Illuminating Game Space Using MAP-Elites for Assisting Video Game Design

Martin Balla [1] and **Adrián Barahona-Ríos** [2] and **Adam Katona** [3] and **Nuria Peña Pérez**[4] and **Ryan Spick** [5]

Abstract. In this paper we demonstrate the use of Multi-dimensional Archive of Phenotypic Elites (MAP-Elites), a divergent search algorithm, as a game design assisting tool. The MAP-Elites algorithm allows illumination in the game space instead of just determining a single game setting via objective based optimization.

We showed how the game space can be explored by generating a diverse set of game settings, allowing the designers to explore what range of behaviours are possible in their games.

The proposed method was applied to the 2D game Cave Swing. We discovered different settings of the game where a Rolling Horizon Evolutionary Algorithm (RHEA) agent behaved differently depending on the selected game parameters. The agent's performance was plotted against its behaviour for further exploration, which allowed visualizing how the agent performed with selected behaviour traits.

1 Introduction

In the context of video games, the search space can be seen as the game space, defined by all the possible combination of game parameters (such as gravity, distance between objects, force applied when jumping, etc.). Search algorithms aim to find optimal game parameters for a particular game evaluation function (fitness function) by exploring the game space. The solution to the search problem is therefore a parameter combination that gets the highest score from the fitness function. However, while search algorithms traditionally focus on finding the best combination of game parameters, finding a set of diverse solutions that lead to "good" games can also be interesting, especially from a design point of view.

Recent literature has explored the game space to find different game variants. Isaksen et. al. used Monte Carlo Tree Search (MCTS) to automatically test the difficulty of various points in the game space [7]. The authors used a player model based on human motor skills (precision, reaction time and actions per second), allowing to retrieve a certain point in the game space linked to the desired difficulty level. Few other studies have focused on finding game variants associated to different agent behaviours instead of level difficulties. Tremblay et. al. compared MCTS with different search algorithms (A* and Rapidly-exploring Random Trees) to explore the player trajectories in platform games [16]. Different game levels were hand crafted and several game solutions were found for each single level. Game space has also been explored to find unique game variants by clustering the behaviours found for an agent playing different game levels [6]. However, these methods do not allow for automatic exploration of different playing behaviours by searching the game space.

Procedural Content Generation (PCG), in the context of video games, refers to the use of software to automatically generate content [14]. Content can be in the form of, for instance, game assets (such as textures, 3D models, etc.) sound, dialog trees or mechanics. Given its programmatic nature, this technique can accelerate the video game development cycle, reducing the development cost. Considering the increasing cost of game development [15], PCG is a very attractive solution, especially for bigger open world games. PCG can also be used to assist creativity, helping developers to generate new ideas faster.

This work proposes the use of a search algorithm to automatically look for multiple, high performing game variants based on different desired user-defined playing behaviour. The Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) algorithm [11] can be used to illuminate the search space. In this way, different variants of the same video game (with different parameters) can be played, resulting in diverse ways of playing the game (different behaviours). In this work a Rolling Horizon Evolutionary Algorithm agent (RHEA) was used to evaluate each point in the game space, obtaining its performance and behaviour with that particular set of parameters. This allowed to map the game space to the desired set of features in the behaviour space. By using MAP-Elites to search the game space, designers could define the behaviour characterization, which results in a diverse set of games, where different play styles are required to achieve the game's objective. In this work we tested MAP-Elites on the game Cave Swing, where we found various levels, requiring many different play styles.

2 Background

2.1 Optimization algorithms

Optimization algorithms have been traditionally used to find the best solution of a given parameter set. To determine the best solution, a fitness function is required, which can be the score in the game or some designed heuristics. As most games have a high number of parameters, it is usually impractical to find the best combination by manual tuning. Parameters can be discrete, where a pre-defined set of values are used, or continuous, where the values are arbitrary within a defined range. Most optimization algorithms require the user to define discrete values, such as Grid Search or NTBEA [10]. Continuous parameter optimization is more complex as more combinations arise. Popular continuous optimization algorithms are Random Search and

[1] Queen Mary University of London, United Kingdom, email: m.balla@qmul.ac.uk
[2] University of York, United Kingdom, email: ajbr501@york.ac.uk
[3] University of York, United Kingdom, email: ak1774@york.ac.uk
[4] Queen Mary University of London, United Kingdom, email: n.penaperez@qmul.ac.uk
[5] University of York, United Kingdom, email: rjs623@york.ac.uk

CMA-ES [5]. One advantage of MAP-Elites is that it can be used for both discrete and continuous optimization problems.

