

Tree Search vs Optimization Approaches for Map Generation

Debosmita Bhaumik*, Ahmed Khalifa†, Michael Cerny Green‡ and Julian Togelius§

Tandon's School of Engineering

New York University

New York, USA

*debosmita.bhaumik01@gmail.com,†ahmed@akhalifa.com‡mcgreentn@gmail.com,§julian@togelius.com

Abstract—Search-based procedural content generation uses stochastic global optimization algorithms to search spaces of game content. However, it has been found that tree search can be competitive with evolution on certain optimization problems. We investigate the applicability of several tree search methods to map generation, and compare them systematically with several optimization algorithms, including evolutionary algorithms. For purposes of comparison, we use a simplified map generation problem where only passable and impassable tiles exist, three different map representations, and a set of objectives that are representative of those commonly found in actual level generation problem. While the results suggest that evolutionary algorithms produce good maps faster, several tree search methods can perform very well given sufficient time, and there are interesting differences in the character of the generated maps depending on the algorithm chosen, even for the same representation and objective.

Index Terms—Procedural Content Generation, Level Generation, Map Layout, Tree Search, Optimization Algorithms

I. INTRODUCTION

Generating levels for games is a research problem with broad relevance across most game genres and many domains outside of games. Video games, from shooters to role-playing games to puzzle games, need level generation in order to create larger and more replayable games, adapt games to players, simplify game development, and enable certain kinds of aesthetics. Domains such as architecture, urban planning, military simulation and logistics need scenario and environment generation for similar reason, and these problems are often very similar to game level generation. In reinforcement learning, level generation allows for creating variable environments which helps with generalization [1]. For these reasons, the past decade has seen considerable interest in research on level generation and other forms of procedural content generation (PCG) from both academia and industry [2].

One particular approach to the generation of levels as well as other types of game content is to use evolutionary algorithms or similar global stochastic optimization algorithms to search for good levels. This approach, called search-based PCG, requires that the levels are represented in such a way that the level space can be efficiently searched, and that there is a fitness function which can reliably approximate the quality of the level [3].

As an alternative to using evolutionary methods for PCG, it has been suggested to use tree search methods such as Monte

Carlo Tree Search (MCTS) [4]. For example, Browne showed that MCTS could be used to effectively search for simple polygon shapes and the Pentominoes¹ puzzle domain [5]. While it seems that both stochastic optimization and tree search can be used for level generation (and many related generative tasks), tree search methods are currently severely understudied. Given the very different ways in which these algorithm types search a space of artifacts, it stands to reason that they should differ sharply in performance depending on the objective and representation.

This paper presents what we consider the first systematic comparison of tree search methods being used for level generation and to compare its performance with evolutionary methods. Section II reviews research with Tree Search and Optimization, the two approaches we experiment with. We go over each technique in detail in Sections III and IV. For ease of analysis we consider a simplified level generation problem, namely generating two-dimensional maps where each cell can be either passable or impassable. Section V describes the 3 different representations we explored with each of the algorithms. In Section VI, we define a series of fitness or objective functions used by all algorithms for all representations which mirror commonly occurring considerations when generating levels for common game genres such as strategy games or shooters, or environments for simulations. Our experiment results in Section VII illuminate not only systematic differences in performance between algorithms of different families, but also which methods and representations work best for particular types of objectives.

II. BACKGROUND

Procedural content generation (PCG) is defined as the automatic generation of desirable artifacts within games, be it game levels, characters, quests and storylines, game elements like trees and rocks, or even entire games themselves [2]. Search-based PCG is a subset of PCG methods that involves search strategies such as tree search and optimization algorithms [3] to generate the content. This section describes previous research in the areas of tree search and evolution as well as procedural content which can be generated using these methods.

¹<http://www.ericharshbarger.org/pentominoes/>

A. Tree Search

Tree search algorithms try to find solutions by starting at a root node and expanding child nodes in a systematic way. Popular techniques include Breadth-First Search, Depth-Search First, A*, and Monte Carlo Tree Search (MCTS) [6]. Tree search agents are commonly used as game-playing agents, like in Super Mario Bros (Nintendo 1985) [7], Go [8], [9], and general video games [10] among many others.

