

RAPID ADAPTATION OF VIDEO GAME AI

Sander Bakkes, Pieter Spronck and Jaap van den Herik
Tilburg University
Tilburg Centre for Creative Computing (TiCC)
P.O. Box 90153, NL-5000 LE Tilburg, The Netherlands
e-mail: {s.bakkes, p.spronck, h.j.vdnherik}@uvt.nl

ABSTRACT

Current approaches to adaptive game AI require either a high quality of utilised domain knowledge, or a large number of adaptation trials. These requirements hamper the goal of rapidly adapting game AI to changing circumstances. In an alternative, novel approach, domain knowledge is gathered automatically by the game AI, and is immediately (i.e., without trials and without resource-intensive learning) utilised to evoke effective behaviour. In this paper we discuss this approach, called ‘rapidly adaptive game AI’. We perform experiments that apply the approach in an actual video game. From our results we may conclude that rapidly adaptive game AI provides a strong basis for effectively adapting game AI in actual video games.

INTRODUCTION

Over the last decades, modern video games have become increasingly realistic with regard to visual and auditory presentation. However, game AI has not reached a high degree of realism yet. Game AI is typically based on non-adaptive techniques [22]. A major disadvantage of non-adaptive game AI is that once a weakness is discovered, nothing stops the human player from exploiting the discovery. The disadvantage can be resolved by endowing game AI with adaptive behaviour, i.e., the ability to learn from mistakes. Adaptive game AI can be established by using machine-learning techniques, such as artificial neural networks or evolutionary algorithms. In practice, adaptive game AI in video games is seldom implemented because machine-learning techniques typically require numerous trials to learn effective behaviour. To allow rapid adaptation in games, we describe a means of rapid adaptation that is inspired by the human capability to solve problems by generalising over previous observations in a restricted problem domain.

The outline of this paper is as follows. First, we discuss the aspect of entertainment in relation to game AI. Then, we discuss our approach to establish rapidly adaptive game AI. Subsequently, we describe an implementation of rapidly adaptive game AI. Next, we describe the experiments that test rapidly adaptive game AI in

an actual video game, followed by a discussion of the experimental results. Finally, we provide conclusions and describe future work.

ENTERTAINMENT AND GAME AI

The purpose of a typical video game is to provide entertainment [22, 15]. Of course, the criteria of what makes a game entertaining may depend on who is playing the game. Literature suggests the concept of immersion as a general measure of entertainment [14, 21]. Immersion concerns evoking an immersed feeling with a video game, thereby retaining a player’s interest in the game. As such, an entertaining game should at the very least not repel the feeling of immersion from the player [12]. Aesthetical elements of a video game, such as graphics, narrative, and rewards, are instrumental in establishing an immersive game-environment. Once established, the game environment needs to uphold some form of *consistency* for the player to remain immersed within it [12].

The task for game AI is to control game characters in such a way that behaviour exhibited by the characters is consistent within the game environment. In a realistic game environment, realistic character behaviour is expected. As a result, game AI that is solely focused on exhibiting the most *challenging* behaviour is not necessarily regarded as realistic. For instance, in a typical first-person shooter (FPS) game it is not realistic if characters controlled by game AI aim with an accuracy of one hundred per cent. Game AI for shooter games, in practice, is designed to make intentional mistakes, such as warning the player of an opponent character’s whereabouts by intentionally missing the first shot [13].

Consistency of computer-controlled characters within a game environment is often established with tricks and cheats. For instance, in the game HALF-LIFE, tricks were used to establish the illusion of collaborative teamwork [12], causing human players to assume intelligence where none existed [13]. While it is true that tricks and cheats may be required to uphold consistency of the game environment, they often are implemented only to compensate for the lack of sophistication in game AI [6]. In practice, game AI in most complex games still is not consistent within the game environment, and exhibits what has been called ‘artificial stupidity’ [13]

rather than artificial intelligence. To increase game consistency, and thus the entertainment value of a video game, in our research we foremost strive to create an optimally playing game AI, as suggested by Buro and Furtak [6]. In complex video games, such as real-time strategy (RTS) games, near-optimal game AI is seen as the only way to obtain consistency of the game environment [12]. Once near-optimal game AI is established, difficulty-scaling techniques can be applied to downgrade the playing-strength of game AI, to ensure that a suitable challenge is created for the player [19].