2.2 MAP-Elites

MAP-Elites is a divergent search algorithm which belongs to a family of algorithms called Quality Diversity (QD) [13]. The goal of Quality Diversity algorithms is to find a diverse set of high quality solutions, instead of a single best solution. By optimizing for both performance and diversity, these algorithms often outperform purely objective based optimization by avoiding getting stuck in local optima. Quality Diversity algorithms rely on a user defined Behaviour Characterization (BC). This BC is used to measure similarity between solutions, which allows the algorithms to directly search for diversity. The BC function is domain dependent. In the maze solving domain [9], where a robot have to navigate in a maze, the BC can be the trajectory taken by the robot, or the final position of the robot. In a six legged robot walking domain [2] the BC can be defined by how much each leg touches the ground. While some QD algorithms like Novelty Search with Local Competition [9] focuses more on quality, MAP-Elites puts more emphasis on diversity. This is achieved by considering multiple dimensions of behaviours separately, instead of just calculating the distance between two behaviours. This property makes MAP-Elites especially useful for exploring the search space, by discovering all kind of different behaviours.

A recent work by Gravina et al. [4] surveys QD algorithms in the context of game design. MAP-Elites have been used for designing levels in Super Mario Bros [17], Dungeon levels [1], Bullet Hell Simulation [8] and to balance player decks in Hearthstone [3]. Our work differs from these methods by tuning the game's parameters directly, not just the parameters of the level generator or using MAP-Elites to output raw levels. We also visualize the behaviour map in the form of heatmaps and agent trajectories to get a better insight of the resulting behaviours.

2.3 Cave Swing

Cave Swing is a tap timing game that consists of travelling along a cave of certain width and height by shooting a rope that can anchor to specific locations. A run of the game is successful when an agent manages to travel all the way along the cave in a certain amount of time while avoiding accidentally hitting any of the borders of the map (see Figure 1).

Only two actions are available for this game, a "null" action and "shooting" action, which throws the rope so that it attaches to the nearest available anchor location. Once a rope is anchored, the agent remains hanging from it until another shooting action takes place. The physics of the game are relatively simple, with the movement of the hanging agent depends on both the pulling force exerted by the rope and an external force. The rope is modelled as an elastic of zero natural length, so its pulling force is determined by its stiffness k. This force is multiplied by a loss factor (see Table 1 for information regarding all game parameters). The external force acting on the agent can be picture as a gravity (G) that pulls the agent on both the horizontal (G_x) and vertical directions(G_y). The rope can only attach to the anchor locations, which are placed at a certain height depending on the map dimensions.

The score of this game is calculated for each time frame, and is therefore available at all times during a run of the game (see Figure 1). The calculation of the score is based on how much the agent has progressed in the x direction, how high it is on the vertical direction and how fast it is completing the level. This is defined in equation 1, for each given time t.

$$Score = xP_x + yP_y - tP_t \tag{1}$$

Where x and y are respectively the horizontal and vertical position at any given time t (in game ticks). Px, Py and Pt are the points for x and y positions and the cost per time spent. If the agent succeeds in the run, it will add to this score a positive bonus, but if it fails a penalty will be subtracted from its score. The values of all the costs are defined in Table 1.

2.4 Rolling Horizon Evolutionary Algorithm

To evaluate the different levels, an agent was required, which played reasonably well. We chose a Statistical Forward Planning agent called Rolling Horizon Evolutionary Algorithm (RHEA)[12].

