

Agent-Based Adaptive Level Generation for Dynamic Difficulty Adjustment in Angry Birds

Matthew Stephenson,^{1,2} Jochen Renz,¹

¹Research School of Computer Science, Australian National University, Canberra, A.C.T. 0200, Australia

²Department of Data Science and Knowledge Engineering (DKE), Maastricht University, Netherlands

Abstract

This paper presents an adaptive level generation algorithm for the physics-based puzzle game Angry Birds. The proposed algorithm is based on a pre-existing level generator for this game, but where the difficulty of the generated levels can be adjusted based on the player's performance. This allows for the creation of personalised levels tailored specifically to the player's own abilities. The effectiveness of our proposed method is evaluated using several agents with differing strategies and AI techniques. By using these agents as models / representations of real human player's characteristics, we can optimise level properties efficiently over a large number of generations. As a secondary investigation, we also demonstrate that by combining the performance of several agents together it is possible to generate levels that are especially challenging for certain players but not others.

Introduction

Procedural level generation (PLG), where levels for a game are created automatically without the need for human designers, is a key area of investigation for video game research (Hendrikx *et al.* 2013; Togelius *et al.* 2011). PLG can be extremely useful for increasing a game's length and replayability, as it allows a large number of levels to be created in a relatively short time. It is also possible to tailor the generated levels towards specific user's playstyles, known as adaptive level generation, which allows for a unique and personalised gameplay experience (Yannakakis and Togelius 2011). Dynamic difficulty adjustment (DDA) is a form of adaptive level generation, where the difficulty of generated levels is modified to better suit the player's current skill based on their performance (Hunicke 2005). This is accomplished by modifying certain generator parameters that control different level features, so that the end result is more likely to achieve the desired amount of challenge for the player.

This paper presents an adaptive level generator for the physics-based puzzle game Angry Birds. This game has been used substantially in AI research over the past few years, primarily for developing agents and level generators, as the game's environment presents more realistic physical

constraints compared to most traditional video games. Successfully generating levels for Angry Birds that are equally as challenging as human-designed levels is a difficult task, but will likely be necessary for Angry Birds agents to improve beyond their current capabilities. Previous level generation methods for Angry Birds used either a heuristic calculation based on level properties or the performance of several agents to help set the difficulty of a level. However, as different players often possess varying levels of ability, many people would likely find these levels too hard or easy to solve. This is also a problem for training and evaluating agents, as levels where most agents either can or cannot solve them yield very little discriminatory information (Stephenson *et al.* 2018a). We therefore suggest an agent-based adaptive generation method for dynamic difficulty adjustment, where the generator adjusts the difficulty of its levels depending on the player's performance. This method can also be used to generate levels that are difficult for one player whilst being easy for another, exploiting the player's own strengths or weaknesses.

The remainder of this paper is organised as follows. Section 2 describes the large amount of background and related work, both for Angry Birds and adaptive level generation in general. Section 3 presents our proposed adaptive generation method. Section 4 describes our conducted experiments and results. Section 5 discusses what these results could mean for both human players and agents, Section 6 concludes this work and outlines future possibilities.

Background

Adaptive Level Generation

While most games that contain some form of PLG typically use generic generation techniques that are not influenced by the player's behaviour, adaptive level generation, also referred to as experience-driven, personalised or player-centred level generation, takes the player's behaviour into account (Shaker *et al.* 2016). Examples of game or level characteristics that could be adjusted for specific players include qualities such as difficulty, engagement, frustration, enjoyment, complexity, learning potential, etc. These properties are indirectly controlled by adjusting certain parameters of the generator. Different players will likely behave or perform differently even when playing the same game. If an

accurate model of the player can be determined, then this can be used to tailor the properties of the generated content towards their individual preferences.

Constructing a model of the player is a difficult and imprecise task, but is essential for adaptive level generation to be effective. Methods for determining player behaviour include analysing their performance across several “test” levels, or using a questionnaire for measuring more intangible qualities. This information can then be used to directly evaluate generated content in the future, allowing us to estimate whether it will be suitable for the player. Another approach, and the one that we will be using in this paper, is to use AI agents to estimate the quality of levels (i.e. agent / simulation-based evaluation functions). Using a collection of agents as representations of different playing styles or abilities allows us to generate levels that are suited to a particular player, or a collection of levels that require several different techniques to solve them.

Examples of genre’s where adaptive level generation has been used to improve the player experience include board games (Marks and Hom 2007), racing games (Togelius *et al.* 2007), action-RPG (Heijne and Bakkes 2017), rogue-like (Stammer *et al.* 2015), tower defence (Sutoyo *et al.* 2015), and platformers (Shaker *et al.* 2012; 2010).

