Statistical Tree-based Population Seeding for Rolling Horizon EAs in General Video Game Playing

Edgar Galván[1,2], Oxana Gorshkova[1],
Peter Mooney[1,2], Fred Valdez Ameneyro[1,2] and Erik Cuevas[2]

[1] Hamilton Institute, Maynooth University, Ireland
[2] Naturally Inspired Computation Research Group,
Department of Computer Science, Maynooth University, Ireland
{edgar.galvan,peter.mooney}@mu.ie; {oxana.gorshkova.2020,
fred.valdezameneyro.2019}@mumail.ie
[3] Centro Universitario de Ciencias Exactas e Ingenieras (CUCEI),
Universidad de Guadalajara, México
erik.cuevas@cucei.udg.mx

Abstract. Multiple Artificial Intelligence (AI) methods have been proposed over recent years to create controllers to play multiple video games of different nature and complexity without revealing the specific mechanics of each of these games to the AI methods. In recent years, Evolutionary Algorithms (EAs) employing rolling horizon mechanisms have achieved extraordinary results in these type of problems. However, some limitations are present in Rolling Horizon EAs making it a grand challenge of AI. These limitations include the wasteful mechanism of creating a population and evolving it over a fraction of a second to propose an action to be executed by the game agent. Another limitation is to use a scalar value (fitness value) to direct evolutionary search instead of accounting for a mechanism that informs us how a particular agent behaves during the rolling horizon simulation. In this work, we address both of these issues. We introduce the use of a statistical tree that tackles the latter limitation. Furthermore, we tackle the former limitation by employing a mechanism that allows us to seed part of the population using Monte Carlo Tree Search, a method that has dominated multiple General Video Game AI competitions. We show how the proposed novel mechanism, called Statistical Tree-based Population Seeding, achieves better results compared to vanilla Rolling Horizon EAs in a set of 20 games, including 10 stochastic and 10 deterministic games.

1 Introduction

General video game playing aims to produce Artificial Intelligence (AI) agents capable of playing a variety of games rather than playing a specific game. This is

one of the grand challenges of AI. Games are excellent benchmark problems that provide opportunities to test and analyse algorithms in AI. The General Video Game AI (GVGAI) framework [30] is a popular setting that offers a collection of games of different characteristics and this has been used extensively.

Statistical forward planning techniques, such as Monte Carlo Tree Search (MCTS) [25] and Rolling Horizon Evolutionary Algorithms (RHEA) [26], are considered the dominant AI techniques in GVGAI competitions [30]. RHEAs have become real contenders to tree search-based methods that have dominated many GVGAI competitions for many years. Vanilla RHEA models sequences of actions encoded in individuals that are improved through an evolutionary process. The algorithm can be stopped at any time and returns the first action of the fittest individual to be executed by the game agent.

Vanilla RHEA ignores valuable information available when testing the individuals in the evolutionary process. This issue has been addressed with the use of a statistical tree approach [29], the employment of a population seeding [6] and the use of a shift buffer mechanism [7], to mention a few. All these methods yielding better results compared to Vanilla RHEA. However, most of the Vanilla RHEA enhancements waste computational power evaluating meaningless individuals and struggle to maintain a diverse population throughout the evolutionary process, which is important to avoid getting stuck in local optima. The lack of diversity in Vanilla RHEA is likely to happen given the small number of individuals forming a population due to the limited time budget available to the algorithm in which potential solutions can be evolved.

The main contribution of this work is the proposal of novel mechanisms for Vanilla RHEA that exploits the theoretical guarantees of the Upper Confidence Tree (UCT) policy, explained in Section 3 and heavily employed by MCTS. We use this to generate tree-based individuals and seed part of the evolved population saving valuable computational time. Valuable information gathered during the rolling horizon simulation of the EAs is used and is stored in an statistical tree.

The remainder of this paper is organised as follows. Section 2 presents the related work. Section 3 outlines the background in Monte Carlo Tree Search, Evolutionary Algorithms (EAs), Rolling Horizon EAs, and GVGAI. Section 4 discusses our proposed approach. Section 5 presents the experimental setup used and Section 6 presents and discusses the results obtained by all the approaches adopted in this study. Section 7 concludes this paper with some remarks.