In the area of PCG, Browne first explored this concept by using a variant of the Upper Confidence Bound for Trees equation (UCT) called *Upper Confidence Bounds for Graphs* (UCG) [5] to develop biominoes, simple polyomino² shapes and the Pentominoes puzzle domain. Summerville et al generated levels for Super Mario Bros (Nintendo 1985) using Markov Chains where the exploration was guided using Monte Carlo Tree Search [12]. Kartal et al used MCTS to generate stories, taking advantage of MCTS' ability to successfully navigate the large search spaces associated with possible character actions and reactions within narratives [13]. Kartal et al also used MCTS to generate *Sokoban* (Imabayashi 1981) levels [14]. At each node in the MCTS tree, the level generator is given choices to take to modify the level, such as deleting/adding objects and moving an agent around within the level to simulate gameplay. Graves experimented using MCTS to generate Angry Birds (Rovio Entertainment 2009) levels [15]. At each node in the tree, the level generator can place/remove structures/pigs or do nothing at all.

B. Optimization

Optimization algorithms are procedures which iterate until optimal solution is found or resources are exhausted. When used in a game-playing fashion (albeit uncommonly), these techniques are usually formatted to create a sequence of actions terminating at a rolling horizon. One such example is the Rolling Horizon Evolution Algorithm (RHEA) [16], which proved to be competitive against tree search agents in the General Video Game Artificial Intelligence Competition [10]. Justesen et al experimented with another technique known as online evolutionary planning in the game Hero Academy (Robot Entertainment 2012), which is a multi-action, turn-based, adversarial game [17].

Optimization search techniques are often popular choices for PCG because of how easy it is to frame the PCG as an single- or multi-point optimization problem, where the fitness functions/objectives can be cleanly mapped to game elements like difficulty, time, physical space, level variety, etc. Ashlock et al. did this several ways, such as optimized puzzle generation for different difficulties [18], [19], or stylized cellular automata evolution for cave generation [20]. McGuinness et al. created a micro-macro level generation process [21], using a wide variety of fitness functions based on level elements. In addition to evolving level elements in GVGAI [22], PuzzleScript [23], bullet-hell games [24], and Super Mario Bros [25], [26] Khalifa et al. [27] offers a literature review of

search based level generation within puzzle games. Shaker et al. [28] evolved levels for *Cut the Rope* (ZeptoLab 2010) using constraint evolutionary search where the fitness measures the playability using playable agents.

III. TREE SEARCH ALGORITHMS

In our experiments, we used the uninformed Breadth First Search and Depth First Search algorithms as well as the informed search A* and Monte Carlo Tree Search algorithms, which are all described below.

A. Breadth First Search

A simple uninformed search algorithm, breadth first search [6] expands by fully exploring a tree level before going deeper by using a queuing system.

B. Depth First Search

Depth-first [6] search always expands one of the nodes at the deepest level of the tree by using a stack system. Only when the search hits a dead end or terminal node the search goes back and expand nodes at shallower levels.

C. A* Search

A* [6] is an informed search algorithm which uses a heuristic function and a priority queue to select the most promising nodes in the search tree first.

D. Monte Carlo Tree Search

MCTS [4] is a stochastic tree-search based algorithm that creates asymmetric trees by expanding the more promising branches of the search space using random sampling. It consists of four phases in its iterative process: *selection*, *expansion*, *simulation*, and *backpropagation*. During *selection*, the algorithm computes which node to be selected using a *selection policy*, which defines how the algorithm will *select* between exploring less visited tree nodes or exploiting nodes with higher estimated reward values, a popular policy being UCT [29]. During *expansion*, a new node is added to the tree as a child of the last selected node which is not fully expanded. The newly created child node is *simulated* forward until it reaches either some terminal state (a win or a loss) or some pre-defined threshold (i.e 500 moves into the future). The node's reward value is calculated from the simulation phase's final state and *backpropogated* through the values of the any parent nodes, from the newly created node to the tree root. The algorithm runs in an iterative fashion, and the updated node values from the last iteration defines how to guide the search in the next iteration.