At each step RHEA constructs a random sequence of actions, which is executed in the forward model (given the current state and an action returns the next state) and a score from the reached state is calculated. The action sequence gets mutated and evaluated again. This process is repeated until a time or an iteration budget is elapsed, and then the first action of the highest-scoring action sequence gets executed. A shift buffer is used, which allows the agent to keep the previously evolved sequence, execute only the first action, shift every element by one position to the left and fill the last position by a random element. We avoid total random mutation, which means that every element of the list cannot be mutated more than once per mutation. The parameters used for the RHEA agent in this experiment can be found in Table 2. The fitness function in our case is the numerical score that Cave Swing provides.

Table 1: Game Parameters. Fixed parameters are those that were kept as a default value. These includes the score-related parameters and parameters that could affect the selected behaviour features. Explored parameters are those that were tuned by the MAP-Elites.

Fixed Parameters	**Default Value**	
Map Width	2500	
Map Height	250	
Anchors Height	100	
Maximum Ticks	500	
Points per x	1000	
Points per y	1000	
Cost per Tick	10	
Success Bonus	-10	
Failure Penalty	1	
Explored Parameters	**Min. Value**	**Max. Value**
Number of Anchors	5	20
Gravity in x	-1	1
Gravity in y	-1	2
Rope Stiffness	0.005	0.1
Loss Factor	0.99	0.99999

3 Methodology

3.1 Experimental Procedure

The code used in this paper is available on Github at: https://github.com/martinballa/MAPElitesCaveSwing. Our implementation has two main parts: The MAP-Elites algorithm and the

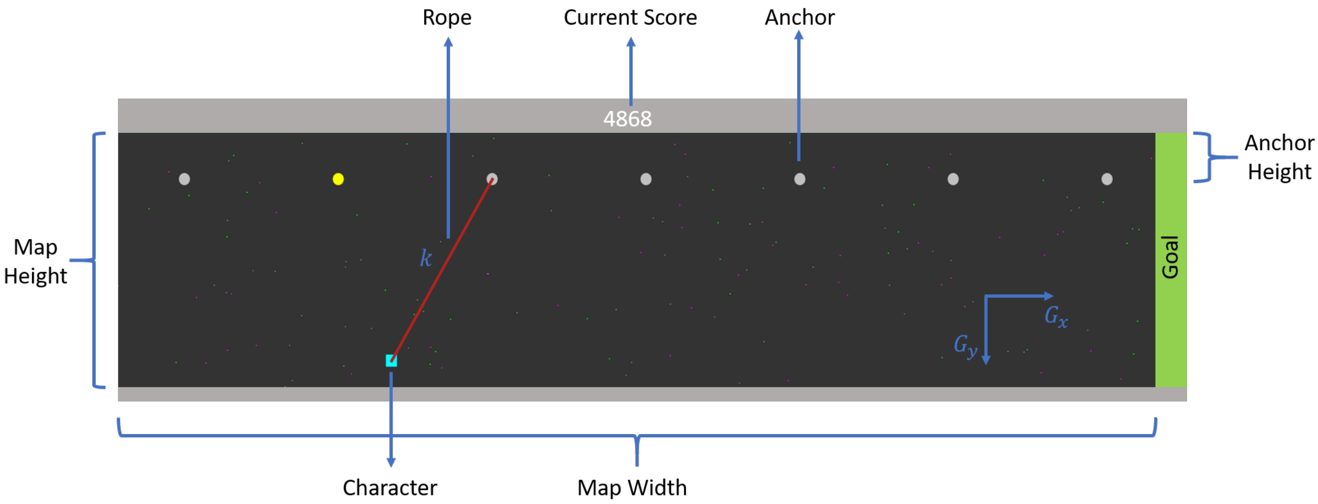

Figure 1: Schematics of the game *Cave Swing*. The agent controlling the character travels along the cave hanging from a rope of stiffness k, which can only attach to the anchors. G_x and G_y are the horizontal and vertical components of the external gravity force acting on the agent. The game finishes if the agent hits any of the grey borders of the map or exceeds a certain amount of time, giving a failure result; or if the agents reaches the goal, giving a successful result.