2   Related Work

Statistical tree enhancements, inspired by tree search algorithms, were devised by Perez et al. [29]. Statistical tree enhancements targets some of the shortcomings of Vanilla RHEA of not using information knowledge obtained from previous game cycles. The idea of this approach is to build a tree with the statistics of the scores acquired during the game simulations. The tree is built while evaluating each individuals action plan. Each action is added as a node to the tree starting

from the root. The root represents a node whose action was chosen on the previous game step. Each node contains information of how many times it has been visited, the action assigned to it and also an accumulated reward. This is called Open Loop Search and it means that the game states (generative models) are not stored in the nodes. This is important because the game states can become obsolete and provide irrelevant information. The tree is built by evaluating the first action of the individual and this action is added as a child to the root node. All subsequent actions are added in the similar manner. When the last action is applied using a forward model, the game state is evaluated and its result is assigned to the individual's fitness. This value is then back-propagated to the root, incrementing the node visits count and augmenting stored rewards with the fitness value.

Some other interesting works have been proposed to improve Vanilla RHEA. Gaina et al. [7] used four approaches to overcome some limitations present in Vanilla RHEA. The first approach, called Bandit-based mutation, aims at using a bandit system for mutations at both individual and gene level. Interestingly, Gaina et al. used this to revert mutations when no improvement in the value of a given individual was observed. The second approach, named statistical tree, keeps statistical information in game trees associated to each individual evolved. The fitness of each individual is employed to update the statistics stored in each node that has been visited during the evaluation of the individual. The third approach proposed by Gaina et al., called shift buffer, is a simple yet effective technique that maximises the information gained during the limited time to make a decision. Instead of discarding information in every generation, the authors simply add a new random action at the end of each individual in the population. The last approach called roll-outs, inspired by the work of Horn et al. [23] and Monte Carlo simulations, consists of rolling a simulation and then selecting a random selection of actions and game simulations using the forward model. The authors reported competitive results compared to tree-based methods. In particular, shift buffer, in combination with other approaches, such as roll-outs, yielded the most promising results.

More recently, Santos and Hauck [31] proposed two ideas combined into one approach: (i) one-step-look-ahead and (ii) a redundant action avoidance policy with the goal of overcoming some limitations present in Vanilla RHEA. The former idea consists of using the forward model to try every possible action available from the current game state with the goal of selecting the best possible action. This was executed after shift buffer was applied, as proposed by Gaina et al. [7] and explained previously. The latter idea proposed by Santos and Hauck seeks to avoid redundant actions. More specifically an action is randomly selected whenever a set of sequence of actions is present in the individuals forming a population. They tested their approach using 20 different games and compared their results against Vanilla RHEA and Shift Buffer obtaining better results overall.

3 Background

3.1 Monte Carlo Tree Search

Monte Carlo Tree Search (MCTS) is a sampling method for finding *optimal decisions* by performing random samples within the decision space and building a tree according to the partial results. These Monte Carlo methods work by approximating future rewards that can be achieved through random samplings. The evaluation function of MCTS relies directly on the outcomes of simulations. The accuracy of this function increases by adding more simulations, thus the optimal search tree is guaranteed to be found with infinite memory and computation [25]. However, in more realistic scenarios (e.g., limited computing resources available), MCTS can produce very good approximate solutions. MCTS has gained a lot of popularity thanks to its recent success in the board game of computer Go [22], where the space of solutions is 10^{170} and up to 361 legal moves. A problem, that until recently, was highly difficult for AI, and considered much more harder than chess, and where MCTS has obtained excellent results, including achieving dan – master – level at the 9 x 9 board game. This is, perhaps, the reason why MCTS has been used heavily in two-player board games.

However, researchers have begun to apply MCTS to other research problems. For instance, MCTS has been explored in constraint satisfaction (e.g., constraint problems [2]) and energy-based problems [12,13].

The Mechanics Behind MCTS MCTS relies on two key elements: (a) that the true value of an action (in our problems, an action could be turn right, turn left, and so on) can be approximated using simulations; and (b) that these values can be used to adjust the policy towards a best-first strategy. The algorithm, explained later in this section, builds a partial tree, guided by the results of previous explorations of that tree. The algorithm iteratively builds a tree until a condition is reached or satisfied (e.g., number of simulations, time given to perform Monte Carlo simulations), then the search is halted and the best performing action is executed. In the tree, each node represents a state, and directed links to child nodes represent actions leading to subsequent states.