Our experiments use the same UCT function mentioned above. However, rather than make C a constant, we calculate C at every tree level using the average score of its successors. To calculate the C constant, we calculate the average score for each of the existing successors of that node then calculate the difference between the maximum average scoring successor and the minimum average scoring successor. If both averages are the same, we add a small value to favor exploration (make sure C is not equal to zero).

²Orthogonally connected sets of squares [11]

IV. OPTIMIZATION ALGORITHMS

In our experiments, we explored hill climbing, simulated annealing, simple evolutionary strategy, and genetic algorithm, which are described below.

A. Hill Climbing

The hill climber algorithm [6] is a single-point optimization algorithm that starts with a random solution and keeps improving the solution until a local optimal solution is found.

At each iteration, the algorithm takes the current solution and finds all possible neighboring solutions in the search space. It then calculates the fitness of these solutions and compares them. If any neighboring state is better than the current one, then the current state is replaced with the neighboring state, otherwise the current state remains.

B. Simulated Annealing

Simulated Annealing [6] is a single-point global optimization algorithm that tries to find a global optimum in the presence of several local optima.

Like the hill climbing algorithm, the process starts with a randomized initial state. Within each iteration, the current selected state's score is calculated using a heuristic function. A neighbor of the current state is then randomly generated, and its score is calculated. If the neighbor state is better than the current state, it is made the current state. Otherwise the new state is accepted with probability less than 1. The probability is calculated by

$$P = \exp(-d/T)$$

where d the absolute difference between the current states score and the new states score and T is temperature. As the value of T is high at the beginning, the probability of accepting the poorer solution is higher at the start of the algorithm. After getting the new current state the same process continues until it finds optimal solution. Before each iteration the temperature is calculated

$$T = T * c$$

where T = temperature, c = cooling rate.

C. Evolutionary Strategy

Evolution Strategy [30] is a nature inspired multi-point optimization algorithm. It applies selection and mutation operators to a population, that contains solutions, to evolve better and better solutions.

$$(\mu/\rho + \lambda) - ES$$

The process begins by initializing a random population of size μ individuals and calculates the fitness the entire population using a fitness function. λ worst individuals are removed from the population. Then selection and mutation operators are used on the population to evolve λ new individuals to fill the places of eliminated individuals and keep the population size same between generations. The ρ best individuals (selected via rank selection) from the old population become parents to new

individuals in the new population. Each time, one individual from the parent collection is randomly selected. Then mutation operator is applied on the selected individual with a predefined chance to mutate and create new offspring, and that offspring is added to the new population. Once the new population is completed, fitness of the new population is computed, and this process continues until desired solution is found.

D. Genetic Algorithm

A Genetic Algorithm [30] is multi-point optimization technique inspired by the Darwinian principle of evolution. Like evolutionary strategy, it uses nature inspired operators like mutation and selection (with the addition of a crossover operation) to generate high quality solutions. Starting with a random population, it selects individuals based on fitness for reproduction.

Within each generative iteration, the population has its fitness calculated. The best $x\%$ individuals from the current population are immediately inserted into the next generation without any crossover or mutation, a process known as *elitism*. For the rest of the population, two parents are selected from the current population based on their fitness using a rank selection algorithm. Then a single point crossover operator is applied with $y\%$ probability for parents to create an offspring. A $z\%$ probability mutation operator is applied on this offspring, which is then inserted into the new population. The reproductive process goes on until the new population is fully created.

V. PROBLEM REPRESENTATION

Our experiments were done using a binary map, where 0 represents empty space and 1 represents a blocked area (or wall). Different types of map representations were used, defined by what percentage of empty and blocked tiles they contained. Starting point maps would be generated randomly, given the percentage of empty space a map would contain. The maps are used as the root node of the search algorithms. For all algorithms, the two possible modifications are to flip a tile (if it is a 0 change it to 1 or if it is a 1 change it to 0) or keep the tile the same. Below the three different operations that could be used by any of the algorithms to modify the map are described.

A. Narrow Representation

The narrow representation approach is defined as changing one specified tile at a time. In this approach, the tiles that the algorithm can modify are randomly ordered, and the algorithm can only modify the map in that particular order. For tree search, this means that each level of the tree represents a different tile being modified. For optimization, this means that a chromosome is represented as a sequence of tile modifications. For both algorithm types, tiles can only be modified (or not) once.

