Robo Cup Simulation League[3] or Neuro-Evolving Robotic Operatives[4], where a swarm of avatars with more or less the same set of skills has to reach a common goal.

2.1 Game description

In ZooOperation, up to five agents each take control over one of a fixed set of avatars where every avatar has unique capabilities; Figure 1 gives an overview over the avatars in the game. To finish a level in this game, all present avatars must reach the designated final destination in the level; Figure 3a shows an example of a level in this game. In order to reach this goal, the avatars have to cross the level while circumventing various obstacles (shown in Figure 2) on the way. These may be harmless obstacles that just obstruct the movement of

[2] `http://www.robocup2017.org`

[3] `http://wiki.robocup.org/Soccer_Simulation_League`

[4] `http://nn.cs.utexas.edu/nero/`

Crocodile, capable of climbing harmless game elements.

Tiger, can run quicker and jump further than any other avatar.

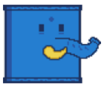

Elephant, can swim and carry other avatars over bodies of water.

Mouse, can change its dimensions, but not its surface area.

Piglet, can serve as a trampoline, giving other avatars a wider jump range.

Fig. 1: Avatars in the game ZooOperation [3].

Ground, walkable ground for topological design of level, sometimes grassy.

Box, movable and walkable obstacle.

Goal, must be reached by all present avatars to finish the level.

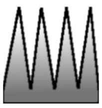

Spikes, dangerous game element, fixed on ground or ceiling. May fall down if an avatar walks underneath.

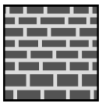

Wall, for topological design of level.

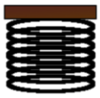

Coil spring, catapults an avatar upwards.

Moving Platform, a walkable platform that moves.

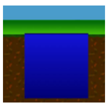

Water, dangerous element for every non-swimming avatar.

Fig. 2: Game elements in the game ZooOperation [3].

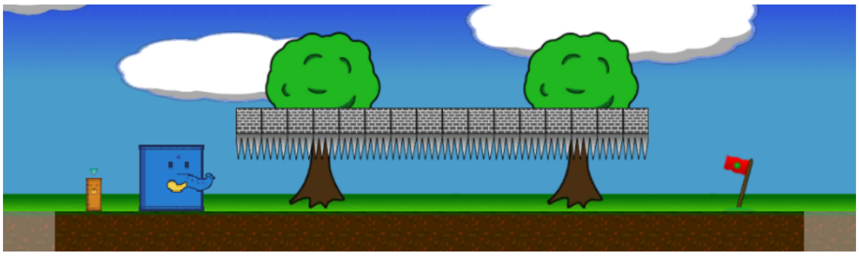

(a) Cooperative level: Tiger needs to clear the way of falling spikes for Elephant to be able to walk to the goal.

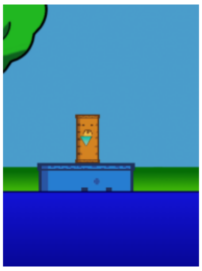

(b) Cooperative task: Elephant carries Tiger over a body of water too wide to jump over.

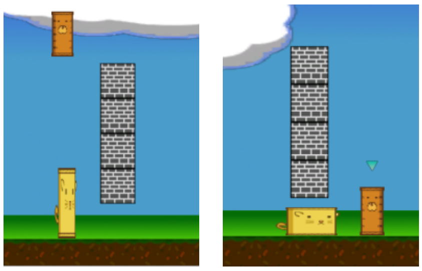

(c) Complex obstacle: Mouse lifts Tiger high enough for it to be able to jump over the wall, then has to change its shape to fit through the gap.

Fig. 3: Different obstacles in ZooOperation to be overcome through cooperation [3].

certain avatars (for instance, being too high to jump over or too narrow to crawl underneath), or deadly obstacles like fixed or falling spikes, deep chasms, or bodies of water. In many cases, it is possible to overcome these obstacles using different strategies for different avatars. For instance, Tiger can jump over a body of water, whereas Elephant swims through it, but the other avatars require a different, cooperative strategy because they can neither jump far enough nor swim. A possible solution to this specific problem is for the avatar to be carried across the water by Elephant (see Figure 3b). Other cooperative obstacles include pathways that have to be cleared by a smaller avatar before a larger one can fit through (Figure 3a) or complex obstacles where special capabilities of different avatars have to be combined to overcome the obstacle (Figure 3c).