Dynamic Difficulty Adjustment Dynamic difficulty adjustment (DDA) is often considered to be one of the simplest and most common forms of game adaption, where the difficulty of a game increases or decreases if the player is performing too well or poorly respectively (Wheat *et al.* 2015). Because a player’s performance in a game can typically be evaluated without the need for questionnaires or overly complex estimations, DDA can usually be implemented in most games without significant issue. Nearly all games feature some form of increasing difficulty as the skill of the player increases, but this element is often lost or overly simplified with most PLG approaches. However, just because estimating the difficulty of a game for a specific player is relatively simple compared to other more complex behavioural characteristics, this certainly doesn’t make the task trivial. The difficulty of a game can often be multi-dimensional in nature, where the same level could be considered hard or easy for various different reasons (Jennings-Teats *et al.* 2010). Players can often have unbalanced skill sets, where they are adept at overcoming certain tasks or challenges more than others. One player may be very good at forward planning, another at making precise actions, another with fast response times, and so on. A successful DDA system should therefore be able to adjust the difficulty of its generated levels in many different ways, and also detect which of these most influence the player’s performance.

Agent-Based Evaluation One approach for evaluating generated content is to utilise AI agents with different strategies to play through the generated levels (Wheat *et al.* 2015; Shaker *et al.* 2010; Togelius *et al.* 2007). By selecting the agent that best models the player’s abilities, we can then use this agent as a player surrogate in the adaptive level generation process. This approach has several benefits. First, agents can often be used to play levels much faster than a normal

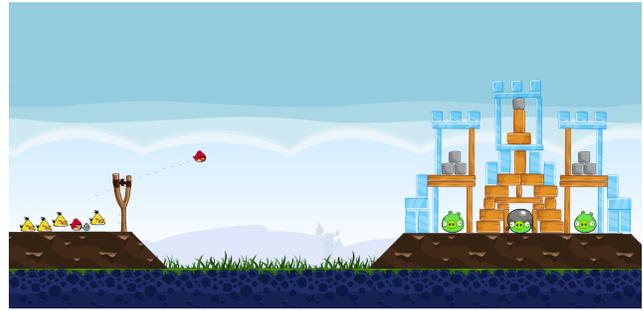


Figure 1: Screenshot of a level from the Angry Birds game.

person, allowing us to evaluate a larger number of levels in a much shorter time (requiring volunteers to playtest hundreds of levels is not very practical). Agents can also typically give more accurate estimations of certain level properties (especially difficulty) than by just analysing the level’s features. Human players are also likely to improve the longer they play, making repeated performances inconsistent between different experiments. The downside of this method is that it naturally requires a large and diverse range of agents to already exist, which Angry Birds thankfully has (agents described in more detail later).

Angry Birds

Angry Birds is a popular physics-based puzzle game where the player’s objective for each level is to kill all pigs using a set number of birds. A typical Angry Birds level like the one shown in Figure 1, requires the player to shoot the birds they have from a slingshot at structures made of blocks that are protecting the pigs. All objects within the level obey simplified physics principles defined by the game’s engine. Blocks can come in several different shapes and materials, and birds can also be one of several different types (all with differing properties). Pigs and blocks can be killed / destroyed by hitting them with either a bird or another object. Points are awarded to the player once the level is solved (all pigs in the level have been killed) based on the number of birds remaining and the total amount of damage caused. The source code for the official Angry Birds game is not currently available, so a modified version of the Unity-based clone known as Science Birds, originally created by Lucas Ferreira (Ferreira and Toledo 2014), was used instead.

Level Generation Several level generators have been presented for Angry Birds in recent years, some of which have attempted to adapt the generated content based on the player’s experience. Previous work by (Kaidan *et al.* 2015; 2016) attempted to measure the difficulty of Angry Birds levels based on their features, and take this into account during the generation process. The generator they present is based on the same genetic algorithm described in (Ferreira and Toledo 2014), but where the fitness function for evaluating generated levels has been modified to take the desired difficulty of the level into account. The first approach simply used the number of pigs within a level as a measure of difficulty (Kaidan *et al.* 2015). The desired number of pigs for

each level would then be adjusted over multiple generations, based on the number of pigs that the player was able to kill in previous levels. An alternative measure of difficulty was proposed in a subsequent paper (Kaidan *et al.* 2016), which attempted to estimate the difficulty of a level based on its overall impact factor. This was calculated based on the ERA relations between objects within the level. In both these prior cases, the fitness function for the generator rewards levels with an estimated difficulty closest to the desired amount, which in turn is based on the player’s performance for previous levels. However, neither of these approaches use agents to evaluate levels and their estimations of difficulty are based solely on a level’s features, which are controlled by only a small number of generator parameters.

Instead of changing the difficulty, some prior generators investigated other level aspects that might influence the player’s experience. The Tanager generator evaluated the immersion and design quality of its generated levels using an on-line user study (Ferreira and Toledo 2018). This was in the form of a questionnaire which asked users to rate both automatically and manually created levels in terms of their enjoyment, engagement and challenge. The Funny quotes generator creates levels based on words or quotes for three levels of difficulty (Jiang *et al.* 2017). Certain generator parameters were manually configured for each difficulty category, based on the results of a user study into the average solve and retry rates of players across different levels. A follow up investigation using a similar version of this generator, modified future levels based on chat comments made by players inside the game (Jiang *et al.* 2018). Words within these comments were used as input parameters when generating future levels (i.e. the generator adjusts the levels it creates based on what the players type to each other).