Like many AI techniques, MCTS has several variants. Perhaps, the most accepted steps involved in MCTS are those described in [4] and are outlined as follows: (a) *Selection:* a selection policy is recursively applied to descend through the built tree until an expandable node has been reached. A node is classified as expandable if it represents a non-terminal state, and also, if it has unvisited child nodes; (b) *Expansion:* normally one child is added to expand the tree subject to available actions; (c) *Simulation:* from the new added nodes, a simulation is run to obtain an outcome (e.g., reward value); and (d) *Back-propagation:* the outcome from the simulation step is back-propagated through the selected nodes to update their statistics.

Simulations in MCTS start from the root state (in our case, from the current time when an action for an electric device should be made) and are divided in two stages. When the state is added in the tree, a *tree policy* is used to select the

actions (the selection step is a key element and it is discussed in detail later in this section). Otherwise, a *default policy* is used to roll out simulations to completion. One element that contributed to enhance the efficiency in MCTS is the selection mechanism proposed in [25]. The main idea of this proposed selection mechanism was to design a Monte Carlo search algorithm that had a small probability error if it were stopped prematurely and eventually converged to the optimal solution given enough time. The selection mechanism nicely balances exploration *vs.* exploitation and this will be explained in the following paragraphs.

Upper Confidence Bounds for Trees As indicated previously, MCTS works by approximating "real" values of the actions that may be taken from the current state. This is achieved through building a search or decision tree. The success of MCTS depends heavily on how the tree is built and the selection process plays a fundamental role in this. One particular selection mechanism that has proven to be very reliable is the UCB1 [1]. Formally, UCB1 is defined as:

$$UCB1 = \bar{X}_j + 2K\sqrt{\frac{2 \cdot \ln(n)}{n_j}} \tag{1}$$

where n is the number of times the parent node has been visited, n_j is the number of times child j has been visited, $\bar{X}_j$ is the mean reward of the node j and $K > 0$ is a constant. In case of a tie for selecting a child node, a random selection is normally used [25]. Thus, this selection mechanism is successful due to its emphasis on balancing both exploitation (first part of Eq. 1) and exploration (second part of Eq. 1). Every time a node is visited, the denominator of the exploration part increases resulting in decreasing its overall contribution. If, on the other hand, another child node of the same parent node is visited, the numerator increases, so the exploration values of unvisited children increase. The exploration term in Eq. 1 guarantees that each child node has a selection probability greater than zero, which is essential given the random nature of the play-outs.

3.2 Evolutionary Algorithms

Evolutionary Algorithms (EAs) [5] refer to a set of stochastic optimisation bio-inspired algorithms that use evolutionary principles to build robust adaptive systems. EAs work with a population of μ-*encoded* potential solutions to a particular problem. Each potential solution, commonly known as an individual, represents a point in the search space, within which the optimum solution lies. The population is evolved by means of genetic operators, over a number of generations, to iteratively produce better results to the problem. Each member of the population is evaluated using a fitness function to determine how good or bad the potential solution is for the problem at hand. The fitness value assigned to each individual in the population probabilistically determines how successful the individual will be at propagating (part of) its code to further generations. Better

performing solutions will be assigned higher values (for maximisation problems) or lower values (for minimisation problems). The evolutionary process is repeated until a stopping condition is met. Normally this is until a maximum number of generations has been executed. The population in the last generation is the result of exploring and exploiting the search space over a number of generations. It contains the best evolved potential solutions to the problem and may also, in some cases, represent the globally optimum solution. EAs have been successfully used in a wide variety of problems including games [11,15,19], energy-based problems [10,18,20,17], evolvable hardware [9,16], file type detection [24], classification of imbalanced datasets [8,14,21], to mention a few examples.

3.3 Rolling Horizon Evolutionary Algorithms

Rolling Horizon Evolutionary Algorithms (RHEAs) [28] are inspired by MCTS as explained in Section 3.1. RHEAs employ roll-outs and a generative model (a simulator) allowing the algorithm to see into the future for a short period of time (generally for a fraction of a second). Using this information, the agent (individual) evolves a plan by means of EAs, explained in Section 3.2, executes an action on the problem by firing the first action of its plan. This continues until a stopping condition is met. In the context of GVGAI, RHEAs are executed until the game is over. Details on how RHEA works can be found in greater detail within [28].

