General Video Game Playing

John Levine[1], Clare Bates Congdon[2], Marc Ebner[3], Graham Kendall[4], Simon M. Lucas[4], Risto Miikkulainen[6], Tom Schaul[7], and Tommy Thompson[8]

1 **Department of Computer and Information Sciences, University of Strathclyde**
john.levine@strath.ac.uk
2 **Department of Computer Science, University of Southern Maine**
congdon@usm.maine.edu
3 **Universitat Griefswald, Institut fur Mathematik und Informatik**
marc.ebner@uni-greifswald.de
4 **School Of Computer Science, University of Nottingham**
gxk@cs.nott.ac.uk
5 **School of Computer Science and Electrical Engineering, University of Essex**
sml@essex.ac.uk
6 **Department of Computer Science, University of Texas at Austin**
risto@cs.utexas.edu
7 **Courant Institute of Mathematical Sciences, New York University**
schaul@cims.nyu.edu
8 **School of Computing and Mathematics, University of Derby**
t.thompson@derby.ac.uk

Abstract

One of the grand challenges of AI is to create general intelligence: an agent that can excel at many tasks, not just one. In the area of games, this has given rise to the challenge of General Game Playing (GGP). In GGP, the game (typically a turn-taking board game) is defined declaratively in terms of the logic of the game (what happens when a move is made, how the scoring system works, how the winner is declared, and so on). The AI player then has to work out how to play the game and how to win. In this work, we seek to extend the idea of General Game Playing into the realm of video games, thus forming the area of General Video Game Playing (GVGP). In GVGP, computational agents will be asked to play video games that they have not seen before. At the minimum, the agent will be given the current state of the world and told what actions are applicable. Every game tick the agent will have to decide on its action, and the state will be updated, taking into account the actions of the other agents in the game and the game physics. We envisage running a competition based on GVGP playing, using arcade-style (e.g. similar to Atari 2600) games as our starting point. These games are rich enough to be a formidable challenge to a GVGP agent, without introducing unnecessary complexity. The competition that we envisage could have a number of tracks, based on the form of the state (frame buffer or object model) and whether or not a forward model of action execution is available. We propose that the existing Physical Travelling Salesman (PTSP) software could be extended for our purposes and that a variety of GVGP games could be created in this framework by AI and Games students and other developers. Beyond this, we envisage the development of a Video Game Description Language (VGDL) as a way of concisely specifying video games. For the competition, we see this as being an interesting challenge in terms of deliberative search, machine learning and transfer of existing knowledge into new domains.

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

Keywords and phrases Video games, artificial intelligence, artificial general intelligence

Digital Object Identifier 10.4230/DFU.Vol6.12191.77

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; pp. 77–83
DAGSTUHL FOLLOW-UPS Dagstuhl Publishing
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Germany

1 Motivation

The field of Artificial Intelligence is primarily concerned with creating agents that can make good decisions [15]. In the context of games, this definition is particularly relevant, as the player of a game will typically have to make a large number of accurate decisions in order to achieve a favourable outcome (a win over an opponent, or a high score). Games have thus always formed a fertile source of benchmarks and challenges for the AI researcher since the earliest days of the field [16].

Great progress has been made with specific games, such as chess, backgammon and poker (to name but three of the many games that AI researchers have effectively mastered). However, each AI program is written to tackle a specific game and the techniques used for one game do not necessarily transfer to other games. Humans can thus still claim to be, in some sense, more intelligent than these programs because they can play multiple games and learn how to play and excel at new ones: their intelligence is more *general*.

In AI, we are always striving for theories concerning intelligence in general rather than intelligence in the specific. If we come up with a theory of good decision making and use a turn-taking game as *an example* of how that theory works, then we have made a better contribution than someone who writes a bespoke computer program just to play that game. In the best case, our theory of intelligence should be able to be applied to any situation in the real world where decisions are taken.

These observations, combined with the success of computer programs at turn-taking games such as chess and Othello, have given rise to the challenge of General Game Playing (GGP) [5], which has now become an annual competition at the AAAI conference [22]. This competition focuses on turn-taking board games. The logic of the game is specified in a language called GDL (Game Description Language) [21] which in turn is inspired by the single agent Planning Domain Definition Language (PDDL) [24] used in the AI planning community.

Turn-taking board games are only one genre of games at which humans excel. A turn-taking board game is usually quite slow-paced, and so deliberative reasoning about the consequences of actions and the likely responses of opponents are possible using seconds or even minutes of CPU time. Humans also excel at other games where decisions have to be made much more quickly and where the environments are noisy, dynamic and not always predictable. Video games are a very good example of this genre and form an interesting challenge for AI researchers.