The adaptive generator presented in this paper is based on the Iratus Aves generator described in (Stephenson and Renz 2017b; 2016a; 2016b), which was also the winner of the 2017 and 2018 Angry Birds level generation competitions (AIBirds 2018). The output of this generator can be partially controlled by changing the values of different input parameters (i.e. the generator’s parameter set). This search-based generator previously used a direct fitness function approach to modify generator parameters over several generations based on desirable level properties. Instead of evaluating a generated level based solely on its observable features, we implement a new agent-based fitness function which uses the performance of several Angry Birds agents. Angry Birds agents have been utilised for a small number of previous generators (Stephenson and Renz 2017b; Ferreira and Toledo 2018), but only to check if a generated level is solvable. This paper uses agents to evaluate and evolve the generated levels in a deeper and more meaningful way.

Agents A wide variety of Angry Birds agents have been developed over the past six years for the AIBirds competition (Renz 2015; Renz *et al.* 2016; 2015). These agents employ a range of different AI techniques, including qualitative reasoning (Walega *et al.* 2016), internal simulation analysis (Polceanu and Buche 2013; Schiffer *et al.* 2016), logic

programming (Calimeri *et al.* 2016), heuristics (Dasgupta *et al.* 2016), Bayesian inferences (Tziortziotis *et al.* 2016; Narayan-Chen *et al.* 2013), and structural analysis (Zhang and Renz 2014). In this paper we selected four different agents to assist with evaluating generated levels. These were the Naive, Datalab, SeaBirds and Eagle’s Wing agents. The Naive agent is the simplest agent available, making it a perfect model of a novice player. The remaining three agents (referred to as “skilled agents”) are some of the best performing agents currently available (Stephenson *et al.* 2018b; Stephenson and Renz 2017a), although they are still well below that of a normal human, with each agent having their own strengths and weaknesses (Stephenson and Renz 2018). Further details about the specific strategies and AI techniques used by each of these agents can be found in (Stephenson *et al.* 2018b).

Methodology

To reiterate our proposed method using previous terminology, we present an agent-based evaluation function for level generation, which allows for dynamic difficulty adjustment in Angry Birds and other similar physics-based puzzle games. To achieve this, we need both a way to evaluate the player’s performance and a way to update the level generator’s output based on this performance. By using a search-based generation approach, we can evolve an initial population of parameter sets for our level generator over many generations using a fitness function (Jennings-Teats *et al.* 2010; Shaker *et al.* 2012).

A general overview of the adaptive level generation process is as follows:

1. Measure the performance of the player and all available agents on a randomly generated collection of levels, and select the agent that best models the player (e.g. lowest root-mean-square error).
2. Randomly create an initial population of parameter sets (individuals) for our level generator.
3. Generate a level for each individual in the population and record each agent’s performance on these levels.
4. Use these agent performance distributions to calculate a fitness value for each individual in the population.
5. Evolve this population using a genetic algorithm (selection, crossover, mutation, elitism, etc.) based on each individual’s fitness value (i.e. create a new generation).
6. Stop once a desired number of generations has been reached; otherwise repeat from step 3.

By following this process, the average fitness of the parameter sets (and levels generated using them) within our population should increase over multiple generations.

Note that step 1 of this process essentially selects an agent to act as a representation of our player in all subsequent steps, and is therefore unnecessary if the player is already an agent (i.e. only needed for human players).

Adaptive Level Generator

As previously mentioned, our proposed adaption method is based on the same Angry Birds level generator described

Generator Input Parameter	Value Range
Number of pigs	1 - 15 (integer)
Number of birds	1 - 8 (integer)
Number of ground structures	1 - 5 (integer)
Number of platform structures	0 - 4 (integer)
Maximum number of TNT	0 - 4 (integer)
Weights for each bird type (x5)	0.0 - 1.0 (float)
Weights for each material (x3)	0.0 - 1.0 (float)
Weights for each block shape (x13)	0.0 - 1.0 (float)

Table 1: Input parameters for our adaptive level generator and their possible value range (minimum - maximum).

in (Stephenson and Renz 2017b; 2016a; 2016b). This generator previously used a fitness function to evaluate and update the probability of selecting different block shapes based on certain features of the generated levels. For our adaption method to be successful, we need to be able to control more level properties than just the frequency of block shapes. We therefore extended the number of input parameters that affect the generated levels, see left column of Table 1. Apart from the increased number of adjustable input parameters, the level generation algorithm itself was not changed.

The first four parameters define the number of pigs, birds and structures (both on the ground and platforms) within the generated level. The fifth parameter determines the maximum number of TNT boxes that the level can contain (could potentially be less than this value depending on the available space). The last three parameters are lists of values that define weightings for each bird type (five options), material (three options) or block shape (thirteen options). Unlike the previous parameters, these weight inputs do not directly define specific level features, but instead influence the probability of selecting their respective elements (i.e. if the weight value of one block shape is twice that of another, then that block shape has twice the chance of being selected during level generation). While all weight inputs are float values between zero and one, integer inputs are limited to within a fixed value range, see right column of Table 1. Each parameter set within our population contains values for each of these generator input parameters (genome length of 26).