APPROACH

For game AI to be consistent with the game environment in which it is situated, it needs the ability to adapt adequately to changing circumstances. Game AI with this ability is called ‘adaptive game AI’. Typically, adaptive game AI is implemented for performing adaptation of the game AI in an online and computer-controlled fashion. Improved behaviour is established by continuously making (small) adaptations to the game AI. To adapt to circumstances in the current game, the adaptation process typically is based only on observations of *current* gameplay. This approach to adaptive game AI may be used to improve significantly the quality of game AI by endowing it with the capability of adapting its behaviour while the game is in progress. For instance, the approach has been successfully applied to simple video games [7, 11], and to complex video games [19]. However, this approach to adaptive game AI requires either (1) a high quality of the utilised domain knowledge, or (2) a large number of adaptation trials. These two requirements hamper the goal of achieving *rapidly* adaptive game AI.

To achieve rapidly adaptive game AI, we propose an alternative, novel approach to adaptive game AI that comes without the hampering requirements of typical adaptive game AI. The approach is coined ‘rapidly adaptive game AI’. We define rapidly adaptive game AI as an approach to game AI where domain knowledge is gathered automatically by the game AI, and is immediately (i.e., without trials and without resource-intensive learning) utilised to evoke effective behaviour. The approach, illustrated in Figure 1, implements a direct feedback loop for control of characters operating in the game environment. The behaviour of a game character is determined by the game AI. Each game character feeds the game AI with data on its current situation, and with the observed results of its actions. The game AI adapts by processing the observed results, and generates actions in response to the character’s current situation. An adaptation mechanism is incorporated to determine how to best adapt the game AI. For instance, reinforcement learning may be applied to assign rewards and penalties to certain behaviour exhibited by the game AI.

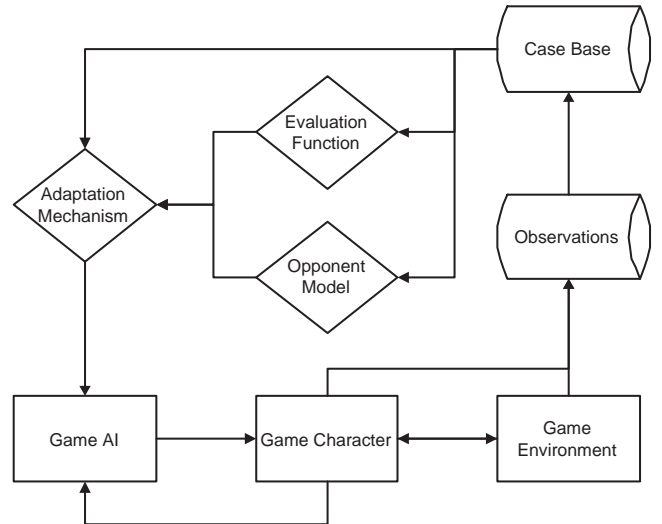


Figure 1: Rapidly adaptive game AI (see text for details).

For rapid adaption, the feedback loop is extended by (1) explicitly processing observations from the game AI, and (2) allowing the use of game-environment attributes which are not directly observed by the game character (e.g., observations of team-mates). Inspired by the case-based reasoning paradigm, the approach collects character observations and game environment observations, and extracts from those a case base. The case base contains all observations relevant for the adaptive game AI, without redundancies, time-stamped, and structured in a standard format for rapid access. To rapidly adapt to circumstances in the current game, the adaptation process is based on domain knowledge drawn from observations of a *multitude* of games. The domain knowledge gathered in a case base is typically used to extract models of game behaviour, but can also directly be utilised to adapt the AI to game circumstances. In our proposal of rapidly adaptive game AI, the case base is used to extract an evaluation function and opponent models. Subsequently, the evaluation function and opponent models are incorporated in an adaptation mechanism that directly utilises the gathered cases.

The approach to rapidly adaptive AI is inspired by the human capability to reason reliably on a preferred course of action with only a few observations on the problem domain. Following from the complexity of modern video games, game observations should, for effective and rapid use, (1) be represented in such a way that stored cases can be reused for previously unconsidered situations, and (2) be compactly stored in terms of the amount of retrievable cases [1]. As far as we know, rapidly adaptive game AI has not yet been implemented in an actual video game.



Figure 2: Screenshot of the SPRING game environment. In the screenshot, airplane units are flying over the terrain.