Classic arcade video games (such as Pac-Man, Asteroids and Space Invaders) are very good examples of the kind of activities we are seeking to cover in this work. In his book *Half-Real* [11], Juul offers a six-part definition of what a game is: (1) games are rule-based; (2) they have variable outcomes; (3) the different outcomes have different values; (4) the player invests effort to influence the outcome; (5) the player is attached to the outcome; and (6) the game has negotiable consequences. Looking specifically at points (3), (4) and (5), our core assumptions in this work are that the value of the outcome achieved by the player of a video game (point 3) is a direct reflection of the amount of intelligence (intellectual effort) brought to bear on the task (point 4); and that this can be further used as an indicator of the player's general level of intelligence (point 5).

In this work, we seek to extend the idea of General Game Playing into the realm of video games, thus forming the area of General Video Game Playing (GVGP). In GVGP, computational agents will be asked to play video games that they have not played before. The agent will have to find out how the game reacts to its actions, how its actions affect its rewards and therefore how to win the game (or maximise its score). At the minimum, the agent will be given the current state of the world and told what actions are applicable. Every game tick the agent will have to decide on what action to take. The state will then be updated, taking into account the actions of the other players in the game and the game physics. We envisage that the rules and the physics of the game can be encoded using a Video Game Description Language (VGDL). We have analysed a selection of classic arcade games with a view to representing them in VGDL and the results of this work are presented in another paper in this volume.

The aims of this work are:

- to define what General Video Game Playing is and our motivation for this investigation;
- to review some of the previous work that has influenced and inspired this work;
- to discuss what video games we would initially like to tackle in this work;
- to speculate on the computational methods that might be appropriate for GVGP;
- to present a proposal for a GVGP framework, based on existing software for the Physical Travelling Salesman Problem (PTSP);
- to outline our proposal for a GVGP competition.

The remainder of the chapter is structured as follows. In the next section, we look at the background behind GVGP, in particular the GGP competition and the work of the Atari 2600 group at Alberta. We then examine what sort of games would like to tackle in GVGP and the mechanisms that agents might use in order to play these games competently. We then examine how we would build a GVGP framework, based on the Physical Travelling Salesman (PTSP) work at the University of Essex. Finally, the paper ends with a proposal for a GVGP competition, with tracks based on the form of the state (frame buffer or object model) and whether or not a forward model of action execution is available.

2 Background

General video game playing is an excellent tool to test artificial intelligence algorithms. If a human player is sitting in front of a video game and is playing the game, then they need to make observations about the state of the game:

- Where is the avatar of the player (i.e. which screen object is 'me')?
- Where are the opponents or enemies?
- What are the current options (direction of next move, firing of a weapon, jump, etc.)?
- Which option will most likely bring the own avatar towards a particular goal?

One of the main advantages of general video game playing is that it is of intermediate complexity. It is more complex than simple board games, e.g. movement options may be continuous and the game may also employ simulated physics. However, general video game playing is not as complex as developing human-like robots.

In developing human-like robots, researchers have resorted to highly constrained settings. Sample tasks include: vacuuming the floor [10], cleaning windows [4], moving objects or persons from location A to location B [8], or giving tours in a museum [19]. Even though these problems have all been beautifully addressed by these researchers the settings are constrained, i.e. algorithms are usually only transferable to similar domains. In some cutting edge areas of robotics (search for the holy grail of artificial intelligence) the goals are not at all clearly defined. Should a robot be able to play chess, learn to play the piano, wash the dishes or go shopping for groceries? Should a robot first learn to play chess and then how to go shopping or vice versa? That aside, robot actuators and input sensors may not allow handling of completely different problem classes such as driving on the road and playing chess.

In contrast, in general video game playing, the goal of the player is clearly defined. Moreover, there are several different goals (one for each game). The goals are not taken from an artificial laboratory setting but correspond to actual goals that have been addressed by numerous game playing enthusiasts around the world. Goals include:

- Navigating a yellow object through a maze while collecting all white dots and at the same time avoiding colored ghosts (Pac Man).
- Landing a space ship on the moon (Lunar Lander).
- Navigating a space ship through a field of asteroids (Asteroids).
- Driving a racing car (Pole Position).
- Climbing to the top of a platform environment while avoiding falling objects (Donkey Kong).

While each game has its unique story and goal which should be achieved by the player, the method of how the player interacts with the game is the same across all these games. All possible moves of the player are defined through the game controller (e.g. the Atari 2600 had a joystick with a single button resulting in 18 possible moves for each time step). The number of possible movements for each time step are sufficiently small and discrete. However, given the continuous nature of the game, i.e. a new decision has to be made for each time step, the game is sufficiently complex.

We believe that addressing the diverse problems of general video game playing will allow us to gain new insights in working towards the holy grail of artificial intelligence, i.e. development of human-like intelligence.

