
An Adaptive Training Framework for Increasing Player Proficiency in Games and Simulations

Simon Demediuk

RMIT University
Melbourne, VIC 3000, AUS
simon.demediuk@rmit.edu.au

William L. Raffe

RMIT University
Melbourne, VIC 3000, AUS
william.raffe@rmit.edu.au

Xiaodong Li

RMIT University
Melbourne, VIC 3000, AUS
xiaodong.li@rmit.edu.au

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.
Copyright is held by the owner/author(s).
CHI PLAY'16 Extended Abstracts, October 16-19, 2016, Austin, TX, USA
ACM 978-1-4503-4458-6/16/10.
<http://dx.doi.org/10.1145/2968120.2987735>

Abstract

To improve a player's proficiency at a particular video game, the player must be presented with an appropriate level of challenge. This level of challenge must remain relative to the player as their proficiency changes. The current fixed difficulty settings (e.g. easy, medium or hard) provide a limited range of difficulty for the player. This work aims to address this problem through developing an adaptive training framework that utilizes existing work in Dynamic Difficulty Adjustment to construct an adaptive AI opponent. The framework also provides a way to measure the player's proficiency, by analysing the level of challenge the adaptive AI opponent provides for the player. This work tests part of the proposed adaptive training framework through a pilot study that uses a real-time fighting game.

Author Keywords

Artificial Intelligence; Dynamic Difficulty Adjustment; Simulation and Games; Simulation Training

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous; I.2.1 Artificial Intelligence: Games; K.3.0 Computers and education: General

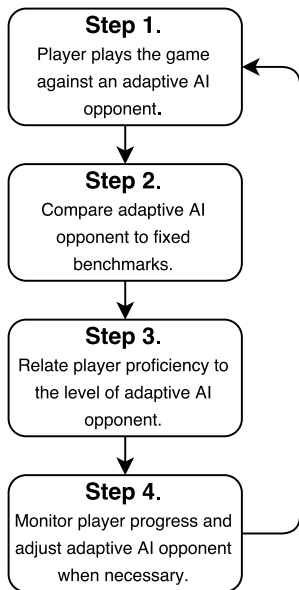


Figure 1: The steps of the adaptive training framework

Introduction

There are a number of reasons why someone might want to improve their proficiency at a particular video game. They may find the game enjoyable, want to train to become a professional in e-Sports, or perhaps they are training a specific skill through simulation (e.g. military simulation training [12]). To improve proficiency at a particular video game, it has been found that players respond best when the game is pitched at an optimum challenge level [11]. When the game is at an appropriate level, the player is challenged and engaged, yet when it is too easy or too difficult players can become frustrated or disengaged.

Currently, the player must select a difficulty level (e.g. easy, medium or hard) prior to commencing the game. This can be a difficult decision, especially if the player has never played the game before. It is also problematic if they select a suitable difficulty setting at the beginning but improve their proficiency while playing, since it may not be possible to alter the difficulty level during the game without restarting. These fixed difficulty settings may also not provide an appropriate level of challenge for all players. In addition, as the player improves, the gap between the difficulty settings may be too large.

There has been a focus in the literature on tailoring a game's difficulty to the level of the player through the use of adaptive Artificial Intelligence (AI) opponents [15] [1] [5] [6] or game environments [16] [10]. This work is sometimes referred to as Dynamic Difficulty Adjustment and involves the process of changing the strategies and behaviour of the adaptive AI opponent or environment, to produce a video game that is more enjoyable, interesting, and most importantly, one that alters its level of challenge as the proficiency of the player changes.

Although Dynamic Difficulty Adjustment provides a way to

tailor the level of challenge to a player, it does not guarantee that the player's proficiency will improve. Since the adaptive AI opponent is always matching the player's proficiency level, when the player does not improve, the adaptive AI opponent similarly does not alter its level of challenge. Therefore, in order to ensure that the player improves over time, it is necessary to monitor their proficiency level and, in some cases, to change how the adaptive AI opponent tailors its level of challenge.

This work-in-progress outlines an adaptive training framework that can enable the development of a challenging game, tailored to each individual player through Dynamic Difficulty Adjustment. This framework also has the ability to monitor a player's progress and to change the adaptation method employed to ensure that the desired proficiency outcomes are met. More specifically, this work will show how a player's proficiency level can be measured, through investigating the adaptive AI opponent that the player plays against, and comparing the adaptive AI opponent to fixed benchmarks. The aim of this work is to provide a way to increase the player's motivation and participation, and to enable them to learn the task at a faster pace than is currently available. This paper also details a pilot study that explores the implementation of parts of this framework.

Adaptive Training Framework

The steps of the adaptive training framework are shown in Figure 1. This framework operates on two primary assumptions: first, that the game will involve a player in competition with an opponent; second, that the player will participate in the game multiple times, until they achieve the required level of proficiency.