B. Wide Representation

The wide representation allows for more freedom in tile modification order. In this representation, the algorithm itself can decide exactly which tiles to modify in any order. For tree search, this means that a node represents a version of the map where a specific tile was flipped. No node may be created that flips a tile which has already been flipped by one of its ancestor nodes. For optimization, this means that a chromosome is essentially a version of the map itself.

C. Turtle Graphics Representation

The turtle graphics representation draws parallels to the Turtle Graphics module in the Logo programming language. In this representation, algorithms are given a random initial position within the map. They are allowed to choose how to modify this tile, and then are given the choice to move in any of the neighboring tiles in the four cardinal directions (unless a direction would take them “out-of-bounds”) and modify the new tile, and the process repeats. For tree search, this means that a node in the tree represents a pair of directional movement and a corresponding modification decision. For optimization, this means that a chromosome is a sequence of pairs containing a directional movement and the corresponding modification decision.

VI. FITNESS/HEURISTIC FUNCTIONS

A heuristic function is a function that ranks choices based on available information. A fitness function is a function that measure how effective a current solution is during an optimization process. There are many heuristic/fitness functions that can be used for level generation, such as difficulty, time, and space. In our experiments we used three functions, chosen because of their differing properties to show how these differences effect algorithmic output. Each of the three functions have a calculate to a value between 0 and 1, with the goal of any of the algorithms to maximize this value in its final solution.

A. Number of empty tiles

This function measures the number of empty tiles in a given map.

$$V = \begin{cases} (e/r1) & \text{if } e < r1 \\ 1 & \text{if } r1 \leq e \leq r2 \\ (t - e)/(t - r2) & \text{if } e > r2 \end{cases} \quad (1)$$

where e = number of empty tiles in map, t = total number of tiles in map, $r1, r2$ = a given acceptable range, and $r1 < r2$.

B. Path length

In this function, the length of the longest path in the map is calculated. The path between every pair of points that can possibly exist in the map is calculated, with the longest path defining the value.

$$V = \begin{cases} p/n & \text{if } p < n \\ 1 & \text{if } p \geq n \end{cases} \quad (2)$$

where p = length of longest path, n = a predefined “goal length” number

C. Connectivity

The connectivity function measures how well-connected empty spaces are to each other. There is inverse relationship between how well-connected a map is and how many unconnected empty regions exist.

$$V = 1/r \quad (3)$$

where r = the number of disconnected empty regions.

VII. EXPERIMENTS

Each algorithm ran 3000 times for each representation and each fitness/heuristic combination using three different empty tile initialization percentages of 25%, 50%, and 75%, for a total of 72 different experimental configurations (the initialization conditions are aggregated for a given configuration). The empty tile initialization percentages are the probabilities of a tile being empty upon map initialization. For example: if empty tile initialization percentage is equal 25%, it means each tile has a 25% chance of being empty. We use different percentages to make sure that non of the algorithms get stuck due to bad starting position. As a final configuration note, each configuration was run on an High Performance Computing Node with an Intel Xeon E-2690 Processor.

Every generation takes at most 60 seconds to generate maps of size 10x10. If no final solution is found within that time, the algorithm stops and returns the best found map. Different parameters are used for each of the fitness/heuristic functions. The empty tiles function tries to find maps in the range between 45 and 65 empty tiles, the path length function tries to find maps with the longest path greater than or equal to 26, while the connectivity function tries to find maps where all empty tiles of the map are connected to each other using orthogonal directions.

We tuned the parameters of each algorithm to make sure they are running efficiently. We used a changing C for the MCTS algorithm (described in Section III-D), a mutation rate of 5% for the hill climber, a cooling rate of 0.99 for the simulated annealing, a μ of 10 and λ of 20 with 5% mutation rate for evolutionary strategy, and a population of size 200 with crossover rate 80% and mutation rate equal to 5% and 0.5% elitism for the genetic algorithm.

In the following subsections, we compare all 72 experiment configurations over all possible combinations of the eight algorithms using the various representations and heuristic/fitness functions across different performance measures. We also compare the generated maps using expressive range analysis techniques to show the regions each algorithm is covering.

