

# The Riddle of Tegelby

Daniel Ashlock, Senior Member IEEE and Christoph Salge, Member IEEE

**Abstract**—At the 2017 Artificial and Computational Intelligence in Games meeting at Dagstuhl, Julian Togelius asked how to make spaces where every way of filling in the details yielded a good game. This study examines the possibility of enriching search spaces so that they contain very high rates of interesting objects, specifically game elements. While we do not answer the full challenge of finding good games throughout the space, this study highlights a number of potential avenues. These include naturally rich spaces, a simple technique for modifying a representation to search only rich parts of a larger search space, and representations that are highly expressive and so exhibit highly restricted and consequently enriched search spaces. We treat the creation of plausible road systems, useful graphics, highly expressive room placement for maps, generation of cavern-like maps, and combinatorial puzzle spaces.

## I. INTRODUCTION

Let us first put the main question into context. Procedural content generation (PCG) [26], [16], [25] is about finding algorithmic methods to generate content. We are looking specifically at PCG in games, where it is used to create a variety of things, such as levels, characters, tress, text, game rules, and even, in a somewhat aspirational fashion, whole games. There are different approaches - such as composing algorithms that produce only acceptable content. As an example, imagine you want to generate some text for a non-player character. You could just define a sentence snippet, such as “*Bring me XXX*” and then replace XXX from a list of things, such as (*a shrubbery, the holy grail, 10 Defias Brotherhood headbands*). These approaches usually rely on some form of grammar or regular expression, and if defined well, only produce acceptable content. See for example [17] or the whole range of more complicated but still compositional approaches to computation narrative [22]. The downside of those approaches is their often limited exploration of possible options. An alternative are generate-and-test methods, also referred to as search based PCG [33]. For those you need a function that can evaluate the produce content, and tell if its acceptable, or even better, how good it is. Given such a function it is then possible to produce a wider range of more experimental content, and keep those that meet the requirements, or just keep the best. Going back to our previous example, you might make sentences for your NPC by stringing random letter from the English alphabet together. You can then test if they form a proper

sentence with some grammar rules, and maybe you even have a function that tells you how good a certain sentence is. One problem with this approach is that you have no guarantee how easy it is to find a feasible solution. What we do know is that the representation of your generator matters [12]. Consider instead of stringing together random letters you string together random words from the English dictionary. While this limits the range of sentences you generator can produce, it also massively enhances the change to find an acceptable one. We know from previous research that the right representation can influence what solutions you are likely to find, and how quick you are likely to find them [13]. It is in this context that The Riddle of Tegelby was posed: “Can you define a representation where every way of filling it in yields a good game?”

While this is trivially true - one could design a search space so small that it just contains one (good) game - we do not currently provide a satisfactory answer to the spirit of this question. We mainly look at systems that provide parts of games, rather than full games, and we mostly look at approaches that increase the chances of finding good solutions, rather than guaranteeing it. What we do provide is a list of example that contain different approaches to finding or designing such search spaces, and we also discuss the properties or tricks used to make them.

One common theme here is the trade off between expressivity of the space, i.e. how different are the things the generator produces, vs. your chance of finding a good solution. We will not look at this quantitatively here, as one could with an expressive range analysis [30]. We rather wanted to give an overview of different ways one could push to the Pareto front of this trade-off - looking for solutions that offer the best of both. We also want to discuss how certain representation achieve this, and take a look on how well those deal with adaptation to existing content and control from a human user. While we are interested in presenting the representation within this document, the vision for this paper is to figure out strategies on how to address the riddle of Tegelby in its entirety.

### A. Overview

The rest of this paper will present different examples for generative systems, and discusses their properties.