2.1 General Game Playing

The General Game Playing competition at AAAI has now been running since 2005 and focuses on turn-taking board games. The programs which enter the competition are not given the details of the games in advance; instead, the rules of the game to be played are specified using GDL, the Game Description Language [21]. Given the unseen nature of the games, there are two common approaches used to construct players, namely deliberative reasoning (e.g. minimax search) or policy learning (e.g. reinforcement learnng). One of the most successful players to date is Cadia Player [2], which is a deliberative approach based on Monte Carlo Tree Search [12, 3].

2.2 The Atari 2600 Games Group

The work described in our paper has a strong connection to the work of the Atari 2600 games group at Alberta, and much of our discussion was influenced by their work [1]. They have built what they call The Arcade Learning Environment (ALE), based on the games for the classic Atari 2600 games console. A variety of different genres of game are available for this console: shooters, adventure, platform, board games, and many others. An emulator for 2600 games is used in the ALE, and the player of these games has to interact with this emulator to find out what action to take next. Two approaches to creating players are investigated, based on reinforcement learning [18] and Monte Carlo Tree Search [12]. Other recent work in this area by Hausknecht *et al.* has investigated the use of evolutionary neural networks for creating general video games players [6].

The aim of the work cited in this section is similar to ours: to create general artificial intelligence agents for playing video games. Our proposal complements this work in two ways:

firstly, by proposing an alternative and potentially more general framework for creating the games, including a Video Game Description Language; and secondly, by proposing a General Video Game Playing competition for evaluating different approaches to this problem.

2.3 Artificial General Intelligence

The field of artificial general intelligence (AGI) addresses the problem of artificial intelligence in its most broad sense. The intelligent agents that it aims to build should be able to (learn to) act well in a very broad range of environments, and under a broad range of constraints. This contrasts with narrower AI research that focuses on a single application domain. A large portion of the recent work in AGI is based on Hutter's theoretical framework [9]. Recently, games were proposed as the specific ingredient for measuring this form of general intelligence [17]; namely, an (infinite) set of unseen games were proposed to play the role of the "broad range of environments".

3 Games for General Video Game Playing

What kind of games are suitable for General Video Game Playing research? Since we want to start our investigations with a feasible challenge, we think that classic 2-dimensional arcade games offer the right balance. If this is successful, we then would envisage moving into 3-dimensional games, which offer more complexity (not least in terms of the effort required to create the game environment). In the companion paper in this volume, we present Video Game Description Language definitions for three diverse arcade games: Space Invaders, Lunar Lander and Frogger. Other suitable candidates might include simple first person shooting games such as Maze War [23], albeit represented in a 2-dimensional form.

All these games are characterised by relatively simple but fast-paced gameplay, with the need for the player to make accurate decisions and adopt high performing strategies in order to win. The emulators for such games can be created more easily than those complex 3-dimensional games but will still allow for sufficient experimentation for the results to be transferrable to more complex game environments.

4 A Competition for General Video Game Playing

One of the main challenges of General Video Game Playing is to create a software framework that allows for games to be designed and representated and different game-playing agents tested via some form of long-running competition. Competitions have been very useful in promoting research in creating AI players of games: examples include the long-running Ms Pac-Man competition [13], the 2K Botprize [7] and the Physical Travelling Salesman Problem [14].

In order to support research in General Video Game Playing we propose using software created for the Physical Travelling Salesman Problem as the basis of a General Video Game Playing framework. In order to do this, we would need to extend the framework to add more agents and more game objects and create a general state model for the system. The framework would be adapted to read 2-dimensional games written in VGDL.

We propose to hold a GVGP competition, in a number of tracks. These tracks will be based on variations in what data is available to the player:

- What data will be provided to the player to characterise the state? A frame buffer or an object model?
- Will the player have access to the entire state or just a first-person perspective?

- Will the player be able to use the simulator as a forward model to be able to do deliberative reasoning?

Since these three features are mutually exclusive, the competition could, in theory at least, run in 8 independent tracks (e.g. frame buffer model, entire state, no forward model). In practice, only some of these variations will be interesting: in particular, the decision to use a the entire state or a first person perspective will be dictated by the genre of game we are playing. However, the use of frame buffer vs object model and the use of a forward model vs no forward model are both interesting dimensions, thus resulting in at least 4 variations for each game.

5 Potential Applications

While creating general video game players is an interesting academic activity in itself, we think that this area of research could find applications in real video games. Our eventual aim would be to enable the creation of high performance AI players without any extra code being written: the game environment would be designed and then handed over to the general AI agent, which would then work out how to play the game competently. In the shorter term, we think that general video game playing techniques could be used both to test game environments, including those created automatically using procedural content generation [20] and to find potential loopholes in the gameplay that a human player could exploit.