Difficulty Estimation

Whilst prior methods for estimating the difficulty of an Angry Birds level relied solely on its observable features, we instead propose a more accurate approach based on agent performance. This allows us to not only better estimate the difficulty of levels overall, but also means that the same level can be given multiple difficulty scores based on different player’s abilities. Angry Birds has two basic measures of success. The first is simply solving each level and the second is achieving a large score for each level, with the score for a level being awarded after it is solved. This score element to solving levels allows for an additional degree of depth when comparing different agents. Perhaps one agent solves a level less often than another agent, but typically achieves a higher score when it does. We therefore proposed two possible dif-

ficulty measures (D_{solve} and D_{score}) of a level (L) for an agent (A_i), see Equations 1 and 2.

$$D_{solve}(A_i, L) = 1 - \frac{\#TimesSolved(A_i, L)}{\#Attempts(A_i, L)} \quad (1)$$

$$D_{score}(A_i, L) = 1 - \frac{AverageScore(A_i, L)}{MaximumScore(L)} \quad (2)$$

Both D_{solve} and D_{score} can be any value between zero and one (normalised). $MaximumScore(L)$ is defined as the theoretical score that could be achieved if all pigs and blocks within L were destroyed using only the first bird.

Essentially, D_{solve} uses the agent’s solve-rate for a level as the measure of difficulty, whilst D_{score} uses the score-rate. Deciding which difficulty measure to use depends on the desired property of the generated levels.

Fitness Function

Now that we can estimate the difficulty of a level for a specific agent, we can define fitness functions that use this to evaluate the parameter sets within our population. The fitness value for each parameter set is based on the difficulty measures of our agents for a level generated using it. Many different fitness functions could be defined that each represent a desired performance distribution of our agent(s), but we will only focus on two in this paper.

The first function defines the fitness of a level in terms of the probability that our agent is able to solve the level each time they attempt it, see Equation 3, where A_s is the specific agent that the generated levels are being adapted for, D_{solve} is the observed solve-rate, and D_{target} is the target / desired solve-rate.

$$Fitness_p(A_s, L) = 1 - abs(D_{solve}(A_s, L) - D_{target}) \quad (3)$$

This allows us to define the desired difficulty of a generated level for a specific agent as a percentage (i.e. if we want an agent to solve each generated level 50% of the times it attempts it, then we simply set D_{target} to 0.5). D_{score} could also be used as our difficulty measure for this function instead of D_{solve} , but trying to define a desired score-rate for a level as a fraction of the total score possible is conceptually harder to understand than simply the desired solve-rate.

The second fitness function is more complex and utilises several different agents. Instead of adapting our generated levels to a fixed solve-rate for a specific agent, it is also possible to adapt our generated levels to be especially hard for our chosen agent when compared to the performance of other agents, see Equation 4, where A is the set of all available agents.

$$Fitness_m(A_s, L) = D_{score}(A_s, L) - \min_{A_n \in A} (D_{score}(A_n, L)) \quad (4)$$

Using this fitness function will favour levels that our specific agent performs poorly in, but where other agents perform better. Essentially, adapting the generated levels using

this fitness function will focus on our specific agent’s weaknesses, while using the inverse of this function will generate levels that focus on its strengths. Using D_{score} as our difficulty measure rather than D_{solve} allows us to still compare the performances of different agents, even when their solve-rates are very similar (i.e. using score-rate gives a more precise measure of performance than solve-rate).

To summarise, $Fitness_p$ gives a higher value to levels that are closer to the desired solve-rate for a specific agent, whilst $Fitness_m$ gives a higher value to levels that a specific agent finds relatively difficult compared to other agents.

Genetic Algorithm

Once a fitness value for each parameter set in our current population has been calculated, a genetic algorithm is used to evolve the population and create the next generation. Individuals are selected from the current population using stochastic universal sampling (Baker 1987). This selection technique reduces the risk of individuals with a large fitness value being overrepresented in the next population (i.e. gives individuals with a lower fitness a greater chance of being chosen). This is desirable, as the uncontrollable stochastic elements of our generator and agents means that the fitness value for each parameter set is likely to be only a rough estimate of its actual fitness. An elitism scheme was also used to select a percentage number of individuals in each generation with the highest fitness value, and include these in the next generation unchanged. Uniform crossover and mutation genetic operators were then used to create the new generation (offspring) from the previously selected individuals of the current generation (parents). Mutations for each parameter set value must be within the possible minimum and maximum range for that parameter, as described in Table 1.