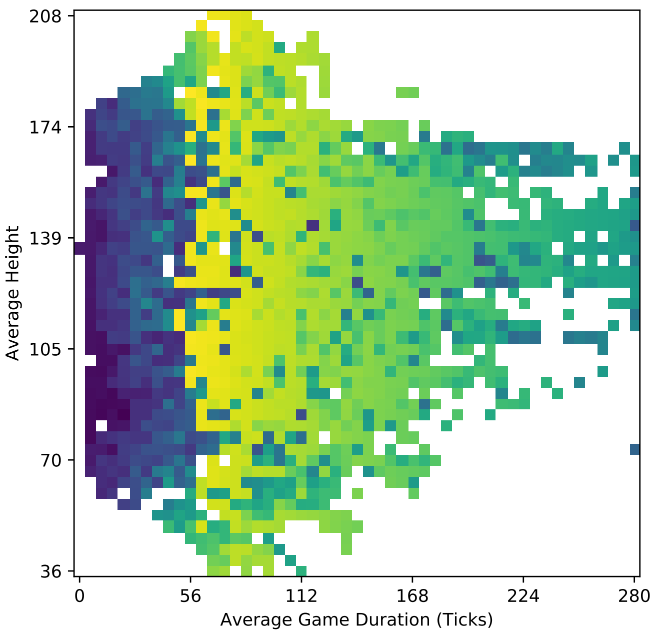

Figure 2: Average game ticks (X axis) against average game height (Y axis). The intensity of the heatmap represents the score achieved by the agent, with the given parameter set (the brighter, the higher scores were achieved by the agent). White regions correspond to areas where no behaviours were found (i.e. either impossible to achieve or not explored by MAP-Elites in 10,000 iterations).

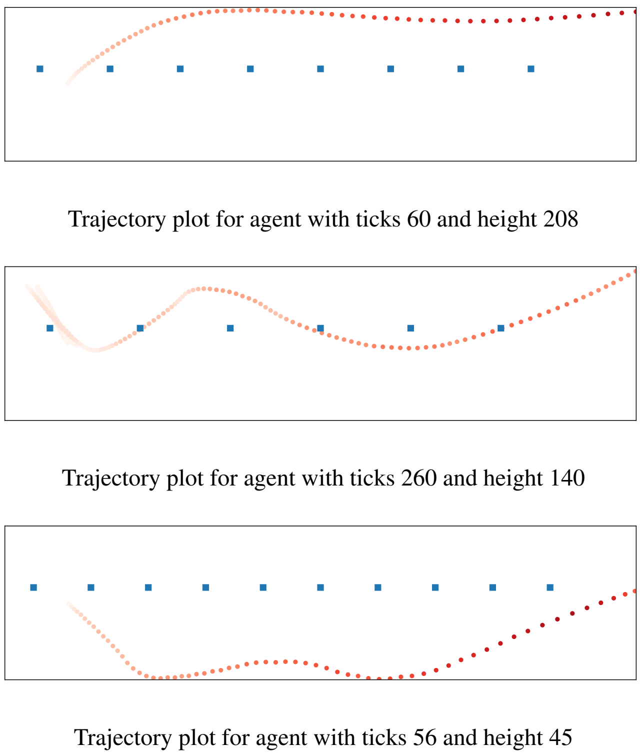

Figure 3: Agent behaviour plots for parameters from specific cells chosen from the behaviour map. Intensity of scatter plot signifies the speed at which the agent was traveling.

Cave Swing with the RHEA agent. For each iteration, the MAP-Elites algorithm picks a set of game parameters, which gets submitted to the game where it gets evaluated. Cave Swing is a fully deterministic game, but the agent relies on random mutations and random initial actions, which make its performance different from run to run. To get a better estimate of the agent's performance the game is played 20 times in our experiments and the collected statistics are averaged. From the collected statistics the MAP-Elites algorithm determines the position of the given game parameters in the map and updates it. We used $10,000$ iterations for the algorithm and a 50 x 50 grid for