References

1 M.G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The arcade learning environment: An evaluation platform for general agents. *Arxiv preprint arXiv:1207.4708*, 2012.

2 Yngvi Björnsson and Hilmar Finnsson. Cadiaplayer: A simulation-based general game player. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(1):4–15, 2009.

3 Cameron Browne, Edward J. Powley, Daniel Whitehouse, Simon M. Lucas, Peter I. Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Trans. Comput. Intellig. and AI in Games*, 4(1):1–43, 2012.

4 M. Farsi, K. Ratcliff, P. J. Sohnson, C. R. Allen, K. Z. Karam, and R. Pawson. Robot control system for window cleaning. In D. A. Chamberlain, editor, *Proceedings of the 11th International Symposium on Automation and Robotics in Construction XI*, pages 617–623. Elsevier Science, 1994.

5 Michael Genesereth and Nathaniel Love. General game playing: Overview of the AAAI competition. *AI Magazine*, 26:62–72, 2005.

6 Matthew Hausknecht, Piyush Khandelwal, Risto Miikkulainen, and Peter Stone. Hyperneat-ggp: A hyperneat-based atari general game player, 2012.

7 Philip Hingston. The 2k botprize. `http://botprize.org/`, 2012. [Online; accessed 27-August-2012].

8 Yuji Hosoda, Masashi Koga, and Kenjiro Yamamoto. Autonomous moving technology for future urban transport. *Hitachi Review*, 60(2):100–105, April 2011.

9 Marcus Hutter. *Universal artificial intelligence: Sequential decisions based on algorithmic probability*. Springer Verlag, 2005.

10 iRobot Corporation. *Roomba. Vacuum Cleaning Robot*. 8 crosby Drive, Bedfard, MA 01730, 2011.

11 Jesper Juul. *Half-Real: Video Games Between Real Rules and Fictional Worlds*. MIT Press, 2005.

12 Levente Kocsis and Csaba Szepesvári. Bandit based monte-carlo planning. In Johannes Fürnkranz, Tobias Scheffer, and Myra Spiliopoulou, editors, *ECML*, volume 4212 of *Lecture Notes in Computer Science*, pages 282–293. Springer, 2006.

13 Simon M. Lucas. Ms pac-man competition. `http://cswww.essex.ac.uk/staff/sml/pacman/PacManContest.html`, 2012. [Online; accessed 27-August-2012].

14 Diego Perez, David Robles, and Philipp Rohlfshagen. The physical travelling salesman problem competition. `http://www.ptsp-game.net/`, 2012. [Online; accessed 27-August-2012].

15 Stuart J. Russell and Peter Norvig. *Artificial Intelligence – A Modern Approach (3rd edition)*. Pearson Education, 2010.

16 Jonathan Schaeffer and Jaap van den Herik, editors. *Chips Challenging Champions: Games, Computers, and Artificial Intelligence*. Elsevier Science Inc., New York, NY, USA, 2002.

17 T. Schaul, J. Togelius, and J. Schmidhuber. Measuring intelligence through games. *Arxiv preprint arXiv:1109.1314*, 2011.

18 Richard Sutton and Andrew Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

19 Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B. Cremers, Frank Delaert, Dieter Fox, Dirk Hähnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulte, and Dirk Schulz. Minerva: A second-generation musuem tour-guide robot. In *Proceedings of IEEE International Conference on Robotics and Automation*, pages 1999–2005. IEEE, 1999.

20 Julian Togelius, Georgios N. Yannakakis, Kenneth O. Stanley, and Cameron Browne. Search-based procedural content generation. In Cecilia Di Chio, Stefano Cagnoni, Carlos Cotta, Marc Ebner, Anikó Ekárt, Anna Esparcia-Alcázar, Chi Keong Goh, Juan J. Merelo Guervós, Ferrante Neri, Mike Preuss, Julian Togelius, and Georgios N. Yannakakis, editors, *EvoApplications (1)*, volume 6024 of *Lecture Notes in Computer Science*, pages 141–150. Springer, 2010.

21 Wikipedia. Game description language — wikipedia, the free encyclopedia. `http://en.wikipedia.org/wiki/Game_Description_Language`, 2012. [Online; accessed 27-August-2012].

22 Wikipedia. General game playing — wikipedia, the free encyclopedia. `http://en.wikipedia.org/wiki/General_Game_Playing`, 2012. [Online; accessed 27-August-2012].

23 Wikipedia. Maze war — wikipedia, the free encyclopedia. `http://en.wikipedia.org/wiki/Maze_War`, 2012. [Online; accessed 27-August-2012].

24 Wikipedia. Planning domain definition language — wikipedia, the free encyclopedia. `http://en.wikipedia.org/wiki/Planning_Domain_Definition_Language`, 2012. [Online; accessed 27-August-2012].