We first present a naturally rich representation that is highly controllable base on Voronoi tilings [15]. Then, we look at a search space that can be restricted to make it naturally rich for generating apoptotic cellular automata [8]. These examples are not generalizable except by a process of discovery, they contain no overarching principle that makes

Daniel Ashlock is at the University of Guelph, dashlock@uoguelph.ca  
Christoph Salge is at the University of Hertfordshire, christophsalge@gmail.com

The authors thank the Canadian Natural Sciences and Engineering Research Council of Canada (NSERC) and the European Commission (Grant INTERCOGAM) for supporting this work.

the spaces rich, but it is still an interesting example of a rich search space.

Beyond having a naturally rich space a technique called *single parent crossover* is presented that can be used to focus search in a rich part of the space. The effect of single parent crossover is to focus search in the part of gene-space where examples lie. The technique acts by permitting immortal population members, which are the examples, participate in a one-sided form of crossover [7], [14].

Finally a new technique called *convex representation* is presented, in which we can pick out a small subset of a search space and use an alternate representation to search only that part of the space. A representation is convex if the weighted average of instances of a representation are, themselves, instances of the representation. This permits search to be restricted, in a transparent fashion, to the convex hull of a collection of examples.

## II. RELATED WORK

Automated level generation in video games can arguably be traced back to a number of related games from the 1980s (Rogue, Hack, and NetHack), collectively called *Roguelike games*. The task is currently of interest to the research community. In [31] levels for 2D sidescroller and top-down 2D adventure games are automatically generated using a two population feasible/infeasible evolutionary algorithm. Answer set programming is another approach to dungeon generation [29]. In [32] multiobjective optimization is applied to the task of search-based procedural content generation for real time strategy maps.

## III. NATURALLY RICH SPACES

First, we will look at an example from a naturally rich space - by that we mean representations where nearly all examples from the space are feasible solutions.

### A. Voronoi Tiling

A problem that arises from time to time is laying out streets for a procedurally generated village or town. We begin with the generation of a system of connected paths like those shown in Figure 2. A Voronoi diagram [15] is created by choosing a collection of points, called tile centers, in the plane and then creating tiles by declaring the tile associated with one of the points to consist of the parts of the plane closer to that tile. Two sorts of Voronoi diagrams are shown in Figure 1. The upper one is a standard tiling, the lower one is modified by weighting the individual points and defining tiles based on distance times weight; this makes the boundaries quadratic curves instead of straight lines, giving the tiles curved sides.

Figure 2 shows three examples of layouts for the roads in a segment of a procedurally generated village. These are the boundaries between Voronoi tiles. If the weighting factors are not too different from one another, then the resulting street map is connected. The three examples given in Figure 2 show that these layouts are easily controllable. The left-most example uses 100 random points. The middle uses a regular




Fig. 1. The upper panel shows a standard Voronoi diagram, the lower one a weighted Voronoi diagram with curved tile sides.

$9 \times 9$  grid of points while the right-most adds 19 random points to the 81 regular points from the middle example.

This space of layouts for streets always yields a connected set of paths, the shape can be controlled by the choice of tile centers, and could be procedurally refined if specific criteria are required. In any case this *Voronoi path network* generation scheme is the first example of a naturally rich space. If complex requirements were needed for a given game, then the space of Voronoi paths could be procedurally searched for refined designs, but it starts eminently rich in good structures.




Fig. 2. Shown are examples of weighted Voronoi street maps using random points (left), a square grid of points (middle) and a square grid of points with 19 added random points (right).

The question is now, why is the Voronoi tiling such a good representation? One answer to this question lies in the specific genotype-phenotype mapping of the Voronoi representation. The genome, i.e. the parameters of the actual search space, are just positions of points, yet they get translated into a phenotype, the actual tiling seen in the figures, in a relatively complex process, as described above. First, this process is very similar to several processes in nature, leading to a certain familiarity. Voronoi tiling is a particular good fit if the artefact generated in question, such as a street map, could in fact be generated by some self-organizing process based on distance based tiling in the real world.

The same process also allows for a certain degree of control over the output. The distance function used can for example be weighted to create curved lines. One could even use it to adapt the output to existing content. Imagine, for example, that they points were in a real existing map, and the tiling would be some kind of separation in nation states. The distance function could then be based on infrastructure, i.e. existing roads, etc., which would then lead to an even richer tiling that adapts to some underlying property.