This framework consists of four steps. The first step draws upon the existing work in the area of Dynamic Difficulty Ad-




Figure 2: Screenshots from the real-time fighting game developed in Universal Fighting Engine

justment. During Step 1, the player interacts with the game, competing against an adaptive AI opponent. This adaptive AI opponent tailors its level of challenge using Dynamic Difficulty Adjustment. To achieve this, the adaptive AI opponent will frequently change its strategies and behaviour, based on their interaction with the player in real-time. Initially, the adaptive AI opponent will aim to achieve a game outcome where the player has a 50% chance of winning the game, since this has been suggested to be the most enjoyable experience for the player [8]. The player having a 50% chance of winning the game occurs when both the player and their opponent are evenly matched.

After the completion of a game, Step 2 consists of evaluating the difficulty level of the adaptive AI opponent. This is achieved by comparing the behaviour of the adaptive AI opponent competing against the player, to the behaviour of the adaptive AI opponent competing against a range of fixed difficulty level benchmarks (e.g. easy, medium and hard). The process by which the adaptive AI opponent is compared to these benchmarks will vary depending on the type of game or the adaption method employed.

Step 3 involves correlating the difficulty level of the adaptive AI opponent to the proficiency level of the player. This step assumes that in the case where the adaptive AI opponent and the player are evenly matched, the player is at the same relative level of proficiency as their adaptive AI opponent.

Now that a metric for measuring the proficiency level of the player has been established, Step 4 monitors how the player's proficiency level changes over a series of games. The aim is for the player to achieve a certain level of proficiency and at this point to cease playing the game. Until this level is achieved, the process will repeat again from Step 1.

However, there may be various problems in this iterative process. For example, the player may improve at a very slow rate, or they may reach their stable learning phase [2] (i.e. the state in which their proficiency is no longer improving) before they have acquired the desired proficiency level. This may occur for a number of reasons, one being that the adaptive AI opponent is either too challenging or not challenging enough. In order to combat these problems, in Step 4 a different game outcome may be set (e.g. a 60% chance of winning, rather than 50%). Alterations such as this, can provide more or less of a challenge to the player, which in turn may change the rate of their improvement. This will change the calculation of the player's proficiency in Step 3, as the player and the adaptive AI opponent are no longer evenly matched.

This framework aims to provide the player with a motivating and engaging game in which to improve their proficiency level at a certain task. This framework also has the advantage of providing a mechanism capable of monitoring the progress of the player, and providing a means to modify the adaptive AI opponent, two factors which will promote progress.

Pilot Study

To assess the adaptive training framework as a means to determine a player's proficiency level, a small pilot study was conducted using members of the research team. A real-time fighting game similar to Street Fighter [14] or Mortal Combat [7] was developed internally using the Universal Fighting Engine (UFE)¹ (see Figure 2). In this game the player controls a character that can move forward or backward, jump and crouch in a two dimensional arena. The player uses actions such as punches, kicks and fireballs to

¹Universal Fighting Engine, www.ufe3d.com

attack the opponent's character or to block incoming enemy attacks. The aim of the fighting game is to reduce the opponent's health to 0, before they reduce your character's health to 0. A player's proficiency in a real-time fighting game is reflected in what action or attack they choose, when they choose it, how quickly they react to the changing game state, and how well they use a combination of attacks at appropriate times to defeat their opponent.

Adaptive AI Opponent

Following the framework, an adaptive AI opponent was developed that changed its feature set based on its current health and the player's current health. A feature set defines the rate at which the adaptive AI opponent selects attacks over blocks (aggressiveness), the time between making decisions based on the current state and the time taken to implement these decisions (reaction time). The fuzzy logic engine provided by UFE² was used to evaluate the actions available to the adaptive AI opponent based on its feature set, and consequently the best action is chosen to be performed.

Five different feature sets were used to represent different difficulty settings; very defensive, defensive, balanced, aggressive and very aggressive. The settings for the feature sets is shown in Table 1. To novice players, the more defensive the adaptive AI opponent behaves, the easier it is for them to win, whereas the more aggressive the adaptive AI opponent behaves, the harder it is for novices players to win.

The adaptive AI opponent begins each game in the balanced feature set. The adaptive AI opponent then increases its difficulty level, changing from its current feature set to the next more aggressive feature set, when its health is lower

<i>Feature set</i>	ΔT Decisions <i>seconds</i>	ΔT Actions <i>seconds</i>	<i>Aggressiveness</i> 0 – 1
V. Defensive	0.4	0.1	0.1
Defensive	0.3	0.1	0.3
Balanced	0.15	0.05	0.45
Aggressive	0.1	0.01	0.6
V. Aggressive	0	0	0.7

Table 1: The settings used for each feature set. The adaptive AI opponent will weigh attacks higher than defensive moves when the aggressiveness is greater than 0.5.

than the player's health. Inversely, the adaptive AI opponent lowers its difficulty level by changing from its current feature set to the next more defensive feature set, when its health is higher than the player's health. The settings for this adaptive AI opponent are described in the following experimental setup.