Experiments and Results

Two experiments were conducted using our proposed adaptive level generation algorithm, for each of the fitness functions previously described. The first experiment investigated our adaptive generator’s ability to create levels with a desired solve-rate ($Fitness_p$) for both a novice player (Naive agent) and an expert player (hyper-agent created from the skilled agents). The second experiment investigated whether our adaptive generator could successfully create levels that skilled agents performed better on than the Naive agent ($Fitness_m$ with the Naive agent as A_s), essentially generating levels that exploited the Naive agent’s limitations.

The specific values used for our genetic and level evolution algorithms are as follows. Parameter sets were adapted over 30 generations, with a population size of 50 individuals. Elitism was set at a rate of 8%, crossover probability at 25%, and mutation probability at 15%. Each agent was given a three minute time limit to play each generated level on a heavily sped up version of the modified Science Birds game. While the number of attempts each agent was able to complete for each generated level can fluctuate depending on the agent’s design and the level’s features, each agent took roughly 2-3 seconds between making each shot. This means that even if a generated level contained the maximum

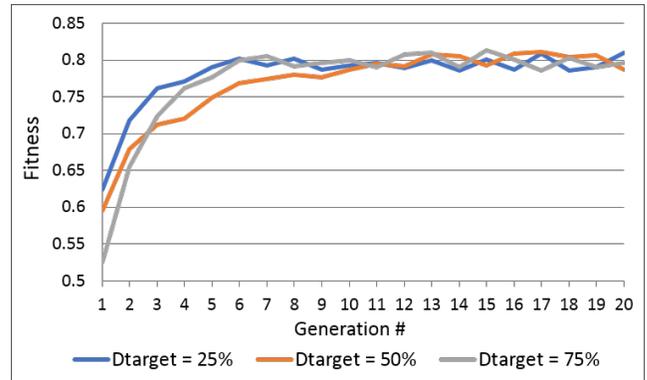


Figure 2: Average $Fitness_p$ value of each generation for the Naive agent.

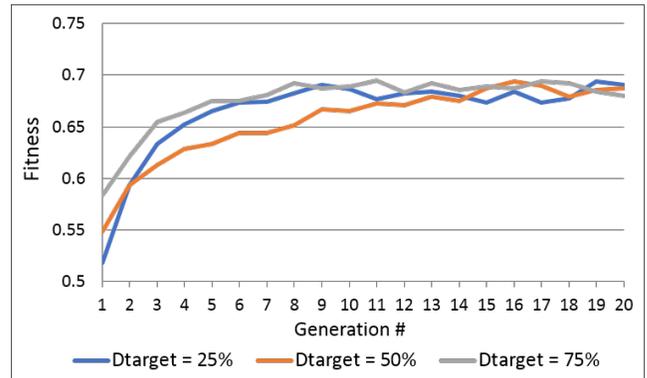


Figure 3: Average $Fitness_p$ value of each generation for the skilled hyper-agent.

number of birds (eight), each agent would still get at least six or seven attempts to solve it.

Due to the high variance in our generator’s output and each agent’s performance, the individual results of our experiments were highly stochastic. We therefore repeated each experiment ten times to reduce any inaccuracy issues, with the results displayed in the following sections being the averaged results over all ten repeated experiments. Please also note that even though our experiments were run over 30 generations the graphs displaying our results only show the first 20 generations, as the average fitness of our generated levels never increased significantly past this point.

Percentage Solvability

This experiment investigated the effectiveness of our $Fitness_p$ function for the Naive agent and different D_{target} values. We also performed the same analysis for a hyper-agent that selects from our three skilled agents (Datalab, SeaBirds and Eagle’s Wing), based on the same score prediction models described in (Stephenson and Renz 2017a). Three D_{target} values were tested, 25%, 50% and 75%. The average $Fitness_p$ values over all parameter sets in each generation, for both the Naive agent and the skilled hyper-agent, are shown in Figures 2 and 3 respectively. From these results

we can see that the average fitness of the generated levels for all agents and D_{target} values increased over the generations tested. However, the rate at which the fitness increased and the optimal fitness value that could be reasonably achieved, appears to be different for each agent and D_{target} pairing.

When using a D_{target} value of 50%, our adaptive generator took longer to reach a high fitness value compared to the other D_{target} values. This was likely because levels that could never be solved and levels that could always be solved had an equal $Fitness_p$ value. Due to the large number of highly variable parameters that can influence the difficulty of a generated level, it would be very easy for our adaptive generator to only produce levels that would probably be impossible to solve, by simply making the number of pigs and structures very high whilst also making the number of birds very small. The opposite can also be done to generate levels that are incredibly easy to solve. Both of these types of levels are also very likely to occur in the randomly generated initial population. As a result of these factors, it is easier for our adaptive generator to create levels for the lower or higher D_{target} values in the earlier generations, as it can initially focus on simply creating either impossibly hard or ridiculously easy levels respectively. Using a D_{target} value of 50% treats both these cases as equally desirable, meaning that our adaptive generator must find a suitable balance between the two. This naturally takes more time to accomplish, but over a large number of generations the average fitness of the generated levels eventually equals that of the other D_{target} values.