A. Performance Comparison

In figure 1, we compare the ability of each algorithm to find an optimal solution. It was expected that ES and GA outperforms other algorithms due to the multiple starting points (chromosomes) compared to single starting point for tree search algorithms, hill climber, and simulated annealing. A surprising observation is that MCTS performs worse on wide representation compared to other tree search algorithms,

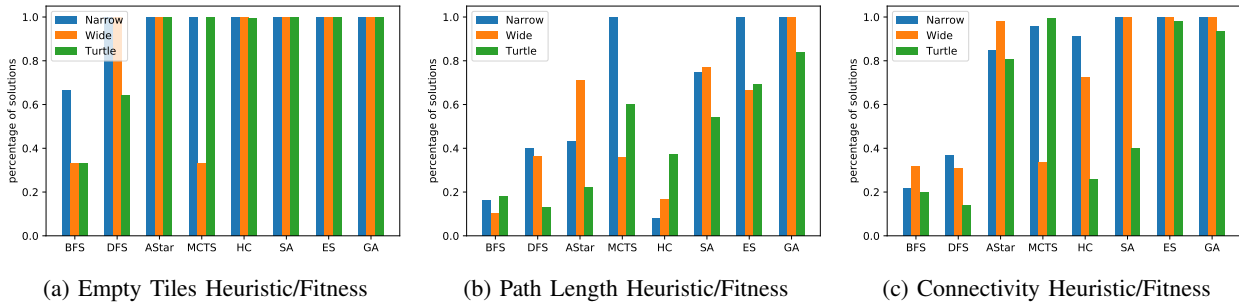


Fig. 1: Solution Percentage

as MCTS usually handles high branching factors better. Another surprising note is that DFS algorithm perform better than BFS algorithm in almost all experiments. On the topic of comparison between single point algorithms, A* and MCTS algorithms find more solutions in most of the cases when compared to hill climber and simulated annealing algorithms. One last observation is that the narrow representation seems to be the easiest of all problem representations to be explored for most of the algorithms, an exception being when using Path Length as heuristic/fitness function.

To understand more about why MCTS performs poorly on wide representation compared to other tree search algorithms, we decided to analyze the maximum depth each of the tree search algorithm is able to reach. Figure 2 shows the maximum depth reached for each of the tree search algorithms. It is obvious that MCTS on wide representation does not expand deep enough in the tree to find a solution. Similarly we can see that DFS is able to go far deeper in the tree, giving it an advantage over BFS as most of the goal maps can be found on a deeper level in the tree.

The last performance metric we analyze is time performance. Figure 3 shows the average time in seconds that each algorithm takes to find a solution. The figure shows that path length heuristic/fitness takes the longest time to find solutions compared to the other two heuristic/fitness, while empty tile heuristic/fitness almost take no time to find solutions. The time needed to compute the path length heuristic/fitness is greater than the other two heuristics/fitness, which explains part of this. Another reason is that increasing the path length might involve including either solid or empty tiles in very particular locations which may not be near local optima. Compared to the empty tile heuristic/fitness, which only cares about increasing the number of empty tiles to the grid, and the connectivity heuristic/fitness, which only cares about adding empty tiles in the correct locations to connect different regions, one could see why the path length heuristic might take more time.

B. Expressive Range Analysis

In this subsection, we analyze generated content using expressive range analysis [31]. For all the generated maps using a certain heuristic/fitness function, we measure the other two heuristic/fitness functions and plot a histogram of their distribution.

Figure 4 shows the expressive range analysis for all algorithms when using the empty tiles heuristic/fitness function. Most of the algorithms tend to have similar generation styles except for BFS using the narrow representation (figure 5a), hill climber with the wide representation (figure 5b), and DFS with the turtle graphics representation (figure 5c). The BFS generated maps with the narrow representation are more connected and have longer paths than most of the other algorithms on the narrow representation, while the hill climber generated maps are more connected and have shorter-to-medium sized paths on the wide representation. Similar to the BFS, the DFS generated maps are more connected with longer paths on the turtle graphic representation.