So, the Voronoi approach is a rich representation, that allows for decent control and adaptive, but is best used to generate things that arise from a similar process as it complex genotype-phenotype mapping. One approach to duplicate its success would be to look at representations that use a process similar to the process generating the artefact, i.e. rather than generating the final artefact one might try to find a mathematical representation of how this thing came about.

### B. Apoptotic Cellular Automata

Another example of a naturally enriched space is that of the apoptotic cellular automata described in [8]. CAs have been applied to the visual arts, used to produce artistic images[11], [24], and their use has been extended to the fields of architecture and urban design[27], [18]. An interesting application has been the use of CA in simulating the emergence of the complex architectural features found in ancient Indonesian structures, such as the Borobudur Temple[28]. Ashlock

and Tsang[11] produced evolved art using 1-dimensional CA rules. CA rules were evolved using a string representation. The CA either underwent slow persistent growth, or planned senescence. The resulting fitness landscapes were rugged with many local optima.

An *apoptotic cellular automaton* is one that, for a given initial condition, fills space and then “dies”, i.e. returns all cells to a quiescent state. The fitness of an automata is the space it fills, save that it is awarded zero fitness if it grows too long. Examples of these automata appear in Figure 3. In [8], random transects of the fitness landscape for these automata were investigated. This yielded a startling discovery. A random transect is taken by starting with a high-fitness automata rule, represented by an array of integers, and then changing the entries, in a random order, to those of another rule. The intermediate rules so generated form the points in the transect. The fitness of automata in the transects were at least two hundred times as large as the fitness of randomly sampled automata.

Again here the process that maps the actual genome, the relatively compact starting condition, to the outcome, is a complex process that is also used to model a range of natural phenomena. Again, the process can be manipulated to adapt the outcome to underlying properties of the environment, and we can even imagine setting some those cells to fixed values beforehand to simulate the integration of existing content.

This is not an example replicable outside of the apoptotic cellular automaton space, though it does enable an example of one of the enriched representations in the next section. The full space of apoptotic cellular automata rules is extremely poor in good optima, but, for reasons we do not understand, there is a huge cluster of them in one part of the fitness landscape. The rules in this paper were 36 dimensional, with one constraint – that the zeroth entry must be zero – to enable death. This means that there are  $2^{35}$  random transects between any two high fitness rules, of which hundreds were located. This gives us a very rich search space for these images that could be used as aliens in a space-invader like game or enemy ships in a game like Galaga. These images




Fig. 3. Ten examples of renderings of evolved apoptotic cellular automata.

can be generated from a 36 byte rule, permitting thousands to be pre-generated and stored in even a small application.

An example of a highly enriched space that occurs outside of games is that of *side effect machines* [23]. A side effect machine is a finite state machines whose transitions are driven by a biological sequence, like DNA. The machine counts how many times it enters each state and the counts form features for classification. Side effect machines are evolvable transducers for changing variable length biological sequences into numerical features. In the course of the research of side effect machines it was found that evolving these machines with a fitness function based on classification accuracy started at high values with evolution improving the fitness by 10-20% over that available in the initial random population.

#### IV. ENRICHED REPRESENTATIONS

In this section we talk about enriched representation, those that are not naturally rich in examples, but can be modified or limited in a way that produces a rich subspace.

##### A. Apoptotic Cellular Automata

The apoptotic cellular automata use a *drawing arena*. If their time histories remain alive when they reach the bottom of the arena, then they did not die in time and are awarded zero fitness. The automata in Figure 3 used a  $201 \times 201$  arena. A problem with this is that these are relatively small pictures. In [7], single parent crossover was used to generate much larger pictures that had a similar appearance to that of a smaller apoptotic cellular automata rendering. The evolutionary algorithm was modified to permit crossover with an example gene called an *ancestor* and the drawing arena was set to a much large size.