IMPLEMENTATION

This section discusses our implementation of rapidly adaptive game AI. We subsequently discuss (A) the game environment in which we implement rapidly adaptive game AI, (B) the established evaluation function, and (C) an adaptation mechanism inspired by the case-based reasoning paradigm. Previously established opponent models [17] will be incorporated in future research.

Game Environment

The game environment in which we implement rapidly adaptive game AI, is the actual video game SPRING [10]. SPRING, illustrated in Figure 2, is a typical and open-source RTS game, in which a player needs to gather resources for the construction of units and buildings. The aim of the game is to use the constructed units and buildings to defeat an enemy army in a real-time battle. A SPRING game is won by the player who first destroys the opponent’s ‘Commander’ unit.

Modern RTS games typically progress through several distinct phases as players perform research and create new buildings that provide them with new capabilities. The phase of a game can be straightforwardly derived from the observed traversal through the game’s tech tree. Traversing the tech tree is (almost) always advantageous, yet there is a cost for doing so in time and game resources. In SPRING, three levels of technology are available. At the start of the game, a player can only construct Level 1 structures and Level 1 units. Later in the game, after the player has performed the required research, advanced structures and units of Level 2 and Level 3 become available.

Evaluation Function

To exhibit behaviour consistent within the game environment presented by modern video games, game AI needs the ability to assess accurately the current situation. This requires an appropriate evaluation function. The high complexity of modern video games makes the task to generate such an evaluation function for game AI a difficult one.

In previous research we discussed an approach to generate automatically an evaluation function for game AI in RTS games [4]. The approach incorporated TD-learning [20] to learn unit-type weights for the evaluation function, to reflect the actual playing strength of each unit type. Our evaluation function for the game’s state is denoted by

$$v(p) = w_p v_1 + (1 - w_p) v_2 \quad (1)$$

where $w_p \in [0 \dots 1]$ is a free parameter to determine the weight of each term v_n of the evaluation function, and $p \in \mathbb{N}$ is a parameter that represents the current *phase of the game*. Our evaluation function incorporates two evaluative terms, the term v_1 that represents the material strength and the term v_2 that represents the Commander safety.

Previous research performed in the SPRING environment has shown that the accuracy of situation assessments are closely related to the phase of the game in which they are made [5]. To distinguish phases of the SPRING game, we map tech levels to game phases and distinguish between when tech levels are “new,” and when they are “mature,” as indicated by the presence of units with a long construction time. This leads us to define the following five game phases.

- **Phase 1:** Level 1 structures observed.
- **Phase 2:** Level 1 units observed that have a build time $\geq 2,500$.
- **Phase 3:** Level 2 structures observed.
- **Phase 4:** Level 2 units observed that have a build time $\geq 15,000$.
- **Phase 5:** Level 3 units or Level 3 structures observed.

Results of experiments to test the established evaluation function showed that just before the game’s end, the established evaluation function is able to predict correctly the outcome of the game with an accuracy that approaches one hundred per cent. Note that this is not a trivial result, for two reasons. First, the evaluation function is tuned to make predictions that are good during a large part of the game, not only at the end, and hence it will trade prediction accuracy at the end of the game for higher prediction accuracy earlier in the game. Second,

if the goal of the game was to destroy all the opponent's units, a correct prediction would be easy to make at the end. However, the goal is to destroy the opponent's Commander, and we found that it sometimes happens that a player who is behind in material strength can win, often because the opponent's Commander makes a high-risk move, such as attacking strong enemy units on its own. An evaluation function that is based on comparison of material strength and Commander safety cannot take such moves into account other than allowing for their general statistical likelihood.

In addition, experimental results showed that the evaluation function predicts ultimate wins and losses accurately before half of the game is played. From these results, we concluded that the established evaluation function effectively predicts the outcome of a SPRING game and that the proposed approach is suitable for generating evaluation functions for highly complex video games, such as RTS games. Therefore, we incorporate the established evaluation function in the implementation of our rapidly adaptive game AI.