To test AI controllers, the game provides a TCP/IP interface that sends the game state to every bound AI controller and allows the AIs to control the avatars and send user-defined messages to be read by all other AIs ("Blackboard"). The controls provided to interact with the game are the same ones a human player might use, that is, the AI can send keystrokes for `up`, `down`, `left`, `right`, `special` (for special skills of the avatar, if applicable). Therefore, it is not necessary to specialise an AI approach for or deeply integrate the approach into the game, but it instead suffices to provide the aforementioned TCP/IP interface that acts as a bilateral translator. This interface needs to (1) translate the game state into a form the AI can understand and (2) translate the designated action of the AI into keystrokes to be sent to the game.

2.2 Automated ZooOperation Controllers

As described above, the game ZooOperation has been developed specifically for testing different AIs. In the project it was developed for, it has already been used to assess and train different controller types and strategies. To illustrate how diverse approaches and controllers for solving levels in ZooOperation may be, we highlight a selection of three strategies already developed and applied to the game; see the project report [3] for a complete overview and detailed description of all strategies developed.

Graph Approach: The graph approach construes a level as a directed graph, where every (physically) coherent traversable area with identical headroom is interpreted as a vertex. An edge is added for every movement in the game that allows the individual avatar to change its position from one of the vertices to another. This relocation can be achieved e.g. by walking, parabolic jumping or falling (to a lower vertex). A vertex in the graph is *reachable* (in a graph-theoretic sense), if and only if it is also reachable in the game, that is, there is a combination of movements that allows the player to manoeuvre the avatar from the starting vertex to the final vertex. Using this approach, a level can be solved by a standard algorithm for finding (shortest) paths in graphs given that the level is solvable without cooperation.

Dynamic Jump: Sometimes, parabolic jumping does not give all possible targets an avatar can reach, as there may be obstacles in the way or the ceiling is low. Using the techniques of dynamic programming, this approach calculates all positions that are reachable from a fixed starting position via a jump or a fall. During this, it populates a table with reachability information for every potential future position of the avatar up to a predefined time limit. It then goes backwards from a destination to find a possible path and the corresponding commands. This has also been used to generate extra edges for the graph approach described above. Figure 4a illustrates the result of one such calculation, that is, the traversals and endpoints an avatar can reach from its actual position using a single dynamic jump.

Motif Search: The prior two approaches use knowledge only in terms of the properties of the characters and the underlying game's physics (how far and high can a character jump, how fast can it run, . . .).

Motif search instead stores obstacles and corresponding solutions in the form of so-called *motifs*. A motif is a tuple of an abstract representation of the obstacle, the sequence of keystrokes that yield a valid solution (also called an *action*

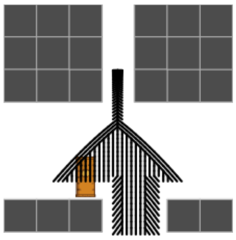

(a) Black lines indicate movements Tiger can make with a dynamic jump from its actual position (endpoints of lines pointing to Tiger's centre of gravity after the jump).

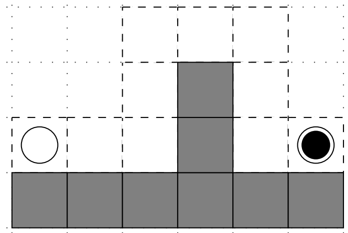

(b) Example motif with start- (circle) and endpoint (double circle), walkable tiles (gray) and traversal (dashed).

Fig. 4: Illustrations of approaches to solving ZooOperation levels [3].

sequence), and the area around the obstacle that is traversed when performing the action sequence. The abstract representation takes the form of a matrix of tiles which encode whether or not a tile is safe to walk through, stand on, etc., and also stores the start and end position of the avatar performing the sequence. These motifs may then be mapped to concrete areas of a level using a distance function on the abstract tiles in the motif and the actual tiles in the level, allowing, for instance, an agent to use the same strategy used to jump over a pit of "dangerous" tiles regardless of whether these tiles are filled with water, spikes or other dangerous elements. Figure 4b is an example of an obstacle's representation by a motif. These motifs can, for instance, be recorded from playthroughs of human or AI players, generated by a machine learning approach, or designed by hand.