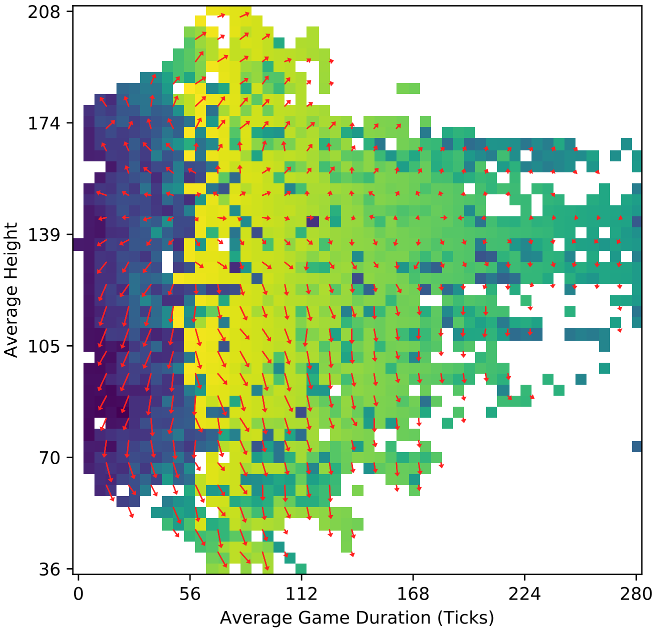

Performance with gravity vectors (2x2 strided average)

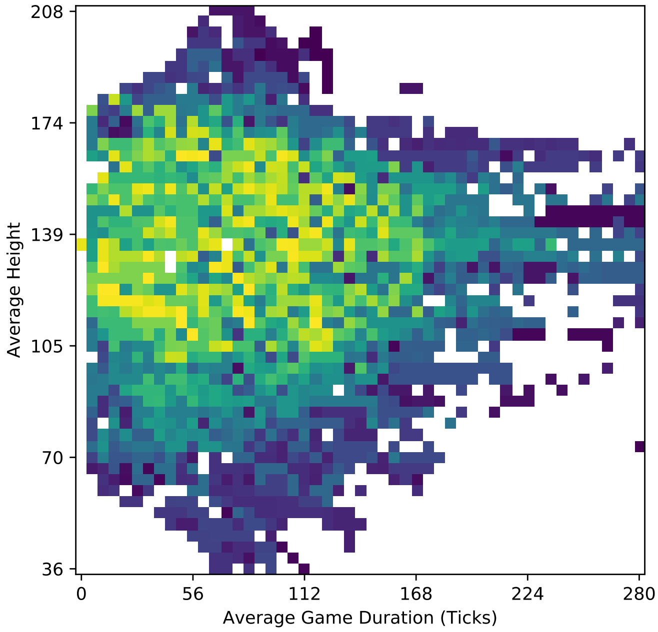

Rope stiffness (Hooke constant)