3.4 General Video Game Artificial Intelligence

In this work, we use the General Video Game AI (GVGAI) framework [30] to test our proposed approach, called Statistical Tree-based Population Seeding (see Section 4). The GVGAI framework provides a corpus of games, as well as means for creating games, that can be used as testbeds. The corpus consists of two-player games and single-player 2D games of different genres, including shooter, maze, survival, puzzle, and so on. All games have their own rules, scoring and winning criteria. For example, in the game "Aliens", the agent (or controller) wins when it kills all of the aliens and it gets points for every killed alien and an obstacle destroyed. A controller (agent) in the game "Survive Zombies" wins if it is able to stay alive by avoiding being eaten by zombies for a particular amount of time. The avatar in this game has a health bar that decreases every time a zombie catches a player. Collecting honey helps restore the health bar. Furthermore, there are obstacles in the players way that must be avoided whilst also running away from zombies.

Different games have different actions that are legal for a controller to execute. In total, there are six actions in the GVGAI framework: left, right, up, down, nil, use and escape. In the planning track of the framework, controllers have access to the current game state; they can copy it and execute actions on the copy by using a forward model, necessary to perform roll-outs, thus simulating the game. Games can be deterministic or stochastic in nature. In stochastic games, a forward model can produce different future game outcomes from the

current game state with each simulation, contrary to deterministic games. All games can be played on five levels. Each level can be played independently, there is no connection between levels or on how the controller performs, as there is no requirement for a controller to win a level to advance to the next level. A higher level can increase a complexity of the game played and can introduce new challenges to the player, and so on. In our experiments, a controller has a budget of 900 forward model calls (equivalent to 40ms) to simulate the game and decide on an action to fire. This is aligned to the GVGAI competitions where controllers face disqualification if they exceed the time allowed for decision making.

4 Statistical Tree-based Population Seeding

We propose a novel mechanism that partly seeds the population of the RHEA using the UCB1 policy on a statistical tree. The statistical tree introduced in [29] is a RHEA improvement that constructs a tree with each state visited when individuals in the population are evaluated. Each state is modelled as a node and is linked to the previous state by an action. Each node stores two values: the times it has been visited and the cumulative reward. After each individual is evaluated the rewards are back-propagated from the corresponding leaf node to the root. In this way, the statistical tree keeps track of the rewards found by the RHEA's evolutionary process. RHEA returns the first action of the fittest individual as the action to be performed by the controller. In the statistical tree, the node linked to the root with the performed action is transformed into the new root node. Any nodes not part of the sub-tree beginning from the new root node are pruned.

When the RHEA algorithm is required again to return a new action, the initial population is seeded with the statistical tree. The first individual is generated following the path in the tree leading to the best reward. The rest of the population is generated following the UCB1 policy from the root node. This action is repeated $m-1$ times, where m is the population size, to generate the remaining $m-1$ individuals. It is important to evaluate each new individual before generating a new one to update the statistical tree. The available computational power to make a decision is limited (900 forward model calls) and as a result of this, the size of the population must be small. This can hinder the population's diversity with the consequence of individuals getting stuck in local optima. To alleviate this, the proposed algorithm attempts to inject a new individual after each generation. The individual is generated following the UCB1 policy on the statistical tree, and replaces the worst individual in the current population only if the newly tree-based generated individual performs better. Regardless of it being added to the population or discarded, the newly generated individual contributes to the construction and statistics of the statistical tree. This step helps to keep interesting individuals emerging in the within population while guiding the search towards promising regions.

5 Experimental Setting

5.1 Controllers

We compare our proposed approached, dubbed (a) Statistical Tree-based Population Seeding, against (b) Vanilla RHEA [28]. Furthermore, we also use (c) the Statistical Tree approach [29], and (d) the Shift Buffer mechanism [7]. The latter three approaches (b – d), explained previously, heavily inspired our proposed approach and show how by borrowing elements from these and integrating new elements such as partially seeding a population, it is possible to get, overall, better results compared to these three approaches.

5.2 Games

Table 1. List of the 20 games used in our studies, including their names with their corresponding indexes divided in two categories: stochastic (G1 – G10) and deterministic (G11 – G20).