Figure 6 displays the expressive range analysis for all the algorithms when using path length heuristic/fitness function. In the narrow representation, BFS, DFS, and A* algorithm have different maps compared to the rest of the algorithms as they have more empty tiles with more connectivity compared to the rest of the algorithms (figure 7a). This is likely due to the fact that these tree search algorithms are visiting tiles in a fixed order compared to optimization algorithms and the random rollouts of the MCTS algorithm. In the wide representation, almost all the algorithms generate similar maps with high connectivity and high number of empty tiles except for the simulated annealing algorithm as the generated maps have the opposite features (figure 7b). In the turtle graphics representation, The graphs show that the tree search algorithms generate more connected maps with more empty tiles compared to the optimization algorithms (figure 7c). Also, one can see that the multi-point optimization algorithms have less empty tiles with less connectivity compared to the single point optimization algorithms.

Figure 8 shows the expressive range analysis for all the algorithms when using connectivity heuristic/fitness function. Using the narrow representation, the tree search algorithms have more empty tiles with shorter path lengths (figure 9a). Similar to the previous expressive range on wide representation, only simulated annealing algorithm has a slight difference from all the other algorithms (figure 9b), where the generated maps have slightly longer paths. In the turtle graphic representation, tree search algorithm maps have more empty tiles but equal path lengths to the rest of the algorithms (figure 9c).




Fig. 2: Tree Depth




Fig. 3: Time to find a solution




Fig. 4: Expressive range analysis for empty tiles heuristic/fitness.

VIII. CONCLUSION

This paper introduces the idea of using a number of tree search algorithms for content generation, specifically for map generation. We compare four different tree search algorithms (BFS, DFS, A*, and MCTS) against four different optimization algorithms (HC, SA, ES, and GA) which are used commonly in map layout generation. We test these algorithms using three different problem representations (narrow, wide, and turtle graphic representation) to see how each algorithm performs using different representations. We also use three different heuristic/fitness functions to guide the algorithms




Fig. 5: Generated maps using empty tiles heuristic/fitness function.

in finding a solution. These functions greatly differ between each other: the empty tile function is directly affected by any tile change, while path length and connectivity functions are indirectly affected depending particularly on which tiles are changed.

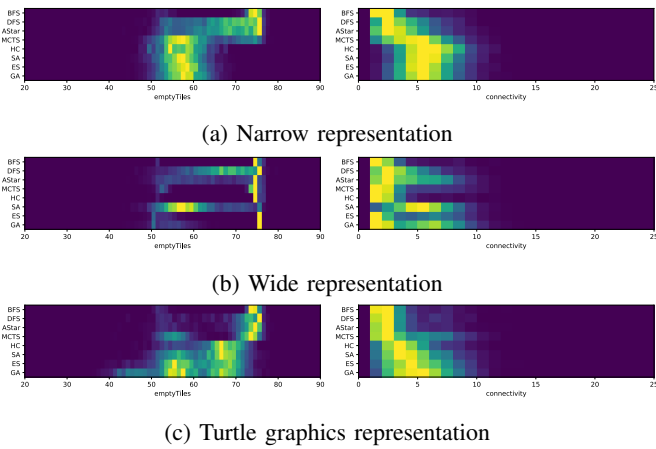


Fig. 6: Expressive range analysis for path length heuristic/fitness.

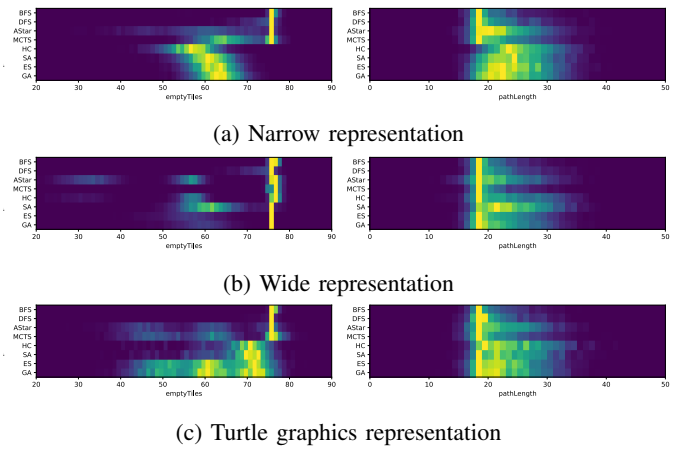





Fig. 8: Expressive range analysis for connectivity heuristic/fitness.