The use of single parent crossover re-injects the content of the ancestor into the population over and over. This makes genes similar to the ancestor very likely and, as a side effect, restricts the search space to points near the ancestor. Because of the dense packing of high fitness genes in the apoptotic cellular automata search space, the algorithm rapidly locates

high fitness genes, giving us the desired enrichment of the search space. Usually, though not always, these genes are also similar in appearance to the ancestor. This permits broad search on small examples – which run far faster – and then selection of a desired appearance to seed larger images through single parent crossover. An example of four such renderings of apoptotic cellular automata are shown in Figure 4.

While we have a ready example available with apoptotic cellular automata, single parent crossover may be used to promote information re-use and focus evolutionary search near an ancestor or ancestors in any representation that uses crossover. It has already been found that single parent techniques substantially enhance the evolution of virtual robots, using multiple ancestors, though this data have not yet been published.

##### B. The *do what's possible* representation

A general technique for generating highly enriched search spaces is the *do what's possible* representation [4]. In a paper accepted to the CoG 2019 it was found that this representation can lay out rooms using a tiny data structure. The original publication on the DWP representation [4] used the representation on a variety of problems. It was capable of solving self avoiding walk problems [1] of unprecedented size. Where  $12 \times 12$  had been a practical limit with even adaptive representations [9], the DWP representation solved  $40 \times 40$  cases of the problem. Applied to the problem of evolving a Gray code [34] the DWP representation enabled the evolution of an 8-bit Gray code with 256 members. The DWP problem was also used to evolve entropically rich binary strings and a type of classification character for DNA in the initial study. A later study refined the systems power for DNA classification [3].

In [10] the representation was used to lay out large numbers of rooms in prototype dungeon maps. The key to this is the use of the self-driving automata shown in Figure 5. This data structure, driving its own transitions with its output, generates a stream of bits that, while deterministic, can have




Fig. 4. Examples of four renderings of apototic cellular automata derived from the same ancestor.




Fig. 5. An example of a self driving automata over the binary alphabet and the first few characters that it emits. The sourceless arrow denotes the initial state. The strings emitted are shown on the states. The automata’s output is used as its input.

quite a complex pattern [4]. Reading these bits in groups to generate integers of desired size permits the self driving automata to execute a serial decision process about where to place rooms. The time needed to locate good examples is short, but does not reach quite the level of enrichment desired.

The room generator, however, places rooms when it can, it “does what’s possible” which means a high fitness room generator is reusable in different situations. The self driving automata generates a stream of proposed rooms which are laid down when they fit and rejected when they do not. The decision process used is the one described in [10]. We used a Voronoi path network, from Section III and applied the room-layout technology. Four examples of layouts by different room layout engines with different evolved self-driving automata are shown in Figure 6.

One modification was made in the representation for this study. The room selection process chooses a room by taking a large integer supplied by the self driving automata modulo the number of existing rooms. This study disallows a room that has failed to have another room placed next to it eight or more times. This directs growth to the active periphery of the the current layout.

In this example the enrichment is again due to manipulating the process that translates the genotype to the phenotype, but this time the feasibility filter specifically filters out actions that would lead to bad results. A search in this genome space would be much more efficient, because the function that translates any solution to an actual room map processes a lot more information and builds a solution that works. As the maps in Figure 6 also show, the solutions are also capable of adapting to existing elements of the environment.

### C. Exploiting convex representation

A representation consisting a a list of real parameters is *convex* if the weighted average of two instances of the representation is an instance of the representation. When a representation is convex, then we may choose example genes and evolve weight vectors for averaging them. We call this process *convex re-representation*. It restricts the search space to the convex hull of the example genes. This both can lower the search dimension and, if the examples are well chosen, substantially enhance the richness of the space. An example showing that such enrichment by use of examples is possible appears [2].