Index	Stochastic	Index	Deterministic
G1	Aliens	G11	Bait
G2	Butterflies	G12	Camel Race
G3	Chopper	G13	Chase
G4	Crossfire	G14	Escape
G5	Dig Dug	G15	Hungry Birds
G6	Infection	G16	Lemmings
G7	Intersection	G17	Missile Command
G8	Roguelike	G18	Modality
G9	Sea Quest	G19	Plaque Attack
G10	Survive Zombies	G20	Wait for Breakfast

Each of the four algorithms (controllers) used in this study were run on 20 games of the GVGAI corpus, on all 5 levels, 40 times each. Each algorithm had to play 4000 times, resulting in 16000 independent runs in total. To perform this large number of simulations, we used a supercomputer with 336 nodes where each node has 2x 20-core 2.4 GHz Intel Xeon Gold 6148 (Skylake) processors, 192 GiB of RAM. Furthermore, to provide a sound comparison of results, we selected ten stochastic games and ten deterministic games. The selection of these is based on performing a uniform sampling from the list of games in Table 1. The selection of these games is the result of analysing dozens of games based on the performance of MCTS expressed in terms of the winning rate for each of the 62 games used in Nelson [27] and based on the similarity of game features among 49 games studied in Bontager et al. [3].

5.3 EAs Configuration and Parameters

The experiments were conducted using a generational approach. The population size was set to 10 and the length of the individual (action plan) set to 14.

All four algorithms used in this study used uniform crossover and mutation operators. Tournament selection (size 3) was used as a selection mechanism in our algorithms. Furthermore, we used elitism (1 individual) to ensure that the best individual was carried over to the next generation. Our algorithms were not limited in the number of generations. Instead, we set a budget of 900 forward model calls every time an action was fired by any of the algorithms used in this study. We use the score as the fitness function in each of the four controllers used in this study.

6 Results and Analysis

All results were converted into Formula-1 (F1) ranking scheme. That is, for each game, controllers are compared by their mean win rate. The controller with the highest average win rate gets 25 points, the second best controller is granted 18 points, and the remaining two controllers get 15 and 12 points respectively. If two or more controllers achieved the same number of victories then they are compared by their scores. The overall results are displayed in Table 2. A breakdown of the results for each game is presented in Table 3 for stochastic games and Table 4 for deterministic games. Vanilla RHEA is ranked 4th and is the worst of all the controllers used in this work. Our proposed Statistic Tree-based Seeding Population algorithm ranked first in 10 games as well as first overall. It is closely followed by the Stat. Tree variant. Both methods were tied in F1 scores in deterministic games, but Stat. Tree-based Population Seeding had a better performance in stochastic games, showing that the proposed method has a potentially greater impact in the latter types of games.

Table 2. F1 total points achieved by each controller used in this study.

Rank	RHEA Approach	F1 Points	Avg. Points
1st	Stat. Tree-based Population Seeding	412	49.67%
2nd	Stat. Tree	382	48.37%
3rd	Shift Buffer	315	45.92%
4th	Vanilla RHEA	291	46.77%

Table 3. F1 ranking of the controllers in the 10 <u>stochastic</u> games. The darker the cell, the better the result.

RHEA Methods	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10	Total
Stat. Tree-based Population Seeding	15	15	25	25	15	15	25	18	25	25	203
Stat. Tree	18	25	12	15	18	25	15	15	15	15	173
Shift Buffer	25	18	18	12	25	18	12	25	12	12	177
Vanilla RHEA	12	12	15	18	12	12	18	12	18	18	147

Table 4. F1 ranking of the controllers in the 10 <u>deterministic</u> games. The darker the cell, the better the result.

RHEA Methods	G11	G12	G13	G14	G15	G16	G17	G18	G19	G20	Total
Stat. Tree-based Population Seeding	25	25	18	25	18	12	18	18	25	25	209
Stat. Tree	18	18	25	12	25	25	25	25	18	18	209
Shift Buffer	12	15	12	15	12	18	15	12	12	15	138
Vanilla RHEA	15	12	15	18	15	15	12	15	15	12	144