Adaptation Mechanism

In our approach, domain knowledge collected in a case base is utilised for adapting game AI. To generalise over observations with the problem domain, the adaptation mechanism incorporates an offline means to index collected games, and performs an offline clustering of observations. To ensure that game AI is effective from the onset of a game, it is initialised with a previously observed, successful game strategy. For online action selection, a similarity matching is performed that considers six experimentally determined features. The adaptation mechanism is algorithmically described below, and is subsequently discussed in detail.

```
//Offline processing
A1. Game indexing; calculate indexes for all
    stored games.
A2. Clustering of observations; group together
    similar observations.

//Initialisation of game AI
B1. Establish the (most likely) strategy of
    the opponent player.
B2. Determine in which parameter-band values
    this opponent strategy can be abstracted.
B3. Initialise game AI with an effective
    strategy observed against the opponent
    with the most similar parameter-band
    values.

//Online action selection
C1. Use game indexes to select the N most
    similar games.
C2. Of the selected N games, select the
    M games that best satisfy the goal
    criterion.
C3. Of the selected M games, select the most
    similar observation.
C4. Perform the action stored for the selected
    observation.
```

Game indexing (A1): We define a game's index as a vector of fitness values, containing one entry for each time step. These fitness values represent the desirability of all observed game states. To calculate the fitness value of an observed game state, we use the previously established evaluation function (denoted in Equation 1). Game indexing is supportive for later action selection, and as it is a computationally-expensive procedure, it is performed offline.

Clustering of observations (A2): As an initial means to cluster similar observations, we apply the standard k -means clustering algorithm [9]. The metric that expresses an observation's position in the cluster space is comprised of a weighted sum of the six observational features that also are applied for similarity matching. Clustering of observations is supportive for later action selection, and as it is a computationally-expensive procedure, it is performed offline.

Similarity matching (A2 and C3): To compare a given observation with another, we define six observational features, namely (1) phase of the game, (2) material strength, (3) commander safety, (4) positions captured, (5) economical strength, and (6) unit count. Similarity is defined by a weighted sum of the absolute *difference* in features values. The weighted sum for both clustering of observations and similarity matching is calculated as follows:

$$\begin{aligned}
 \textit{similarity} &= ((1 + \textit{diff}(\textit{phase_of_the_game})) \\
 &* (0.5 * \textit{diff}(\textit{unit_count}))) \\
 &+ \textit{diff}(\textit{material_strength}) \\
 &+ \textit{diff}(\textit{commander_safety}) \\
 &+ \textit{diff}(\textit{positions_captured}) \\
 &+ \textit{diff}(\textit{economical_strength})
 \end{aligned}$$

As observations are clustered, calculating the similarity between observations is relatively computationally-inexpensive. This is important, as similarity matching must be performed online.

Initialisation of game AI (B1-B3): To intelligently select the strategy initially followed by the game AI, we first determine which strategy the opponent is likely to use. In our game environment, perfect information of the opponent is available when loading the opponent game AI (i.e., the settings of all opponent parameters, indicating the strategic preferences, are known). We use this information to determine the initial strategy of the game AI. In a typical game-playing setting, however, such information is not directly available, but has to be established via statistical learning techniques such as opponent modeling. This will be investigated in future work. When the opponent strategy has been established, we determine in which parameter bands [8] the opponent strategy can be abstracted. We define three bands for each parameter, 'low', 'medium' and 'high'. We subsequently initialise the game AI with an effective strategy observed against the most similar opponent. We consider a strategy effective when in previous play

it achieved a set goal criterion (thus, the game AI will never be initialised with a predictably ineffective strategy), and consider opponents strictly similar when the observed values of the parameter bands are identical.

Action selection (C1-C4): Using the established game indexes, we select the N games with the smallest accumulated fitness difference with the current game, up until the current observation. Subsequently, of the selected N games, we perform the game action of the most similar observation of the M games that satisfy a particular goal criterion. The goal criterion can be any metric to represent preferred behaviour. For instance, a preferred fitness value of 0 can represent challenging gameplay, as this implies that players are equally matched. Naturally, we have to consider that performing actions associated to similar observations may not yield the same outcome when applied to the current state. Therefore, to estimate the effect of performing the retrieved game action, we straightforwardly compensate for the difference in metric value between the current and the selected observation.

EXPERIMENTS

This section discusses experiments that test our implementation of rapidly adaptive game AI. We first describe the experimental setup and the performance evaluation, and then the experimental results.