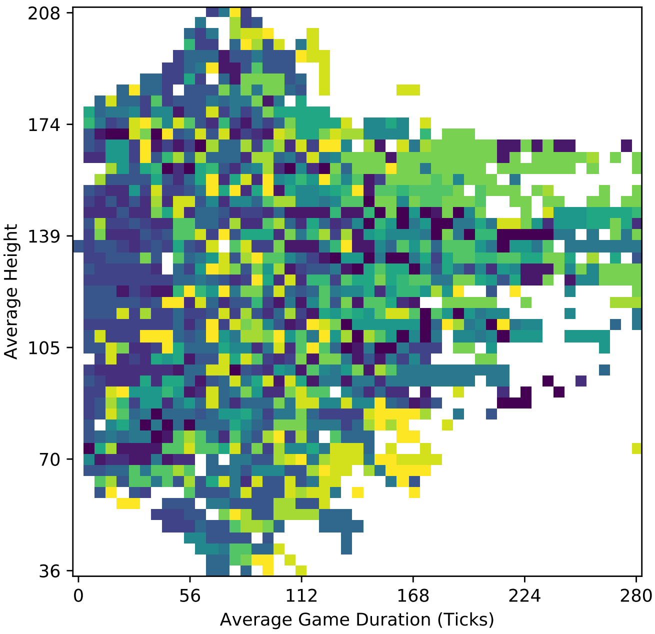

Number of anchors

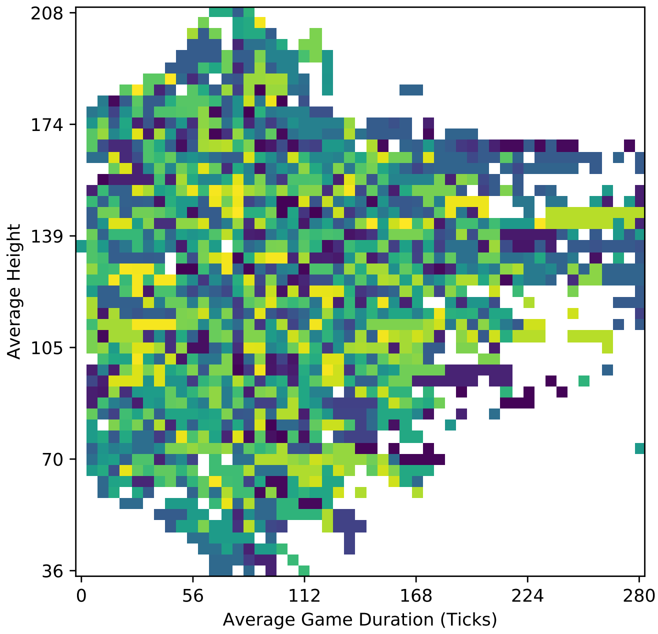

Loss factor

Figure 4: Average game ticks (X axis) against average game height (Y axis). The heatmap represents the performance of the agent with the gravity gradient represented by arrows visualizing the intensity and direction of gravity (3a), the rope stiffness (given by the Hooke constant)(3b), the number of anchors (3c) and loss factor (3d)

the map in this paper.

Each cell in the map produced by the algorithm belonged to a different behaviour. The horizontal and the vertical dimensions were determined by the value of the chosen BC. For demonstrating the effectiveness of our approach, we selected the following behaviours

Table 2: RHEA Parameters. All of these parameters were fixed for the experiment.

RHEA Parameters	Default Value
Sequence Length	200
Number of Evaluations	20
Mutation Rate	10
Shift Buffer	true
Total Random Mutation	false

for further analysis:

1. **Game duration**, which was calculated as the average number of ticks per game run.
2. **Average Height** (vertical position averaged over the run) of the agent in pixels throughout the evaluation runs.

Some parameters of Cave Swing would modify the score function used for evaluation, to avoid this, we fixed those parameters. We decided to also fix the agent's parameters as it would result in additional stochasticity in the evaluations. Please refer to table 1 regarding which parameters were explored.

Table 3: parameters used for the MAP-Elites algorithm

MAP-Elites parameters	
Mutation rate	4
Initial random evaluations	1000
Total evaluations	10000
Grid size	50 x 50

3.2 Data Analysis

The resulting output of $10,000$ iterations of the MAP-Elite algorithm (with 1000 of these being initial random permutations) was computed and the corresponding data collected.

The behaviour maps obtained were explored visually using heatmap style graphs. Each cell of the heatmap contains the parameter set for that agent's evaluation in the game. If a behaviour could be found for that set of parameters, the cell contains a coloured value, which shade was changed depending on the performance value being represented. Those cells where a behaviour could not be found remained in white. This allowed to easily identify regions of high versus low performance within the behaviour space with a contrasting gradient of colours.

Moreover, in order to facilitate the exploration of the found behaviours, a graphical interface was designed by making each of these heatmaps interactive. In this way, each of the cells in the map return its associated parameters by clicking them and the agent playing that particular game variant is shown.

4 Results

The proposed MAP-Elites algorithm successfully found relationships between the parameter space and several points in the behaviour space, with a 57.6% of the cells containing a solution. Note that this percentage depends on the selected ranges of the BC and the number of bins.

Figure 2 shows that the agent's performance distribution depends on both of the behaviour features selected. In this map, the horizontal axis represents the average duration of the game and the vertical axis represents the average height. The brighter the value, the better the performance.

By visually inspecting Figure 2, it can be observed that for short game duration, the performance tended to be very poor. In general, this correspond to runs when the agent did not manage to succeed in the game and the failure penalty was applied. Interestingly, the region determined by medium game duration seems to show the best performance values with longer runs of the game being more detrimental to the agent's performance. This region of medium game duration and high performance values is also related to a larger range of average height values for the agent. Specially for longer games, the variability of the average height reduces considerably.