2.3 ZooOperation Competition at KI 2017

ZooOperation will be used in a competition at KI 2017 intended to measure the path planning, coordination and puzzle solving capabilities of submitted AI agents. Participants can upload their AI controllers which will then face two types of levels:

– *small* levels with a single obstacle that may or may not require cooperation to overcome, and
– *regular* levels that combine multiple challenges, an example of which is depicted in Figure 5.

The submissions are ranked according to the number of regular levels they finished. In case of a draw, we use the number of small levels finished as a secondary ranking criterion. Any remaining ties will be broken using the time needed to finish the level, measured in terms of the number of game ticks elapsed. Apart from the tertiary ranking criterion, the controllers are not required to make quick

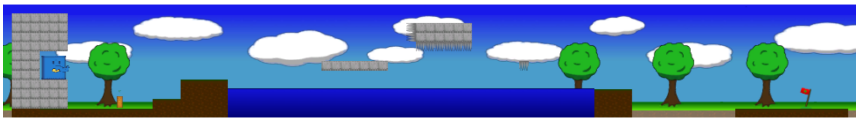

Fig. 5: A *regular* level from the ZooOperation competition: In order to finish the level, Tiger has to carry Elephant so it can master the stairs. Then Elephant must carry Tiger so it can cross the big body of water. During this "boat trip", Tiger must jump on and over the obstacles so it is neither pushed in the water nor killed by the spikes.

real-time decisions, but instead have a maximum of eight minutes per level to solve it. With only a loose time restriction, complete and perfect information and a deterministic game engine, this competition (in contrast to other competitions such as GVGAI[5] and Geometry Friends[6]) stresses the cooperation and problem solving aspects of (cooperative) platformers. Thus, it is possible to include multiple approaches which may differ in their reaction speed, and judge them by their general capability of solving a level in the game rather than the time needed to calculate a solution.

At the same time, the continuous environment provides a challenge different from *grid-based* problems such as, for instance, the Wumpus World [13]. In addition to the selected approaches from [3] used as illustrating examples in Section 2.2, we encourage submissions using diverse strategies and controllers, the possibility of which is ensured by the TCP/IP interface (Section 2.1).

Technically, the competition backend is realised via a web server that provides a user account system backed by a relational database. Here, users may upload AIs to the server where they will be enqueued and tested within a sandbox container using the Docker framework. The test results are then extracted and stored in the database to be displayed as leaderboards. The Docker framework was chosen for this task due to its high scalability and automatic load balancing between containers, allowing for a dynamic reallocation of resources depending on how intensely users strain the system through frequent uploads and tests. On top of that, Docker containers do not expose the underlying server system, impeding malicious action such as a modification of the competition framework.

3 Discussion

The competition at KI 2017 is of course only one of the possible competition setups and scoring schemes that can be based on the ZooOperation software. Since the setup directly characterises the challenges the game provides and steers the focus of the competition, different setups can be employed in order to investigate other aspects of AI. In the following, we list and discuss the different

[5] General Video Game Playing Competition, `http://gvgai.net`

[6] Cooperative physics puzzles, `http://gaips.inesc-id.pt/geometryfriends/`

experimentation scenarios we envision in context of AI research using the ZooOperation software.

Multi Agent Systems: Cooperation generally requires communication among agents. Additionally, to reason whether avatar B is capable of helping avatar A to overcome an obstacle, the controller of A needs a model of B as well. In order to focus on this aspect, the software can restrict information on other agents so that communication between agents is enforced, controlled, or restricted.

Knowledge Representation and Reasoning: The motif approach already uses abstract knowledge to represent partial solutions. Thus, it seems reasonable to examine whether an even more abstract representation, such as a hierarchical knowledge base [1] or a representation using defeasible (conditional) rules to form a conditional knowledge base with respective semantics (see, e.g., [4,15,6,9]) yields satisfactory results, too. In order to specifically analyse the knowledge representation aspects, one could restrict the information passed to the agent accordingly, for example by passing all information through an interface that prohibits or redacts specific information.