Comparing individual D_{target} values, it also appears that adapting generated levels for the skilled hyper-agent took longer to reach a high fitness value when compared to the Naive agent. The maximum fitness value that could be achieved also appeared to be less for the skilled hyper-agent, only around 0.69 compared to the Naive agent’s maximum fitness of around 0.81. This was likely due to the skilled hyper-agent have more strategies and behaviours that must be “learned” by our level adaption algorithm (i.e. combinations of multiple level properties probably required to construct levels of a suitable difficulty).

Relative Solvability

This experiment investigated the effectiveness of our $Fitness_m$ function for evolving levels that the Naive agent performed poorly on relative to the performance of more skilled agents. Generated levels should not only be hard for the Naive agent (which could easily be achieved using the $Fitness_p$ function and setting D_{target} to a very small value) but should also be easier for the skilled agents to solve with a larger score. If successful, this would essentially create levels that require a certain degree of skill to perform well on, an idea that is often represented within traditional human-designed levels. The average $Fitness_m$ values over all parameter sets in each generation are shown in Figure 4. Please note that as $Fitness_m$ is calculated using the D_{score} measure, which is based on the theoretical maximum score that could possibly be obtained for a level and is often significantly higher than any realistically achievable score, the $Fitness_m$ values for levels are significantly lower than the

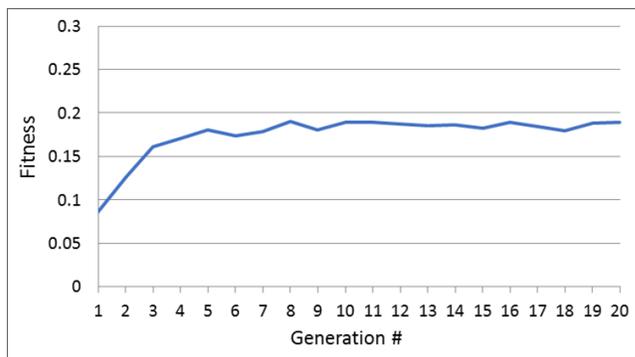


Figure 4: Average $Fitness_m$ value of each generation for the Naive agent.



Figure 5: Generated level with a high $Fitness_m$ value.

previous $Fitness_p$ values.

From these results we can see that the average fitness of the generated levels increased slightly over the generations tested. This result is promising, as it means that it is possible to generate levels that favour certain agents over others. By manually comparing the parameter sets of the generated levels with the highest $Fitness_m$ values it would seem that, apart from simply being harder overall, levels that the Naive agent struggled the most with compared to the skilled agents contained more TNT boxes and bird types with difficult to use abilities (yellow, blue and white birds). This makes sense, as our Naive agent doesn’t directly target TNT and doesn’t vary the tap time for activating bird’s abilities (unlike the skilled agents). This observation is also backed up by a previous investigation into deceptive Angry Birds level design (Stephenson and Renz 2018).

An example of a level that was generated using an evolved parameter set based on our $Fitness_m$ function is shown in Figure 5. This level had a (relatively) high fitness value of 0.62, indicating that at least one of the skilled agents was able to significantly outperform the Naive agent.

Discussion

Using an agent-based adaption method to adjust the difficulty of generated levels has many potential uses. Being able to generate personalised content for human players has previously been shown to increase user engagement

and overall enjoyment in games (Togelius and Yannakakis 2016), but what we are most interested in discussing here is how adaptive level generation can be used to help improve agent development. Angry Birds agents that attempt to use some form of reinforcement learning to solve unknown levels have become increasingly popular over the past few years at the annual AIBirds competition (AIBirds 2018; Stephenson *et al.* 2018b), but have so far failed to demonstrate any of the exceptional performance this technique has exhibited for many other games. In fact, many of these agents often rank among the lowest performing Angry Birds agents currently available. One of the main reasons for this poor performance is believed to be a lack of available levels for training purposes, something that PLG can help address (Justesen *et al.* 2018). We believe that the adaptive generation method proposed in this paper can potentially be used to improve the performance of reinforcement learning agents better than simply using randomly generated levels, and that adaptive generation can also be used to evaluate and help identify weaknesses within non-learning agents as well.

Firstly, for generating levels with a fixed percentage solve-rate for a specific agent ($Fitness_p$). When training an agent it is often desirable to focus on levels that the agent can occasionally solve, while still leaving plenty of areas to improve upon. Levels that the agent currently performs very well on every time do not give much new information to learn, whilst levels that the agent can never solve also give little information for the opposite reason. This issue is especially important in a game like Angry Birds as reward is only given to the agent when it solves a level, making any accumulated score from previous shots meaningless if the level is not also solved (i.e. delayed reward). Generating adapted levels that a learning agent can currently solve some of the time (e.g. using a D_{target} value of 50%) will likely help the agent improve quicker. Although this hypothesis is yet to be demonstrated, it seems to us like a reasonably intuitive idea.