Experimental Setup

The experiment conducted involved two members of the research team, who stated that they were novices of the real-time fighting genre. They each played a series of 10 games against the adaptive AI opponent. Each game was a best out of 3 rounds format, where in each round both characters start with 1000 health points and the characters fight each other in close combat until one character's health reaches 0.

Work done by [4] suggests that an adaptive AI opponent which wins or loses within a $\pm 10\%$ health margin is an appropriate level of challenge for an evenly matched game. Due to the limited and naive nature of a health based adaption system, the closest game outcome that was achieved was a $\pm 20\%$ health differential. To achieve this outcome

²Fuzzy Logic Engine <http://www.ufe3d.com/doku.php/ai:start>

Δ Health	Feature Set
$> 40\%$	V. Defensive
$> 20\%$	Defensive
$20\% > 0 > -20\%$	Balanced
$< -20\%$	Aggressive
$< -40\%$	V. Aggressive

Table 2: The feature set used based on the health difference between the adaptive AI opponent and the player.

the adaptive AI opponent changed its feature set as mentioned above, following the Table 2.

The player's proficiency level was calculated after each game and their average proficiency level over the 10 games was estimated. To confirm this proficiency level the players then played a best out of 5 game series against a fixed AI using the feature set that was one level higher than their estimated proficiency level. Additionally, they played another best out of 5 game series against a fixed AI using the feature set that one level lower than their estimated proficiency level.

Results

To see how well the adaptive AI opponent performed its task of adapting to the individual player, the average resultant health of each game was recorded. Figure 3 shows these results with a positive percentage being a win for the player and a negative percentage being a loss for the player. The value of the percentage represents the average health differential for the game. Even though the game outcome margin that adaptive AI opponent aimed for, was not as small as that used by [4], the adaptive AI opponent was still able to adapt appropriately to each player in nearly all games. To determine the proficiency of the player, some




Figure 3: The average resultant health differential for each game. Dashed lines represent the target game outcome of the adaptive AI opponent.

benchmark proficiency levels were established by playing the adaptive AI opponent against five different fixed difficulty AIs: Very Aggressive, Aggressive, Balanced, Defensive and Very Defensive. The time spent by the adaptive AI opponent in each feature set when adapting to the fixed difficulty AI was recorded. These times were averaged over the 5 games and normalised, with the result being saved as a five dimensional vector.

In addition, the time spent by the adaptive AI opponent in each feature set during the 10 games played against each human player was also recorded and normalised, and saved as a five dimensional vector. Using the cosine similarity function, the vectors of the adaptive AI opponent when it played against each human player, were compared to the vector of the adaptive AI opponent when it played the fixed difficult AI benchmarks. This enabled the proficiency of each player to be calculated for each game, with the results shown in Figure 4. Based on the proficiency level calculated for each player, there was an expectation that each player would win more games against the fixed AI when its




Figure 4: The proficiency level of each player for each game. Dashed lines represent the average proficiency of each player.

difficulty level was lower than their average proficiency, and less games against the fixed AI when its difficulty level was greater than theirs. For Player 1, their expected proficiency on average was between balanced and aggressive. Player 2's proficiency was on average between defensive and balanced. To confirm this, Player 1 played five games both against a fixed aggressive AI and a fixed balanced AI (see Table 3). Player 2 played five games against both a fixed balanced AI and a fixed defensive AI (see Table 4). The results of these games confirm that this method is capable of measuring a player's proficiency level.

	Wins	Losses
Aggressive	2	3
Balanced	3	2

Table 3: Player 1's results when playing against the fixed AI opponents.

	Wins	Losses
Balanced	1	4
Defensive	3	2

Table 4: Player 2's results when playing against the fixed AI opponents.

Conclusion and Future Work

Although this work did not test framework's ability to influence the player's rate of learning, this small pilot study shows that even with a non-ideal adaptive AI opponent, the player's proficiency level can still be determined. This gives confidence in the framework and its ability to calculate the proficiency level of each player for each game. In this work an average proficiency level for each player over the 10 games was calculated and used to determine their proficiency level. This number of games was insufficient to determine the rate of learning for the player. Over a larger number of games this average proficiency level could be

calculated using a sliding window method and this would show how the average proficiency of the player changes over time. This would consequently allow for the player's rate of learning to be calculated and thus enable further work in investigating how this rate changes when the adaption method is altered between games.