AI Generalisability: *Motif search* is only one of the possible approaches that use abstraction in order to generalise from previously learned behaviour. In recent years, the computational intelligence in games community has put considerable effort into finding generalisable AI approaches [11]. The games, however, tend to be extremely different and do not produce observable patterns across different AIs [8]. Using a similarity measure on levels, the need for generalisability could be scalarised in an experiment scenario in order to identify issues where general game AI breaks.

AI Robustness In order to investigate approaches based on uncertain reasoning and belief revision, the information provided to AI players could be limited or otherwise modified in order to create scenarios in which employing the respective techniques becomes inevitable. For example, the physics of the game could be be unknown to the AIs (as is the case in the Angry Birds AI Competition[7] for instance) or be subject to undisclosed changes (e.g. by randomly changing gravity). Another possible scenario is one where the characters in the level are controlled by AIs that are unfamiliar with one another.

Measuring Game Characteristics: Measuring game characteristics such as difficulty and required strategy depth are open issues of interest within the computational intelligence in games community [8,2]. We plan to extend future work in this regard by investigating measures that can identify levels of the same difficulty class based on empirical results of different AI agents playing ZooOperation.

Involving Human Players: As artificial ZooOperation controllers and human players steer the characters in the same way, it is possible to include human players in the experiments. Possible scenarios include the comparison of behaviour patterns of AI and human players as well as identifying

[7] `https://aibirds.org/angry-birds-ai-competition.html`

explanatory components for black-box agents by asking human players that have gained expertise through repeated games. Furthermore, the challenge could be extended to include human-computer-interaction by building mixed AI and human teams that need to communicate and collaborate (cf. [5]).

4 Conclusion

In this paper, we made a case for using computer games to test and compare approaches from artificial intelligence and, specifically, approaches from knowledge representation and reasoning. We presented the game ZooOperation and illustrated how platform games in general, and this game in particular, can serve as a proving ground to investigate a variety of questions. Additionally, we provided a description of the ZooOperation competition at KI 2017 as an example for investigating the path planning, coordination and puzzle solving capabilities of AI agents along with an overview of suitable solutions. We already invited other researchers of AI to use ZooOperation as test bed for their approaches and / or to compete with each other in the ZooOperation competition. We discussed how the presented game, and computer games in general, can be used to investigate and rank different approaches to AI in terms of further properties, be it communication, knowledge representation and reasoning, or robustness and generalisability of AIs. This underpins our general claim that computer games provide a viable proving ground for judging and comparing AI-approaches in addition to their formal properties.

This, of course, was only the first step, and future work can be broken into two major parts: First, applying established AI approaches to the task of solving platform games like ZooOperation as a simulation of tasks in complex environments, and comparing them with each other based on their performance in these simulations. Second, as indicated in the discussion, this general framework can help in improving interactions between human users and automated systems through researching, for example, general notions of difficulty of (scalably) complex tasks (as seen in platform games), or the performance and results of mixing human and artificial players (involving humans as sparring partners or as members of mixed teams).

Acknowledgements

We thank the anonymous reviewers for their valuable hints that helped us improve the paper. This work was supported by DFG-Grant KI1413/5-1 as part of the priority program “New Frameworks of Rationality” (SPP 1516) and DFG research unit FOR 1513 on “Hybrid Reasoning for Intelligent Systems” to Gabriele Kern-Isberner. Christian Eichhorn is supported by Grant KI1413/5-1, Richard Niland is supported by FOR 1513.