Secondly, for generating levels that are relatively hard for a specific agent compared to other agents ($Fitness_m$). Similar to using the $Fitness_p$ function, training on levels where our reinforcement learning agent performs poorly compared to other agents, that perhaps even use different AI techniques, might help to improve learning efficiency. Using this approach has the advantage that it can create levels which emphasise the learning agent’s most obvious weaknesses more than others, ensuring that the learning agent’s more pressing limitations are attended to first (i.e. ensures that the learning agent is at least on an equal performance to other agents before attempting to improve beyond this). This approach could also potentially be used for non-learning agents, allowing us to identify flaws in our strategies that need improving the most (i.e. understand where other agents are outperforming us). Another use is for benchmarking multiple or new agents, where it is often desirable to test on a collection of levels that produce a large variation in prior agent performance (Stephenson *et al.* 2018a). This could be achieved by generating a small subset of levels with a high $Fitness_m$ value for each previous agent, and then combining these subsets together to give our benchmark set for a new or improved agent.

Conclusions and Future Work

In this paper we have presented an adaptive level generator for Angry Birds, that uses agents to adjust the difficulty of the generated levels based on the player’s performance. Levels are generated using a search-based approach, with several different adjustable parameters. A genetic algorithm and fitness function based on agent performance is then used to evolve the generator’s parameters over multiple generations. Levels can be generated for specific player solve-rates, or that are especially hard for the current player relative to the performance of others. Several experiments were conducted that demonstrated the effectiveness of our adaptive generator for both these requirements on a variety of agents. This approach can be used to create personalised levels for human players, as well as improving the usefulness of generated levels for training and evaluating agents.

While the experiments presented in this paper demonstrate the effectiveness of our adaptive generator (at least when using agents as human surrogates), there are several areas that could be improved in the future. The most obvious improvement would be to increase the number of generator input parameters that can be adjusted, as well as testing our method on a greater number of agents. We could also attempt to integrate the Iratus Aves level generator with other Angry Birds generators, increasing the variety of levels that could be created. It would also be good to analyse our adaptive generator’s ability to cope with learning agents, whose performance might improve or change over time. We also hope to be able to investigate our hypothesis that adaptive level generation can improve the generality and effectiveness of reinforcement learning agents even more so than regular level generation algorithms. The approaches presented in this paper can also be easily extended to other physics-based puzzle games with similar mechanics.

References

- AIBirds. AIBirds homepage. <https://aibirds.org>, 2018. Accessed: 2018-10-25.
- James E. Baker. Reducing bias and inefficiency in the selection algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application*, pages 14–21, 1987.
- F. Calimeri, M. Fink, S. Germano, A. Humenberger, G. Ianni, C. Redl, D. Stepanova, A. Tucci, and A. Wimmer. Angry-HEX: An artificial player for Angry Birds based on declarative knowledge bases. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(2):128–139, 2016.
- S. Dasgupta, S. Vaghela, V. Modi, and H. Kanakia. s-Birds Avengers: A dynamic heuristic engine-based agent for the Angry Birds problem. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(2):140–151, 2016.
- L. Ferreira and C. Toledo. A search-based approach for generating angry birds levels. In *Computational Intelligence and Games (CIG), 2014 IEEE Conference on*, pages 1–8, 2014.
- L. N. Ferreira and C. F. M. Toledo. Tanager: A generator of feasible and engaging levels for angry birds. *IEEE Transactions on Games*, 10(3):304–316, 2018.
- N. Heijne and S. Bakkes. Procedural zelda: A pcg environment for player experience research. In *Proceedings of the 12th Inter-*