Figure 3 shows the game simulations corresponding to the sets of parameters associated to 3 different cells in the behaviour map. As it can be observed, the trajectory of the agent in Figure 2b shows an oscillating behaviour as the agent travelled up and down around the anchors and therefore achieved a medium average height. This behaviour led to a long game duration. Figure 2a shows a very different behaviour. The trajectory of the agent was restricted to the top of the map and the duration of the game was relatively short. The agent seemed to have gained momentum, propelling itself towards the top of the map and being able to travel at fast speed. Figure 2c shows a similar behaviour, however the agent propelled itself to the bottom of the map. In order to understand how each of the five explored parameters affected the behaviour space four different heatmaps were built (see Figure 4).

Figure 3a adds to the previously described behaviour-performance relationship the corresponding direction and magnitude of the gravity force. This figure suggests that the average height has a dependency on the gravity direction, as high average height values correspond to gravity values that would have pulled the agent upwards and low values correspond to gravity values pulling the agent downwards. This could explain why the agent propelled itself downwards or upwards in Figures 2a and 2c. Visual inspection of Figure 3a also suggests that long games correspond to points in which the magnitude of the gravity was very low, which would have difficulted the agent's horizontal movement and decreased its height range.

Figures 3b to 3d use the color intensity to represent the value of the parameter selected instead of the performance.

Figure 3b suggests that the stiffness of the rope also played a relevant role in determining the agent's behaviour. Elite solutions with high rope stiffness were only found for low to medium game duration and medium average height, but more diverse solutions were found for more elastic ropes.

Figure 3c displays the relationship between the number of anchors and the behaviour map and Figure 3d presents the relationship between the different values of the loss factor and the behaviour space. It can be observed that the agent's behaviour does not seem to depend on the value of these parameters.

5 Discussion

In general, the obtained results show that it is possible to explore the behaviour space by using MAP-Elites. Visualization of the data seems to be quite useful to explore which parameters are giving place to interesting behaviours and how well a certain agent is able to perform in these conditions. In this way, using MAP-Elites can be useful for game designers to find many variants of a game/level. Instead of just highlighting the best solution, where the agent scores the highest, this method illuminates many different variants. The variants can be visualized, depending on their behaviour features, which would allow game designers to explore all the possible behaviours their game could achieve.

The game selected for this work (Cave Swing) is relatively simple and is governed by non-complex physics. Therefore, not many behaviour emerge from the AI. This game was therefore ideal for a proof-of-concept but the proposed method could be more interesting with more complex games, where the relationship between the game parameters and the agent's behaviour is less clear. For this behaviour characterization only 2 features were used at a time, but MAP-Elites is not limited to having a maximum number of features. As the number of features increase the visualization becomes more complex, but it is still possible to represent it as done by Cully et al [2].

The used agent for the evaluation had a fixed goal determined by the game's score, so the agent did not have any reason to follow a particular play style, that would not result in a high-score. A better evaluation would be to try either more agents or change the objective of the agent. Our method could be used to tune the agent and fix the level, or tune both simultaneously.

6 Conclusion and Future Work

This work presents an exploration of the use of MAP-Elites to a relatively low-parameter game space in order to discover interesting combinations of game parameters. The parameters were evaluated by a RHEA agent by trying to achieve the highest score in the produced game. The visualization of the data collected through the use of the MAP-Elites algorithms shows a high level of detail when comparing the features of the game space, which we chose to optimize for the grid of the chosen behaviour values for both average height and game duration. Furthermore, a plot of the agent playing a game with a chosen set of behaviours (speed/ticks) was also presented to understand how the game parameters affect the agent's movement. In this work we have shown the MAP-Elites can be effectively used to explore the parameters of a game. In this paper we used a simplistic game as a proof-of-concept and a Statistical Forward Planning agent, but any parameterised game with a reasonable agent could be used for this purpose.

As future work, the parameter tuning could be extended to tune more parameters in the game and also the agent to play the game. With MAP-Elites we were able to find parameters for many diverse levels, which could be used to optimize agents to play well a large set of different levels. As Cave Swing is fairly simple, applying MAP-Elites to more complex games would be more interesting, which could result in more complex behaviours.