(a) BFS/DFS/A* generated maps using narrow representation compared to the rest of the algorithms generated maps.




(b) Simulated annealing generated maps using wide representation compared to the rest of the algorithms generated maps.




(c) Tree search algorithms generated maps using turtle graphics representation compared to the optimization algorithms generated maps.

Fig. 7: Generated maps using path length heuristic/fitness function.


All 72 experimental configurations were ran 3000 times each using different initialized map ratios. Each experiment took at most 60 seconds. The results showed that A* and MCTS are a worthy adversary to single point optimization (HC and SA) but not as powerful as multi point optimization (ES and GA). One interesting result was that MCTS does not perform well on wide representation, especially because MCTS is usually able to handle problems with branching factors. By analyzing the maximum depth that MCTS was able to reach, we found that MCTS is not reaching deep in the tree where most of the solutions exist. That is also the same reason why DFS works better than BFS in almost all the



(a) Tree search algorithms' generated maps using narrow representation compared to the optimization algorithms' generated maps.



(b) Simulated annealing generated maps using wide representation compared to the rest of the algorithms generated maps.



(c) Tree search algorithms generated maps using turtle graphics representation compared to the optimization algorithms generated maps.

Fig. 9: Generated maps using connectivity heuristic/fitness function.

representations. By comparing the generation time, we notice that the path length heuristic/fitness function is the slowest of all the techniques with fewer solutions, suggesting that it is a more difficult goal to reach compared to the other two functions. It is surprising that the simulated annealing was the fastest to find solutions using the path length fitness function with a high percentage of success to find a solution when compared to ES and GA.

We display the expressive range of each of these experimental configurations and describe how tree search algorithms have their own generative styles compared to optimization

algorithms. BFS and DFS usually find solutions that are totally different from the rest due to the different way of node expansion (not following the heuristic function). Notably, some of the single-point optimization algorithmic styles differ from the multi-point algorithms.

To conclude, we suggest that using tree search algorithms are a good alternative to generate content as shown in this work and previous work [5], [12]–[15]. We would like to take this work further and investigate multi-point tree search algorithms, treating the domain as a graph with multiple starting points similar to the work by Browne [5]. Another direction is to test these techniques on different domains and generative problems and see how well they translate (narrative generation, character generation, sprite generation, rule generation, etc). Lastly, we want to experiment with the relationship between different map initialization methods and heuristic/fitness functions while applying it on an actual game.

ACKNOWLEDGEMENTS

Ahmed Khalifa acknowledges the financial support from NSF grant (Award number 1717324 - “RI: Small: General Intelligence through Algorithm Invention and Selection.”). Michael Cerny Green acknowledges the financial support of the SOE Fellowship from NYU Tandon School of Engineering.