Experimental Setup

To test our implementation we start collecting observations of games where two game AIs are posed against each other. Multiple SPRING game AIs are available. We found one open-source game AI, which the author called ‘AAI’ [18]. We enhanced this game AI with the ability to collect game observations in a case base, and the ability to disregard radar visibility so that perfect information on the environment was available. As opposing player, we used the original AAI game AI. We found 27 parameters that define the strategic behaviour of the game AI.¹

For collecting observation, we simulate different players competing with different players, by for each game pseudo-randomising the strategic parameters of both players. The collection process was as follows. During each game, game observations were collected every 127 game cycles, which corresponds to the update frequency of AAI. With the SPRING game operating at 30

¹The concerning parameters determine the game strategy on a high, strategic level, and not so much on a low, tactical level. Three examples of these parameters are AIRCRAFT_RATE (determines how many airplane units the AI will build), MAX_MEX_DEFENCE_DISTANCE (maximum distance to base where the AI defends metal extractors), and MAX_SCOUTS (maximum of units scouting at the same time). The authors happily provide a full list of the parameters on request.




Figure 3: Two game AI’s pitted against each other on the map ‘SmallDivide’. In the screenshot, the white player has captured the centre position.

game cycles per second, this resulted in game observations being collected every 4.233 seconds.

We acknowledge that the amount of offline storage should be low for our approach to be considered practical for implementation in a game-production setting. We therefore store game observations in a lightweight fashion, by of each game observation only abstracting the position and unit-type of each unit. This abstraction, of approximately 3 KB per observation, provides a powerful basis for deriving observational features. Accordingly, a case base was built from 213.005 observations of 325 games, resulting in case base consisting of 679 MB of uncompressed observational data. Approaches are available to further reduce the size of the case base, such as offline data compression and subsequent online data decompression [2] and automatic condensation of the case base [3], however these lie outside the scope of the present research.

We collect observations from games played on the map ‘SmallDivide’. This map is also used for adaptation experiments. The map, illustrated in Figure 3, is a symmetrical map without water areas. All games are played under identical starting conditions.

For offline clustering of observations, k is set to ten per cent of the total number of observations. Before the game starts, the initial strategy is determined. While the game is in progress, action selection is performed at every phase transition. The parameter N for action selection is set to 50, and the parameter M is set to 5. The game action is expressed by the configuration of the 27 parameters of strategic behaviour.

Performance Evaluation

To evaluate the performance of the rapidly adaptive game AI, we determine to what extent it is capable of adapting effectively to game circumstances. We perform three different experiments. First, we test to what extent the rapidly adaptive game AI is capable of adapting to the original AAI game AI, set to play in a medium playing strength. Second, we test to what extent the rapidly adaptive game AI is capable of adapting to previously unobserved opponents, which is simulated by pitting the game AI against the original AAI game AI, initialised with randomly generated strategies. Third, we test to what extent the rapidly adaptive game AI is capable of upholding a tie, when pitted against the original AAI game AI, also set to play in a medium playing strength.

For each of the first two experiments, we perform a trial where the rapidly adaptive game AI is set to win the game (i.e., obtain a positive fitness value). For the third experiment, we set the rapidly adaptive game AI to uphold a tie (i.e., maintain a fitness value of 0, while never obtaining a fitness value less than -10, or greater than 10). To measure how well the rapidly adaptive game AI is able to maintain a fitness value of 0, the variance in fitness value is calculated. A low variance implies that the rapidly adaptive game AI has the ability to consistently maintain a predefined fitness value.

All experimental trials are repeated 25 times, except the trial where the rapidly adaptive game AI is pitted against randomly generated opponents, which is repeated 100 times.

Results

Table 1 gives an overview of the results of the first and second experiments performed in the SPRING game. Figure 4 displays the obtained fitness value as a function over time of two typical experimental runs. The results reveal that when pitted against the original AAI game AI, set to play in a medium playing strength, the rapidly adaptive game AI effectively obtains a victory (80% of the experimental runs). This result indicates that rapidly adaptive game AI is effective in play against the original AAI game AI.

In addition, the results reveal that when pitted against the original AAI game AI, initialised with randomly generated strategies, the rapidly adaptive game AI improves on the effectiveness of strictly randomised games played without adaptation mechanism (i.e., a verified effectiveness of 50%). The obtained improvement in effectiveness is 8%. This result indicates that even in randomised play against the AAI game AI, the rapidly adaptive game AI is able to find effective strategies in the case-base.

Table 2 gives an overview of the results of the third experiment. The results reveal that the rapidly adaptive game AI is capable of upholding a tie for a relatively

Opponent	#Games	Goal Ach.	Goal Ach.(%)	Impr.(%)
Original AAI	25	20	80%	30%
Random	100	58	58%	8%

Table 1: Effectiveness of rapidly adaptive game AI.