The Mann-Whitney U test at 5% significance level was used to compare controllers with each other to determine if the difference in scores and wins is statistically significant. Table 5 shows the count of stochastic and deterministic games where the algorithm in the row significantly outperforms the algorithm in the column in win rates. Table 6 is similar to this but now considers those significant scores between any two algorithms. Table 7, listed in the Appendix, shows the details of wins for each game attained by each of the four methods used in this work. The average scores are not shown due to space constraints. Based on these tables, it can be seen that both statistical tree-based RHEA versions achieved the highest number of games where they outperformed the other two algorithms. Our proposed approach, Statistical Tree-based Population Seeding, had a significantly superior score in 13 games in total, and significantly superior win rates in 9 games, compared to the other three algorithms (controllers). The three games in which the Statistical Tree-based Population Seeding significantly outperformed the Vanilla RHEA in win rates are *Wait for breakfast*, *Missile command* and *Chopper*. Regarding scores, our approach was significantly better than Vanilla RHEA in *Wait for breakfast*, *Roguelike*, *Crossfire*, *Chopper* and *Intersection*. *Missile command* is the only game in which both the Statistical Tree-based Population Seeding and the Stat. Tree controllers significantly improve over the Vanilla RHEA in both scores and win rates. *Dig-dug* is the only game in which the Stat. Tree-based Population Seeding performed significantly worse than the Vanilla RHEA regarding scores. This is not the same situation for win rates. Stat. Tree-based Population Seeding also struggled with the game *Lemmings*. One possible reason is that exploration in this game can be extremely costly and rewards are sparse. The exploratory nature of Stat. Tree-based Population Seeding ends up actually impacting it negatively. In general, Stat. Tree-based Population Seeding is a promising approach that improves over the Vanilla RHEA, being clearly superior to the Shift Buffer and more robust than the Stat. Tree in many cases.

Table 5. Count of games where the controller in the row has significantly <u>better win rates</u> than the controller in the column.

	Stat. Tree-based Seeding	Stat. Tree	Shift Buffer	Vanilla RHEA
Stat. Tree-based Seeding	–	2	4	3
Stat. Tree	2	–	3	3
Shift Buffer	0	0	–	0
Vanilla RHEA	0	2	1	–

Table 6. Count of games where the controller in the row has significantly <u>better scores</u> than the controller in the column.

	Stat. Tree-based Seeding	Stat. Tree	Shift Buffer	Vanilla RHEA
Stat. Tree-based Seeding	–	3	4	6
Stat. Tree	4	–	4	5
Shift Buffer	2	0	–	5
Vanilla RHEA	1	1	0	–

7 Conclusions

In this paper we proposed a novel mechanism named Statistical Tree-based Population Seeding to be employed in Rolling Horizon Evolutionary Algorithms (RHEAs) with the goal of overcoming some limitations of the vanilla RHEA. This mechanism uses a Statistical Tree that keeps track of information gathered by the individuals contained in the population. Each of these individuals contribute to the information stored in this Statistical Tree every time a potential solution is rolled in a game simulation. We then use this tree to generate an individual to be seeded in the population in order to have diversity in the population as well as speeding up the process of evolution required in the limited time (a fraction of a second) to 'fire' an action to be executed by the game agent. We demonstrated how our approach is significantly superior to Vanilla RHEA as well as being better, in average, to the other two controllers, named Statistical Tree and Shift Buffer, used in this study.

Acknowledgements

References

1. P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Mach. Learn.*, 47(23):235256, May 2002.
2. S. Baba, Y. Joe, A. Iwasaki, and M. Yokoo. Real-time solving of quantified csps based on monte-carlo game tree search. In *IJCAI*, pages 655–661, 2011.
3. P. Bontrager, A. Khalifa, A. Mendes, and J. Togelius. Matching games and algorithms for general video game playing. In *AIIDE*, 2016.
4. C. Darken and M. Mateas, editors. *Proceedings of the Fourth Artificial Intelligence and Interactive Digital Entertainment Conference, October 22-24, 2008, Stanford, California, USA*. The AAAI Press, 2008.
5. A. E. Eiben and J. E. Smith. *Introduction to Evolutionary Computing*. Springer Verlag, 2003.

6. R. D. Gaina, S. M. Lucas, and D. P. Liebana. Population Seeding Techniques for Rolling Horizon Evolution in General Video Game Playing. In *Proceedings of the Congress on Evolutionary Computation*, pages 1956–1963, June 2017.

7. R. D. Gaina, S. M. Lucas, and D. Perez-Liebana. Rolling horizon evolution enhancements in general video game playing. In *2017 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 88–95, 2017.

8. E. Galván and M. Schoenauer. Promoting semantic diversity in multi-objective genetic programming. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO 19, page 10211029, New York, NY, USA, 2019. Association for Computing Machinery.

9. E. Galvan-Lopez. Efficient graph-based genetic programming representation with multiple outputs. *International Journal of Automation and Computing*, 5(1):81–89, Jan 2008.

10. E. Galván-López, T. Curran, J. McDermott, and P. Carroll. Design of an autonomous intelligent demand-side management system using stochastic optimisation evolutionary algorithms. *Neurocomputing*, 170:270–285, 2015.