REFERENCES

- [1] N. Justesen, R. R. Torrado, P. Bontrager, A. Khalifa, J. Togelius, and S. Risi, “Illuminating generalization in deep reinforcement learning through procedural level generation,” in *NeurIPS Workshop on Deep Reinforcement Learning*, 2018.
- [2] N. Shaker, J. Togelius, and M. J. Nelson, *Procedural content generation in games*. Springer, 2016.
- [3] 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*, vol. 3, no. 3, pp. 172–186, 2011.
- [4] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, “A survey of monte carlo tree search methods,” *IEEE Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 1, pp. 1–43, 2012.
- [5] C. Browne, “Uct for pcg,” in *2013 IEEE Conference on Computational Intelligence in Games (CIG)*. IEEE, 2013, pp. 1–8.
- [6] S. J. Russell and P. Norvig, *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited., 2016.
- [7] J. Togelius, S. Karakovskiy, and R. Baumgarten, “The 2009 mario ai competition,” in *IEEE Congress on Evolutionary Computation*. IEEE, 2010, pp. 1–8.
- [8] S. Gelly, Y. Wang, R. Munos, and O. Teytaud, “Modification of {UCT} with patterns in {M} onte-{C} arlo {G} o,” 2006.
- [9] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot *et al.*, “Mastering the game of go with deep neural networks and tree search,” *nature*, vol. 529, no. 7587, p. 484, 2016.
- [10] D. Perez-Liebana, S. Samothrakis, J. Togelius, T. Schaul, and S. M. Lucas, “General video game ai: Competition, challenges and opportunities,” in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [11] S. W. Golomb, *Polyominoes: puzzles, patterns, problems, and packings*. Princeton University Press, 1996, vol. 16.
- [12] A. J. Summerville, S. Philip, and M. Mateas, “Mcmcts pcg 4 smb: Monte carlo tree search to guide platformer level generation,” in *Eleventh Artificial Intelligence and Interactive Digital Entertainment Conference*, 2015.
- [13] B. Kartal, J. Koenig, and S. J. Guy, “Generating believable stories in large domains,” in *Ninth Artificial Intelligence and Interactive Digital Entertainment Conference*, 2013.
- [14] B. Kartal, N. Sohre, and S. Guy, “Generating sokoban puzzle game levels with monte carlo tree search,” in *The IJCAI-16 Workshop on General Game Playing*, 2016, p. 47.
- [15] M. Graves *et al.*, “Procedural content generation of angry birds levels using monte carlo tree search,” Ph.D. dissertation, 2016.
- [16] 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*. ACM, 2013, pp. 351–358.
- [17] N. Justesen, T. Mahlmann, and J. Togelius, “Online evolution for multi-action adversarial games,” in *European Conference on the Applications of Evolutionary Computation*. Springer, 2016, pp. 590–603.
- [18] D. Ashlock, “Automatic generation of game elements via evolution,” in *Computational Intelligence and Games (CIG), 2010 IEEE Symposium on*. IEEE, 2010, pp. 289–296.
- [19] D. Ashlock, C. Lee, and C. McGuinness, “Search-based procedural generation of maze-like levels,” in *IEEE Transactions on Computational Intelligence and AI in Games*. IEEE, 2011, pp. 260–273.
- [20] D. Ashlock, “Evolvable fashion-based cellular automata for generating cavern systems,” in *Computational Intelligence and Games (CIG), 2015 IEEE Conference on*. IEEE, 2015, pp. 306–313.
- [21] C. McGuinness and D. Ashlock, “Decomposing the level generation problem with tiles,” in *Evolutionary Computation (CEC), 2011 IEEE Congress on*. IEEE, 2011, pp. 849–856.
- [22] A. Khalifa, D. Perez-Liebana, S. M. Lucas, and J. Togelius, “General video game level generation,” in *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. ACM, 2016, pp. 253–259.
- [23] A. Khalifa and M. Fayek, “Automatic puzzle level generation: A general approach using a description language,” in *Computational Creativity and Games Workshop*, 2015.
- [24] A. Khalifa, S. Lee, A. Nealen, and J. Togelius, “Talakat: Bullet hell generation through constrained map-elites,” in *Proceedings of The Genetic and Evolutionary Computation Conference*. ACM, 2018.
- [25] M. C. Green, A. Khalifa, G. A. Barros, A. Nealen, and J. Togelius, “Generating levels that teach mechanics,” in *Proceedings of the 13th International Conference on the Foundations of Digital Games*. ACM, 2018, p. 55.
- [26] A. Khalifa, M. Green, G. A. Barros, and J. Togelius, “Intentional computational level design,” in *Proceedings of The Genetic and Evolutionary Computation Conference*. ACM, 2019.
- [27] A. Khalifa and M. Fayek, “Literature review of procedural content generation in puzzle games,” <http://www.akhalifa.com/documents/LiteratureReviewPCG.pdf>, 2015.
- [28] N. Shaker, M. Shaker, and J. Togelius, “Evolving playable content for cut the rope through a simulation-based approach,” in *AIIDE*, 2013.
- [29] L. Kocsis, C. Szepesvári, and J. Willmson, “Improved monte-carlo search,” *Univ. Tartu, Estonia, Tech. Rep.*, vol. 1, 2006.
- [30] J. Brownlee, *Clever Algorithms: Nature-Inspired Programming Recipes*, 2011.
- [31] G. Smith and J. Whitehead, “Analyzing the expressive range of a level generator,” in *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*. ACM, 2010, p. 4.