In [20] the authors use a two-dimensional cellular automata with a majority based rule to generate cavern-like maps. A *fashion based cellular automaton* is an automata whose rule is based on a competition matrix that gives the score that a cell in one state obtains against a cell in



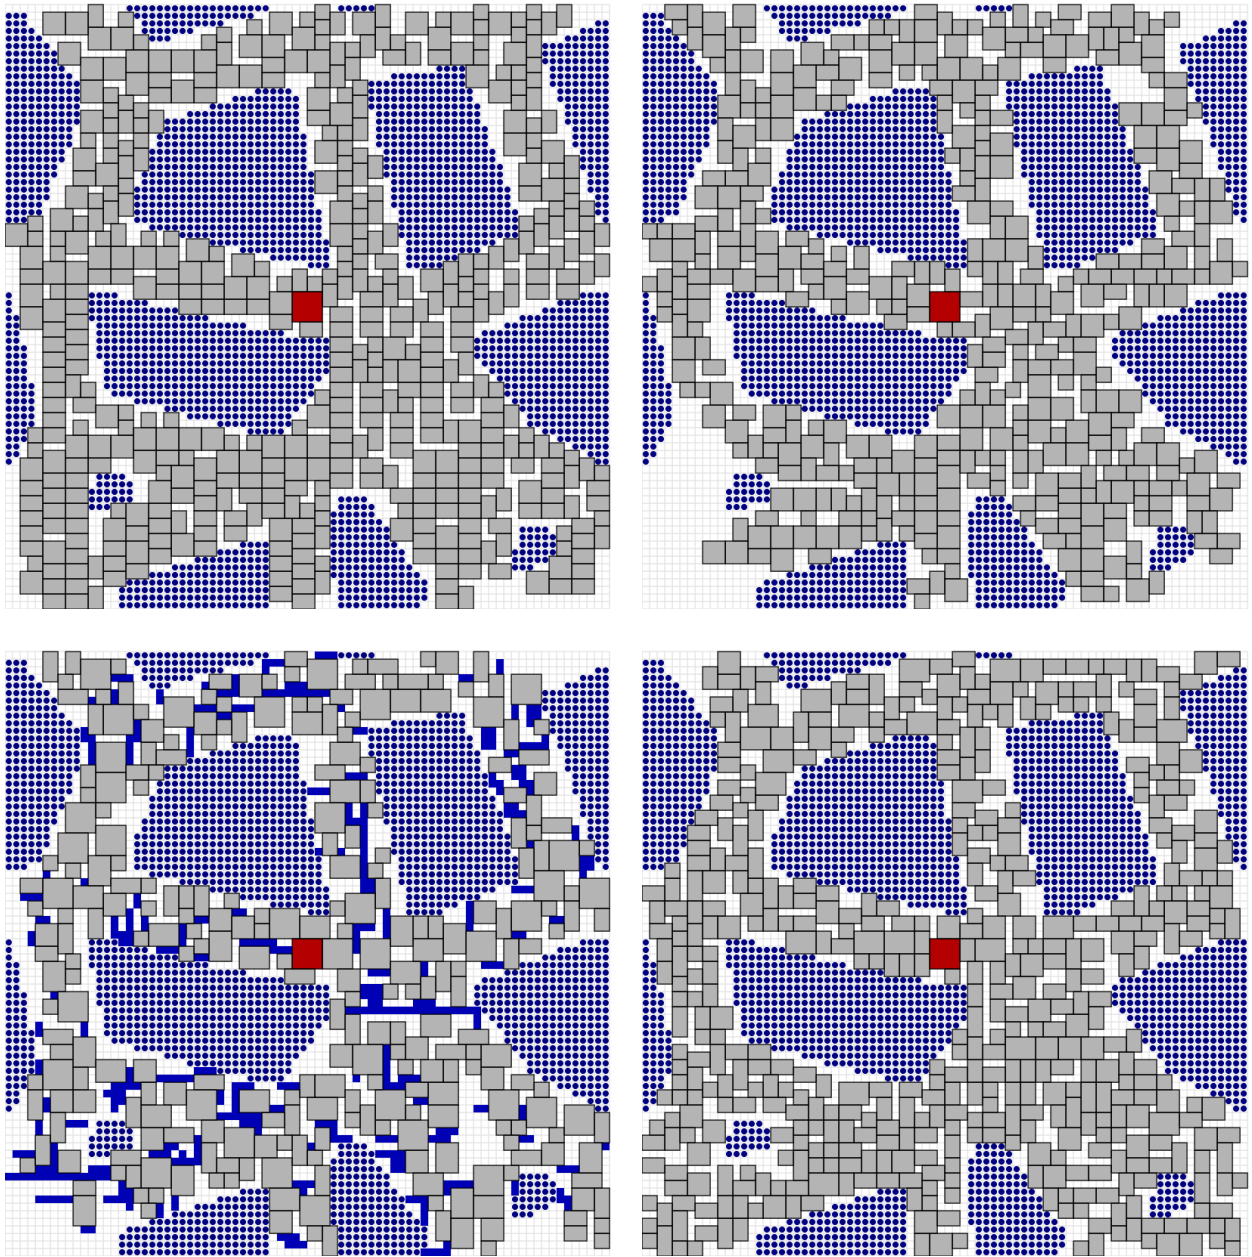


Fig. 6. Four different room layouts made with do-what’s possible room generators. Grey rectangles are rooms, blue lines are corridors. These room assignments were placed in a set of passages generated as the borders of a Voronoi diagram. The red room is the first room laid down.

another. Updating consist of compiling the scores, based on current neighbors, of each cell and then adopting the state of the highest scoring member of a cells neighborhood. This representation was used to make cavern maps with diverse appearances. In [5], [21] pairs of automata were evolved so that the weighted averages of the pairs all generate good maps.

An example demonstrating, via continuously changing weighted average from left to right, the intermediate rules between two co-evolved examples if shown in Figure 7. The representation for fashion based cellular automata rules is a square matrix or real parameters, the score values, of size

$N \times N$  where  $N$  is the number of states. The weighted average of two such matrices is such a matrix. The morph between two rules demonstrates that convex re-representation is an excellent way to enrich a search space by restricting it to a good region while retaining diversity of possible solutions.

## V. CONSTRUCTIVELY RICH SPACES

Figure 8 shows an example of a puzzle located in an automatic content generation effort. The player is presented with the polyominos, represented in the figure by colored regions, and a board with the numbers on it. To work the puzzle, the player puts the polyominos on the board. Their score is the sum of the scores for the individual polyominos,




Fig. 7. A rendering of a fashion based cellular automata that traverses the weighted averages of two rules from left to right.

which are themselves the better of the sum or product of the numbers under the polyomino. The score of 300 in the figure can be verified by the reader and is thought to be optimal.




Fig. 8. An example of a polyomino math puzzle.

The software that located these puzzles is as yet unpublished and is a terrible answer to the problem of finding an enriched search space. Worse, a majority of the puzzles located achieved their best scores when some of the pieces were not used – this is the *common* outcome of procedural content generation on this problem. This, however, is not the reason we include these puzzles.

It is clear that, once you know the size of the polyominos, then the numbers may be optimized by simply assigning them to groups with the same sizes as the polyominos. The score can then be optimized by a search that changes that assignment. Suppose that we have an optimal assignment. Then placement of that assignment under the polyominos will yield a board where the optimal solution uses all the polyominos. This still does not give us an enriched space.

First of all, note that the order of the numbers under a given polyomino is irrelevant to the score. In addition, once there is one way to place polyominos to fill a rectangle, there is often a huge combinatorial space of such placements

[6]. This means that locating one puzzle where the optimal solution uses all the pieces gives us a combinatorial space of thousands to millions of equivalent puzzles with very different number layouts.