11. E. Galván-López, D. Fagan, E. Murphy, J. Swafford, A. Agapitos, M. O'Neill, and A. Brabazon. Comparing the performance of the evolvable πgrammatical evolution genotype-phenotype map to grammatical evolution in the dynamic ms. pac-man environment. In *Evolutionary Computation (CEC), 2010 IEEE Congress on*, pages 1–8, July 2010.

12. E. Galván-López, C. Harris, L. Trujillo, K. Rodriguez-Vazquez, S. Clarke, and V. Cahill. Autonomous demand-side management system based on monte carlo tree search. In *2014 IEEE International Energy Conference (ENERGYCON)*, pages 1263–1270, May 2014.

13. E. Galván-López, R. Li, C. Patsakis, S. Clarke, and V. Cahill. Heuristic-based multi-agent monte carlo tree search. In N. G. Bourbakis, G. A. Tsihrintzis, and M. Virvou, editors, *5th International Conference on Information, Intelligence, Systems and Applications, IISA 2014, Chania, Crete, Greece, July 7-9, 2014*, pages 177–182. IEEE, 2014.

14. E. Galván-López, E. Mezura-Montes, O. Ait ElHara, and M. Schoenauer. On the use of semantics in multi-objective genetic programming. In J. Handl et al., editors, *Parallel Problem Solving from Nature – PPSN XIV: 14th International Conference, Edinburgh, UK, September 17-21, 2016, Proceedings*, pages 353–363. Springer, 2016.

15. E. Galván-López and M. O'Neill. On the effects of locality in a permutation problem: The sudoku puzzle. In P. L. Lanzi, editor, *Proceedings of the 2009 IEEE Symposium on Computational Intelligence and Games, CIG 2009, Milano, Italy, 7-10 September, 2009*, pages 80–87. IEEE, 2009.

16. E. Galván-López, R. Poli, and C. A. C. Coello. Reusing code in genetic programming. In M. Keijzer, U. O'Reilly, S. M. Lucas, E. Costa, and T. Soule, editors, *Genetic Programming, 7th European Conference, EuroGP2004, Coimbra, Portugal, April 5-7, 2004, Proceedings*, volume 3003 of *Lecture Notes in Computer Science*, pages 359–368. Springer, 2004.

17. E. Galván-López, M. Schoenauer, and C. Patsakis. Design of an autonomous intelligent demand-side management system by using electric vehicles as mobile energy storage units by means of evolutionary algorithms. In *2015 7th International Joint Conference on Computational Intelligence (IJCCI)*, volume 1, pages 106–115, 2015.

18. E. Galván-López, M. Schoenauer, C. Patsakis, and L. Trujillo. Demand-side management: Optimising through differential evolution plug-in electric vehicles to partially fulfil load demand. In J. J. Merelo, A. Rosa, J. M. Cadenas, A. D. Correia, K. Madani, A. Ruano, and J. Filipe, editors, *Computational Intelligence*, pages 155–174, Cham, 2017. Springer International Publishing.

19. E. Galván-López, J. M. Swafford, M. O'Neill, and A. Brabazon. Evolving a ms. pacman controller using grammatical evolution. In C. D. Chio, S. Cagnoni, C. Cotta, M. Ebner, A. Ekárt, A. Esparcia-Alcázar, C. K. Goh, J. J. M. Guervós, F. Neri, M. Preuss, J. Togelius, and G. N. Yannakakis, editors, *EvoApplications (1)*, volume 6024 of *Lecture Notes in Computer Science*, pages 161–170. Springer, 2010.

20. E. Galván-López, A. Taylor, S. Clarke, and V. Cahill. Design of an automatic demand-side management system based on evolutionary algorithms. In *Proceedings of the 29th Annual Symposium on Applied Computing, SAC '14*, pages 525 – 530, Gyeongju, Korea, 26-28 Mar. 2014. ACM.

21. E. Galván-López, L. Vázquez-Mendoza, and L. Trujillo. Stochastic semantic-based multi-objective genetic programming optimisation for classification of imbalanced data. In O. Pichardo-Lagunas and S. Miranda-Jiménez, editors, *Advances in Soft Computing*, pages 261–272, Cham, 2017. Springer International Publishing.

22. S. Gelly, Y. Wang, R. Munos, and O. Teytaud. Modification of UCT with patterns in Monte-Carlo Go. Technical Report 6062, INRIA, France, Nov. 2006.