ACKNOWLEDGEMENTS

This work was supported by the EPSRC Centre for Doctoral Training in Intelligent Games & Game Intelligence (IGGI) [EP/L015846/1].

REFERENCES

[1] A. Alvarez, S. Dahlskog, J. Font, and J. Togelius, 'Empowering quality diversity in dungeon design with interactive constrained map-elites', in *2019 IEEE Conference on Games (CoG)*, pp. 1–8, (Aug 2019).

[2] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret, 'Robots that can adapt like animals', *Nature*, **521**(7553), 503, (2015).

[3] Matthew C. Fontaine, Scott Lee, L. B. Soros, Fernando De Mesentier Silva, Julian Togelius, and Amy K. Hoover, 'Mapping hearthstone deck spaces through map-elites with sliding boundaries', in *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO 19, p. 161169, New York, NY, USA, (2019). Association for Computing Machinery.

[4] D. Gravina, A. Khalifa, A. Liapis, J. Togelius, and G. N. Yannakakis, 'Procedural content generation through quality diversity', in *2019 IEEE Conference on Games (CoG)*, pp. 1–8, (Aug 2019).

[5] Nikolaus Hansen, Sibylle D Müller, and Petros Koumoutsakos, 'Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (cma-es)', *Evolutionary computation*, **11**(1), 1–18, (2003).

[6] Aaron Isaksen, Dan Gopstein, Julian Togelius, and Andy Nealen, 'Discovering unique game variants', in *Computational Creativity and Games Workshop at the 2015 International Conference on Computational Creativity*, (2015).

[7] Aaron Isaksen, Daniel Gopstein, and Andrew Nealen, 'Exploring game space using survival analysis.', in *FDG*, (2015).

[8] Ahmed Khalifa, Scott Lee, Andy Nealen, and Julian Togelius, 'Talakat: Bullet hell generation through constrained map-elites', in *Proceedings of The Genetic and Evolutionary Computation Conference*, pp. 1047–1054, (2018).

[9] Joel Lehman and Kenneth O Stanley, 'Evolving a diversity of virtual creatures through novelty search and local competition', in *Proceedings of the 13th annual conference on Genetic and evolutionary computation*, pp. 211–218. ACM, (2011).

[10] Simon M Lucas, Jialin Liu, and Diego Perez-Liebana, 'The n-tuple bandit evolutionary algorithm for game agent optimisation', in *2018 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1–9. IEEE, (2018).

[11] Jean-Baptiste Mouret and Jeff Clune, 'Illuminating search spaces by mapping elites', *arXiv preprint arXiv:1504.04909*, (2015).

[12] Diego Perez, Spyridon Samothrakis, Simon Lucas, and Philipp Rohlfshagen, 'Rolling horizon evolution versus tree search for navigation in single-player real-time games', in *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*, GECCO '13, pp. 351–358, New York, NY, USA, (2013). ACM.

[13] Justin K Pugh, Lisa B Soros, and Kenneth O Stanley, 'Quality diversity: A new frontier for evolutionary computation', *Frontiers in Robotics and AI*, **3**, 40, (2016).

[14] Julian Togelius, Emil Kastbjerg, David Schedl, and Georgios N Yannakakis, 'What is procedural content generation?: Mario on the borderline', in *Proceedings of the 2nd international workshop on procedural content generation in games*, p. 3. ACM, (2011).

[15] Julian Togelius, Noor Shaker, and Mark J Nelson, 'Procedural content generation in games: A textbook and an overview of current research', *Togelius, N. Shaker, M. NelsonBerlin: Springer*, (2014).

[16] Jonathan Tremblay, Alexander Borodovski, and Clark Verbrugge, 'I can jump! exploring search algorithms for simulating platformer players', in *Tenth Artificial Intelligence and Interactive Digital Entertainment Conference*, (2014).

[17] Vivek R Warriar, Carmen Ugarte, John R Woodward, and Laurissa Tokarchuk, 'Playmapper: Illuminating design spaces of platform games', in *2019 IEEE Conference on Games (CoG)*, pp. 1–4. IEEE, (2019).