This is a constructively rich design space and the idea exemplified here can be applied to other sorts of puzzles. A good analogy could be made with a biological neutral network [19] – the many forms of the initial puzzle form an enormous neutral network in the search space. This means there is an operator that can take one already found good solution and produce a large range of equally good solutions. This allows one, once a single good solution was found, to define a subspace that only contains good solutions. This means that we have managed once to answer the Riddle of Tegelby successfully, albeit in a very limited domain.

## VI. CONCLUSIONS

This study has given a number of methods, some quite special purpose and others general, for creating enriched search spaces for game content. Based on the example so far, it looks like the rich search spaces we saw so far all shared a relatively complex translation process from their minimal representation to the final artefact. These non-deterministic processes were able to process additional information to ensure the feasibility of the produced artefact, and could often be manipulated in understandable ways to influence the output. This allowed for more human control on the output on and more option to adapt to existing content.

The other main method discussed here are way to generate subspaces from bigger subspaces, by finding operators that modify representations to produce other artefacts of similar quality, and thereby defining a good subspace, or concentrating of limiting the search to subspaces that have a higher rate of good solutions.

It is to be hoped that this paper gives the games research community some grist for their research mills. Among the things to do are the following. Locate more game-content types that are rich as a matter of chance or nature, like the Voronoi path networks. Work to understand the capabilities and limits of single parent techniques, convex representations, and self-enriching representations like the do-whats-possible representation. At a lower priority level, construct other spaces like the polyomino math puzzle where one positive example has a billion cousins that are all good games. Beyond this, there are doubtless other ways to attack

the Riddle of Togelby, we hope to hear back that our colleagues have found more of them.