23. H. Horn, V. Volz, D. Prez-Libana, and M. Preuss. Mcts/ea hybrid gvgai players and game difficulty estimation. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–8, 2016.

24. A. Kattan, E. Galván-López, R. Poli, and M. O'Neill. Gp-fileprints: File types detection using genetic programming. In *Genetic Programming, 13th European Conference, EuroGP 2010, Istanbul, Turkey, April 7-9, 2010. Proceedings*, pages 134–145, 2010.

25. L. Kocsis and C. Szepesvári. Bandit based monte-carlo planning. In *ECML*, pages 282–293, 2006.

26. D. P. Liebana, S. Samothrakis, S. M. Lucas, and P. Rohlfshagen. Rolling horizon evolution versus tree search for navigation in single-player real-time games. In C. Blum and E. Alba, editors, *Genetic and Evolutionary Computation Conference, GECCO '13, Amsterdam, The Netherlands, July 6-10, 2013*, pages 351–358. ACM, 2013.

27. M. J. Nelson. Investigating vanilla mcts scaling on the gvg-ai game corpus. In *2016 IEEE Conference on Computational Intelligence and Games (CIG)*, pages 1–7, 2016.

28. D. Perez, S. Samothrakis, S. Lucas, and P. 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, page 351358, New York, NY, USA, 2013. Association for Computing Machinery.

29. D. Perez Liebana, J. Dieskau, M. Hunermund, S. Mostaghim, and S. Lucas. Open loop search for general video game playing. In *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, GECCO 15, page 337344, New York, NY, USA, 2015. Association for Computing Machinery.

30. D. Perez-Liebana, J. Liu, A. Khalifa, R. D. Gaina, J. Togelius, and S. M. Lucas. General video game ai: A multitrack framework for evaluating agents, games, and content generation algorithms. *IEEE Transactions on Games*, 11(3):195–214, 2019.

31. B. Santos, H. Bernardino, and E. Hauck. An improved rolling horizon evolution algorithm with shift buffer for general game playing. In *2018 17th Brazilian Sym-*

posium on Computer Games and Digital Entertainment (SBGames), pages 31–316, Los Alamitos, CA, USA, nov 2018. IEEE Computer Society.

Appendix

Table 7. Mean win rate achieved by each controller in each of the 20 games used in this study.

Game name	Mean wins (± standard deviation)			
	Stat. Tree-based Seeding	Shift Buffer	Stat. Tree	Vanilla RHEA
aliens	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)
bait	0.11 (±0.31)	0.04 (±0.20)	0.09 (±0.28)	0.07 (±0.25)
butterflies	0.92 (±0.27)	0.92 (±0.26)	0.98 (±0.14)	0.88 (±0.32)
camelRace	0.05 (±0.21)	0.03 (±0.18)	0.04 (±0.20)	0.03 (±0.17)
chase	0.03 (±0.18)	0.03 (0.17)	0.05 (±0.21)	0.035 (±0.18)
chopper	1.00 (± 0.00)	0.99 (0.09)	0.975 (±0.15)	0.98 (±0.14)
crossfire	0.08 (± 0.27)	0.05 (0.21)	0.05 (±0.21)	0.06 (±0.23)
digdug	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)
escape	0.38 (± 0.48)	0.2 (±0.40)	0.175 (±0.37)	0.315 (±0.46)
hungrybirds	0.30 (± 0.45)	0.27 (±0.44)	0.33 (±0.47)	0.285 (±0.45)
infection	0.96 (± 0.18)	0.985 (±0.12)	0.995 (±0.07)	0.955 (±0.20)
intersection	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)	1.00 (± 0.00)
lemmings	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)
missilecommand	0.72 (± 0.44)	0.64 (±0.48)	0.74 (±0.43)	0.545 (±0.49)
modality	0.26 (± 0.44)	0.255 (±0.43)	0.265 (±0.44)	0.26 (±0.43)
plaqueattack	0.97 (± 0.17)	0.92 (±0.27)	0.965 (±0.18)	0.945 (±0.22)
roguelike	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)	0.00 (± 0.00)
seaquest	0.88 (± 0.32)	0.75 (±0.43)	0.87 (±0.33)	0.88 (±0.32)
survivezombies	0.50 (± 0.49)	0.415 (±0.49)	0.455 (±0.49)	0.475 (±0.49)
waitforbreakfast	0.75 (± 0.43)	0.67 (±0.47)	0.69 (±0.46)	0.64 (±0.48)