In future work, large scale trials will be conducted to fully test the adaptive training framework, and investigate if the framework significantly improves the rate of learning a task. For these trials a more effective Dynamic Difficulty Adjustment system, such as Challenge Sensitive Action Selection [3] will be implemented into the real-time fighting game. Challenge Sensitive Action Selection is a Reinforcement Learning [13] based approach, where the action selected in each state is dictated by the difficulty that the adaptive AI opponent is trying to achieve. The actions for each state are ranked based on their outcome, and as the AI difficulty increases, so does the rank of the action selected. [3] demonstrates that this system is capable of closely adapting to a player, while not requiring a limited number of fixed feature sets as used in this work. Additionally, with this method a larger range of fixed AIs could be developed, i.e producing a fixed AI for each ranked action.

These trials will also enable the formulation of a method to relate the adaptive AI opponent's proficiency to the player's proficiency when they are not evenly matched. A ranking system such TrueSkill [9] is being investigated as a robust way to measure, rank and represent the player's proficiency level. Additionally, implementing a Dynamic Difficulty Adjustment method such as Challenge Sensitive Action Selection will also allow for a more accurate measurement of the difficulty of the adaptive AI opponent, removing the limitations of the cosine similarity method used in this work.

REFERENCES

1. Amy L Alexander, Tad Brunyé, Jason Sidman, and Shawn A Weil. 2005. From gaming to training: A review of studies on fidelity, immersion, presence, and buy-in and their effects on transfer in pc-based simulations and games. In *The interservice/industry training, simulation, and education conference (IITSEC), NTSA, Orlando, Florida*.
2. Gustavo Andrade, Geber Ramalho, Alex Sandro Gomes, and Vincent Corruble. 2006. Dynamic Game Balancing: An Evaluation of User Satisfaction. *AIIDE* 1 (2006), 3–8.
3. Gustavo Andrade, Geber Ramalho, Hugo Santana, and Vincent Corruble. 2005a. Challenge-sensitive action selection: an application to game balancing. In *Intelligent Agent Technology, IEEE/WIC/ACM International Conference on*. IEEE, 194–200.
4. Gustavo Andrade, Geber Ramalho, Hugo Santana, and Vincent Corruble. 2005b. Extending reinforcement learning to provide dynamic game balancing. In *Proceedings of the Workshop on Reasoning, Representation, and Learning in Computer Games, 19th International Joint Conference on Artificial Intelligence (IJCAI)*. 7–12.
5. Phillipa M Avery and Zbigniew Michalewicz. 2010. Adapting to human gamers using coevolution. In *Advances in Machine Learning II*. Springer, 75–100.
6. Alexander Baldwin, Daniel Johnson, and Peta A Wyeth. 2014. The effect of multiplayer dynamic difficulty adjustment on the player experience of video games. In *CHI'14 Extended Abstracts on Human Factors in Computing Systems*. ACM, 1489–1494.
7. Ed Boon and Tobias John. 1992. *Mortal Combat*. Game [Arcade]. (8 October 1992). Midway Games, Chicago, Illinois, U.S.
8. Shengbo Guo, Scott Sanner, Thore Graepel, and Wray Buntine. 2012. Score-based bayesian skill learning. In *Machine Learning and Knowledge Discovery in Databases*. Springer, 106–121.
9. Ralf Herbrich, Tom Minka, and Thore Graepel. 2006. Trueskill: A Bayesian skill rating system. In *Advances in Neural Information Processing Systems*. 569–576.
10. Robin Hunicke and Vernell Chapman. 2004. AI for dynamic difficulty adjustment in games. In *Challenges in Game Artificial Intelligence AAAI Workshop*. 91–96.
11. Thomas W Malone and Mark R Lepper. 1987. Making learning fun: A taxonomy of intrinsic motivations for learning. *Aptitude, learning, and instruction* 3, 1987 (1987), 223–253.
12. Roger Smith. 2009. The long history of gaming in military training. *Simulation & Gaming* (2009).
13. Richard S Sutton and Andrew G Barto. 1998. *Reinforcement learning: An introduction*. Vol. 1. MIT press Cambridge.
14. Takashi Nishiyama. 1987. *Street Fighter*. Game [Arcade]. (30 August 1987). CAPMCOM, Osaka, Japan.
15. Georgios N Yannakakis and John Hallam. 2009. Real-time game adaptation for optimizing player satisfaction. *Computational Intelligence and AI in Games, IEEE Transactions on* 1, 2 (2009), 121–133.
16. Georgios N Yannakakis and Julian Togelius. 2011. Experience-driven procedural content generation. *Affective Computing, IEEE Transactions on* 2, 3 (2011), 147–161.