- national Conference on the Foundations of Digital Games*, pages 11:1–11:10, 2017.
- M. Hendrikx, S. Meijer, Joeri Van Der Velden, and A. Iosup. Procedural content generation for games: A survey. *ACM Trans. Multimedia Comput. Commun. Appl.*, 9(1):1–22, 2013.
- Robin Hunicke. The case for dynamic difficulty adjustment in games. In *Proceedings of the 2005 ACM SIGCHI International Conference on Advances in Computer Entertainment Technology*, pages 429–433, 2005.
- M. Jennings-Teats, G. Smith, and N. Wardrip-Fruin. Polymorph: A model for dynamic level generation. In *Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 138–143, 2010.
- Y. Jiang, T. Harada, and R. Thawonmas. Procedural generation of angry birds fun levels using pattern-struct and preset-model. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 154–161, 2017.
- Y. Jiang, P. Paliyawan, T. Harada, and R. Thawonmas. An audience participation angry birds platform for social well-being. In *GAME-ON'2018*, pages 1–6, 2018.
- N. Justesen, R. R. Torrado, P. Bontrager, A. Khalifa, J. Togelius, and S. Risi. Illuminating generalization in deep reinforcement learning through procedural level generation, 2018. arXiv:1806.10729v2.
- M. Kaidan, C. Y. Chu, T. Harada, and R. Thawonmas. Procedural generation of angry birds levels that adapt to the player's skills using genetic algorithm. In *2015 IEEE 4th Global Conference on Consumer Electronics (GCCE)*, pages 535–536, 2015.
- M. Kaidan, T. Harada, C. Y. Chu, and R. Thawonmas. Procedural generation of angry birds levels with adjustable difficulty. In *IEEE Congress on Evolutionary Computation*, pages 1311–1316, 2016.
- J. Marks and V. Hom. Automatic design of balanced board games. In *Proceedings of the Artificial Intelligence and Interactive Digital Entertainment International Conference*, pages 25–30, 2007.
- A. Narayan-Chen, L. Xu, and J. Shavlik. An empirical evaluation of machine learning approaches for Angry Birds. In *IJCAI Symposium on AI in Angry Birds*, 2013.
- M. Polceanu and C. Buche. Towards a theory-of-mind-inspired generic decision-making framework. In *IJCAI Symposium on AI in Angry Birds*, 2013.
- J. Renz, X. Ge, S. Gould, and P. Zhang. The Angry Birds AI competition. *AI Magazine*, 36(2):85–87, 2015.
- J. Renz, X. Ge, R. Verma, and P. Zhang. Angry Birds as a challenge for artificial intelligence. In *AAAI Conference on Artificial Intelligence*, pages 4338–4339, 2016.
- J. Renz. AIBIRDS: The Angry Birds artificial intelligence competition. In *AAAI Conference on Artificial Intelligence*, pages 4326–4327, 2015.
- S. Schiffer, M. Jourenko, and G. Lakemeyer. Akbaba: An agent for the Angry Birds AI challenge based on search and simulation. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(2):116–127, 2016.
- N. Shaker, G. Yannakakis, and J. Togelius. Towards automatic personalized content generation for platform games. In *Proceedings of the Sixth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 63–68, 2010.
- N. Shaker, G. N. Yannakakis, J. Togelius, M. Nicolau, and M. O'Neill. Evolving personalized content for super mario bros using grammatical evolution. In *Conference on Artificial Intelligence and Interactive Digital Entertainment*, pages 75–80, 2012.
- N. Shaker, J. Togelius, and M. J. Nelson. *Procedural Content Generation in Games: A Textbook and an Overview of Current Research*. Springer, 2016.
- D. Stammer, T. Gnther, and M. Preuss. Player-adaptive spelunky level generation. In *2015 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 130–137, 2015.
- M. Stephenson and J. Renz. Procedural generation of complex stable structures for Angry Birds levels. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8, 2016.
- M. Stephenson and J. Renz. Procedural generation of levels for Angry Birds style physics games. In *Twelfth AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment (AIIDE-16)*, pages 225–231, 2016.
- M. Stephenson and J. Renz. Creating a hyper-agent for solving angry birds levels. In *AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 2017.
- M. Stephenson and J. Renz. Generating varied, stable and solvable levels for angry birds style physics games. In *IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8, 2017.
- M. Stephenson and J. Renz. Deceptive angry birds: Towards smarter game-playing agents. In *Proceedings of the 13th International Conference on the Foundations of Digital Games*, pages 13:1–13:10, 2018.
- M. Stephenson, D. Anderson, A. Khalifa, J. Levine, J. Renz, J. Togelius, and C. Salge. A continuous information gain measure to find the most discriminatory problems for ai benchmarking. *CoRR*, abs/1809.02904:1–8, 2018.
- M. Stephenson, J. Renz, X. Ge, and P. Zhang. The 2017 AIBIRDS Competition, 2018. arXiv:1803.05156v1.
- R. Sutoyo, D. Winata, K. Oliviani, and D. M. Supriyadi. Dynamic difficulty adjustment in tower defence. *Procedia Computer Science*, 59:435 – 444, 2015.
- J. Togelius and G. N. Yannakakis. Emotion in games: Theory and praxis. pages 155–166. Springer, 2016.
- J. Togelius, R. De Nardi, and S. M. Lucas. Towards automatic personalised content creation for racing games. In *IEEE Symposium on Computational Intelligence and Games*, pages 252–259, 2007.
- J. Togelius, G. N. Yannakakis, K. O. Stanley, and C. Browne. Search-based procedural content generation: A taxonomy and survey. *IEEE Transactions on Computational Intelligence and AI in Games*, 3(3):172–186, 2011.
- N. Tziortziotis, G. Papagiannis, and K. Blekas. A bayesian ensemble regression framework on the Angry Birds game. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(2):104–115, 2016.
- P. A. Walega, M. Zawidzki, and T. Lechowski. Qualitative physics in Angry Birds. *IEEE Transactions on Computational Intelligence and AI in Games*, 8(2):152–165, 2016.
- D. Wheat, M. Masek, C. P. Lam, and P. Hingston. Dynamic difficulty adjustment in 2d platformers through agent-based procedural level generation. In *2015 IEEE International Conference on Systems, Man, and Cybernetics*, pages 2778–2785, 2015.
- G. N. Yannakakis and J. Togelius. Experience-driven procedural content generation. *IEEE Transactions on Affective Computing*, 2(3):147–161, 2011.
- P. Zhang and J. Renz. Qualitative spatial representation and reasoning in Angry Birds: The extended rectangle algebra. In *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning*, pages 378–387, 2014.