	Average	Standard deviation
Time to uphold tie	37 min.	15 min.
Variance in fitness value	2.04	0.59

Table 2: Upholding a tie with rapidly adaptive game AI.

long time (37 minutes on average), while at the same time maintaining a relatively low variance in the fitness value that is strived for.

DISCUSSION

In the experiments that test our implementation of rapidly adaptive game AI, we observed that the game AI was well able to achieve a victory when pitted against the original AAI game AI, set to play in a medium playing strength. We noticed that the rapidly adaptive game AI was able to find in the case base a strategy that could effectively defeat the original AAI game AI. As the original AAI game AI is not able to adapt its behaviour, the rapidly adaptive game AI could exploit its discovery indefinitely. Note that in some cases, the rapidly adaptive game AI did not win the game, despite it exhibiting strong behaviour. Such outliers cannot be avoided due to the inherent randomness that is typical to video games. For instance, in the SPRING game, the most powerful unit is able to destroy a Commander unit with a single shot. Should the Commander be destroyed in such a way, the question would arise if this was due to bad luck, or due to an effective strategy of the opponent. For game AI to be accepted as effective players, one could argue, recalling the previously mentioned need for consistent AI behaviour, that game AI should not force a situation that may be regarded as the result of lucky circumstances.

In addition, we observed that even in randomised play against the AAI game AI, the rapidly adaptive game AI is able to find effective strategies in the case base, and was thereby able to improve on the randomised performance by 8%. This is a satisfactory result. As randomised play may be considered a simulated way to test the game AI against previously unobserved opponents, naturally, the question remains if the performance in randomised play can be further enhanced. We discuss two approaches to enhance the performance in randomised play.

First, note that our case-base currently consists of observations collected over 325 games. For randomised play, determined by 27 pseudo-randomised behavioural parameters, it would be beneficial to collect more games in the case base in order to increase the probability




Figure 4: Obtained fitness values as a function over time, when pitted against the original AAI game AI. The figure displays a typical experimental result of (1) the rapidly adaptive game AI set to win the game, and (2) the rapidly adaptive game AI set to uphold a tie.

of it containing effective game strategies. As rapidly adaptive game AI can be expected to be applied in the playtesting phase of game development, and predictably in multi-player games, the case base in practical applications is expected to grow rapidly to contain a multitude of effective strategies.

Second, we observed that the final outcome of a SPRING game is largely determined by the actions performed in the beginning of the game. This exemplifies the importance of initialising the game AI with effective behaviour. In order to do so, one needs to accurately determine the opponent one will be pitted against. In video-game practice, (human) game opponents do not exhibit behaviour as random as in our experimental setup, but will typically exhibit behaviour that can be abstracted into a limited set of opponent models. Previous research has shown that even in complex RTS games such as SPRING, accurate models of the opponent player can be established [17]. We will therefore follow the expert opinion that game AI should not so much be focussed on directly exploiting current game observations, but should rather focus on effectively applying models of the opponent in actual game circumstances [16].

In addition, we found the rapidly adaptive game AI to be able to uphold a tie for a relatively long time, while at the same time maintaining a relatively low variance in the fitness value that is strived for. This ability may be regarded as a straightforward form of difficulty scaling. If a metric can be established that represents the preferred level of challenge for the human player, then in theory the rapidly adaptive game AI would be capable of scaling the difficulty level to the human player. Such a capability provides an interesting challenge for future research.

CONCLUSIONS AND FUTURE WORK

In this paper we discussed an approach to establish rapidly-adaptive game AI. In the approach, domain knowledge is gathered automatically by the game AI, and is immediately (i.e., without trials and without resource-intensive learning) utilised to evoke effective behaviour. In our implementation of the approach, game observations are collected in a case base. Subsequently, the case base is used to extract an evaluation function, and gathered cases are directly utilised by an adaptation mechanism. Results of experiments that test the approach in the SPRING game show that rapidly adaptive game AI can effectively obtain a victory and is capable of upholding a tie for a relatively long time. From these results, we may conclude that the established rapidly adaptive game AI provides a strong basis for effectively adapting game AI in actual video games. For future work, we will extend the established rapidly adaptive game AI with a means to scale the difficulty level to the human player. Subsequently, we will investigate how our approach to rapidly adapting game AI can be improved by incorporating opponent models.