References

1. Apeldoorn, D., Kern-Isberner, G.: When Should Learning Agents Switch to Explicit Knowledge? In: Benzmüller, C., Sutcliffe, G., Rojas, R. (eds.) GCAI 2016. 2nd Global Conference on Artificial Intelligence. EPiC Series in Computing, vol. 41, pp. 174–186. EasyChair Publications (2016)
2. Apeldoorn, D., Volz, V.: Measuring Strategic Depth in Games Using Hierarchical Knowledge Bases. In: Computational Intelligence in Games Conference (CIG). pp. 9–16. IEEE Press, New York, NY, USA (2017)
3. Buttkus, M., Fecke, M., Gärtner, D., Junge, J., Majchrzak, K., Nehrke, A., Schendekehl, T., Schlüter, C., Shao, X.: ZooOperation: Spielende kooperierende Agenten. Projektgruppenbericht der PG 596, Faculty of Computer Science, TU Dortmund University, Dortmund, DE (2016), (in German)
4. Dubois, D., Prade, H.: Possibility Theory and Its Applications: Where Do We Stand? In: Kacprzyk, J., Pedrycz, W. (eds.) Springer Handbook of Computational Intelligence, pp. 31–60. Springer Berlin Heidelberg, Berlin, DE (2015)
5. Eger, M., Martens, C., Córdoba, M.A.: An Intentional AI for Hanabi. In: CIG'2017 - IEEE Conference on Computational Intelligence and Games. pp. 68–75. IEEE CIG, IEEE Computer Society, New York, NY, USA (2017)
6. Eichhorn, C., Kern-Isberner, G.: Qualitative and Semi-Quantitative Inductive Reasoning with Conditionals. KI – Künstliche Intelligenz 29(3), 279–289 (2015)
7. Holmgård, C., Togelius, J., Henriksen, L.: Computational Intelligence and Cognitive Performance Assessment Games. In: Computational Intelligence in Games Conference (CIG). pp. 485–492. IEEE Press, Santorini, Greece (2016)
8. Horn, H., Volz, V., Perez-Liebana, D., Preuss, M.: MCTS/EA hybrid GVGAI players and game difficulty estimation. In: Computational Intelligence in Games Conference (CIG). pp. 278–285. IEEE Press, Santorini, Greece (2016)
9. Kern-Isberner, G.: Conditionals in Nonmonotonic Reasoning and Belief Revision – Considering Conditionals as Agents. No. 2087 in Lecture Notes in Computer Science, Springer Science+Business Media, Berlin, DE (2001)
10. Ontanon, S., Synnaeve, G., Uriarte, A., Richoux, F., Churchill, D., Preuss, M.: A Survey of Real-Time Strategy Game AI Research and Competition in StarCraft. IEEE Transactions on Computational Intelligence and AI in Games 5(4), 293–311 (2013)
11. Perez, D., Samothrakis, S., Togelius, J., Schaul, T., Lucas, S.M., Couëtoux, A., Lee, J., Lim, C.U., Thompson, T.: The 2014 General Video Game Playing Competition. IEEE Transactions on Computational Intelligence and AI in Games 8(3), 229–243 (2015)
12. Prada, R., Lopes, P., Catarino, J., Quitério, J., Melo, F.S.: The Geometry Friends Game AI Competition. In: CIG'2015 - IEEE Conference on Computational Intelligence and Games. pp. 431–438. IEEE CIG, IEEE Computer Society, Tainan, Taiwan (2015)
13. Russell, S.J., Norvig, P.: Artificial Intelligence: A Modern Approach. Prentice Hall series in artificial intelligence, Prentice Hall (2010)
14. Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., Hassabis, D.: Mastering the game of Go with deep neural networks and tree search. Nature 529(7587), 484–489 (2016)

15. Spohn, W.: The Laws of Belief: Ranking Theory and Its Philosophical Applications. Oxford University Press, Oxford, UK (2012)
16. Togelius, J.: AI Researchers, Video Games Are Your Friends! In: Merelo, J.J., Rosa, A., Cadenas, J.M., Correia, A.D., Madani, K., Ruano, A., Filipe, J. (eds.) 7th International Joint Conference on Computational Intelligence (IJCCI 2015). pp. 3–18. Springer International Publishing, Cham, CH, Lisbon, Portugal (2016)
17. Turing, A., Bates, M.A., Bowden, B., Strachey, C.: Digital computers applied to games. In: Bowden, B. (ed.) Faster than Thought, chap. 25, pp. 286–310. Sir Isaac Pitman & Sons, Ltd., London, UK (1953)