ACKNOWLEDGEMENTS

This research is funded by a grant from the Netherlands Organization for Scientific Research (NWO grant No 612.066.406) and is performed in the framework of the ROLEC project.

REFERENCES

- [1] Agnar Aamodt and Enric Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications*, 7(1), March 1994.

- [2] Samir Abou-Samra, Claude Comair, Robert Champagne, Sun Tjen Fam, Prasanna Ghali, Stephen Lee, Jun Pan, and Xin Li. Data compression/decompression based on pattern and symbol run length encoding for use in a portable handheld video game system. US Patent 6416410, 2002.
- [3] F. Angiulli and G. Folino. Distributed nearest neighbor-based condensation of very large data sets. *IEEE Transactions on Knowledge and Data Engineering*, 19(12):1593–1606, 2007.
- [4] Sander Bakkes and Pieter Spronck. *AI Game Programming Wisdom 4*, chapter Automatically Generating Score Functions for Strategy Games, pages 647–658. Charles River Media, Hingham, MA., U.S.A., 2008.
- [5] Sander Bakkes, Pieter Spronck, and Jaap van den Herik. Phase-dependent evaluation in rts games. In Mohammad Mehdi Dastani and Edwin de Jong, editors, *Proceedings of the 19th Belgian-Dutch Conference on Artificial Intelligence (BNAIC)*, pages 3–10. Universiteit Utrecht, The Netherlands, 2007.
- [6] Michael Buro and Timothy M. Furtak. RTS games and real-time AI research. In *Proceedings of the BRIMS Conference*. Arlington VA, 2004.
- [7] Pedro Demasi and Adriano J. de O. Cruz. Online coevolution for action games. *International Journal of Intelligent Games and Simulation*, 2(3):80–88, 2002.
- [8] Roger Evans. *AI Game Programming Wisdom*, chapter Varieties of Learning, pages 571–575. Charles River Media, 2002.
- [9] J. A. Hartigan and M. A. Wong. A k -means clustering algorithm. *Applied Statistics*, 28(1):100–108, 1979.
- [10] Stefan Johansson, Jelmer Cnossen, and Tomaz Kunaver. Spring game engine. <http://spring.clan-sy.com/>, 2007.
- [11] S Johnson. *AI Game Programming Wisdom 2*, chapter Adaptive AI: A Practical Example, pages 639–647. Charles River Media, Inc., Hingham, MA, 2004.
- [12] Ronni Laursen and Daniel Nielsen. Investigating small scale combat situations in real-time-strategy computer games. Master’s thesis, Department of computer science, University of Aarhus, Denmark, 2005.
- [13] L. Liden. *AI Game Programming Wisdom 2*, chapter Artificial Stupidity: The Art of Making Intentional Mistakes, pages 41–48. Charles River Media, Inc., Hingham, MA, 2004.
- [14] Lev Manovich. *The Language of New Media*. The MIT Press, Cambridge, Massachusetts, U.S.A., 2002.
- [15] Alexander Nareyek. AI in computer games. *ACM Queue*, 1(10):58–65, 2004.
- [16] Steve Rabin. *AI Game Programming Wisdom 4*, chapter Preface - What’s happening to learning?, pages ix – xi. Charles River Media, Inc., 2008. ISBN 1-584-505230.
- [17] Frederik Schadd, Sander Bakkes, and Pieter Spronck. Opponent modeling in real-time strategy games. In Marco Roccetti, editor, *Proceedings of the GAME-ON 2007*, pages 61–68, 2007.
- [18] Alexander Seizinger. AI:AAI. Creator of the game AI ‘AAI’, <http://spring.clan-sy.com/wiki/AI:AAI>, 2006.
- [19] Pieter Spronck, Marc Ponsen, Ida Sprinkhuizen-Kuyper, and Eric Postma. Adaptive game AI with dynamic scripting. *Machine Learning*, 63(3):217–248, 2006.
- [20] Richard S. Sutton. Learning to predict by the methods of temporal differences. *Machine Learning*, 3:9–44, 1988.
- [21] Laurie N. Taylor. Video games: Perspective, point-of-view, and immersion. Masters thesis, Graduate Art School, University of Florida, U.S.A., 2002.
- [22] Paul Tozour. *AI Game Programming Wisdom (ed. Rabin, S.)*, chapter The Perils of AI Scripting, pages 541–547. Charles River Media, 2002.