## REFERENCES

- [1] D. Ashlock. *Evolutionary Computation for Optimization and Modeling*. Springer, New York, 2006.
- [2] D. Ashlock, J. A. Brown, and L. Sulancva. Exploiting fertility to enable automatic content generation to ameliorate user fatigue in interactive evolutionary computation. In *Proceedings of the 2018 Congress on Evolutionary Computation*, pages 1392–1398, 2018.
- [3] D. Ashlock, S. Gillis, and W. Ashlock. Infinite string block matching features for dna classification. In *Proceedings of the 2017 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology*, pages 1–8, Piscataway NJ, 2017. IEEE Press.
- [4] D. Ashlock, S. Gillis, A. McEachern, and J. Tsang. The do what’s possible representation. In *Proceedings of the IEEE 2016 Congress on Evolutionary Computation*, pages 1586–1593, Piscataway NJ, 2016. IEEE Press.
- [5] D. Ashlock and M. Kreitzer. Evolving diverse cellular automata based level maps. In *Proceedings of the 6th International Conference in Software Engineering for Defence Applications*, pages 10–23, 2019.
- [6] D. Ashlock and C. McGuinness. Graph-based search for game design. *Game and Puzzle Design*, 2(2):68–75, 2016.
- [7] D. Ashlock and S. McNicholas. Single parent generalization of cellular automata rules. In *Proceedings of the 2012 IEEE Congress on Evolutionary Computation*, pages 179–186, Piscataway NJ, 2012. IEEE Press.
- [8] D. Ashlock and S. McNicholas. Fitness landscapes of evolved cellular automata. *IEEE Transactions on Evolutionary Computation*, 17(2):198–212, 2013.
- [9] D. Ashlock and J. Montgomery. An adaptive generative representation for evolutionary computation. In *Proceedings of the IEEE 2016 Congress on Evolutionary Computation*, pages 1578–1585, Piscataway NJ, 2016. IEEE Press.
- [10] D. Ashlock and C. Salge. Automatic generation of level maps with the do what’s possible representation. Accepted to the 2019 IEEE Conference on Games, 2019.
- [11] D. Ashlock and J. Tsang. Evolved art via control of cellular automata. In *IEEE Congress on Evolutionary Computation, 2009*, pages 3338 – 3344, May 2009.
- [12] Dan Ashlock, Sebastian Risi, and Julian Togelius. Representations for search-based methods. In *Procedural Content Generation in Games*, pages 159–179. Springer, 2016.
- [13] Daniel Ashlock and Eun-Youn Kim. The impact of cellular representation on finite state agents for prisoner’s dilemma. In *Proceedings of the 7th Annual Conference on Genetic and Evolutionary Computation, GECCO ’05*, pages 59–66, New York, NY, USA, 2005. ACM.
- [14] W. Ashlock and D. Ashlock. Single parent genetic programming. In *Proceedings of the 2005 Congress on Evolutionary Computation*, volume 2, pages 1172–1179, Piscataway NJ, 2005. IEEE Press.
- [15] F. Aurenhammer. Voronoi diagrams a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345405, 1991.
- [16] Kate Compton. So you want to build a generator, feb 2016.
- [17] Kate Compton, Ben Kybartas, and Michael Mateas. Tracery: an author-focused generative text tool. In *International Conference on Interactive Digital Storytelling*, pages 154–161. Springer, 2015.
- [18] M. Devetakovic, L. Petrusovski, M. Dabic, and B. Mitrovic. Les folies cellulaires an exploration in architectural design using cellular automata. *12th Generative Art Conference*, pages 181–192, 2009.
- [19] Marc Ebner, Peter Langguth, Jürgen Albert, Mike Shackleton, and R A Shipman. On neutral networks and evolvability. 2001.
- [20] Lawrence Johnson, Georgios N. Yannakakis, and Julian Togelius. Cellular automata for real-time generation of infinite cave levels. In *Proceedings of the 2010 Workshop on Procedural Content Generation in Games, PCGames ’10*, pages 10:1–10:4, New York, NY, USA, 2010. ACM.
- [21] M. Kreitzer, D. Ashlock, and R. Pereira. Automatic generation of diverse cavern maps with morphing cellular automata. Accepted to the 2019 IEEE Conference on Games, 2019.
- [22] Ben Kybartas and Rafael Bidarra. A survey on story generation techniques for authoring computational narratives. *IEEE Transactions on Computational Intelligence and AI in Games*, 9(3):239–253, 2017.
- [23] A. McEachern and D. Ashlock. Shape control of side effect machines for DNA classification. In *Proceedings of the 2014 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology*, pages 1–8, Piscataway NJ, 2014. IEEE Press.
- [24] G. Monro. Emergence and generative art. *Leonardo - MIT Press*, 42(5):476–477, 2009.
- [25] Noor Shaker, Julian Togelius, and Mark J Nelson. *Procedural content generation in games*. Springer, 2016.
- [26] Tanya X Short and Tarn Adams. *Procedural Generation in Game Design*. CRC Press, 2017.
- [27] V. Singh and N. Gu. Towards an integrated generative design framework. *Design Studies*, in press, 2011.
- [28] H. Situngkir. Exploring ancient architectural designs with cellular automata. *BFI Working Paper No. WP-9-2010*, 2010.
- [29] Anthony J Smith and Joanna J Bryson. A logical approach to building dungeons: Answer set programming for hierarchical procedural content generation in roguelike games. In *Proceedings of the 50th Anniversary Convention of the AISB*, 2014.
- [30] Gillian Smith and Jim Whitehead. Analyzing the expressive range of a level generator. In *Proceedings of the 2010 Workshop on Procedural Content Generation in Games*, page 4. ACM, 2010.
- [31] N. Sorenson and P. Pasquier. Towards a generic framework for automated video game level creation. In *Proceedings of the European Conference on Applications of Evolutionary Computation (EvoApplications)*, volume 6024, pages 130–139. Springer LNCS, 2010.
- [32] Julian Togelius, Mike Preuss, and Georgios N. Yannakakis. Towards multiobjective procedural map generation. In *PCGames ’10: Proceedings of the 2010 Workshop on Procedural Content Generation in Games*, pages 1–8, New York, NY, USA, 2010. ACM.
- [33] Julian Togelius, Georgios Yannakakis, Kenneth Stanley, and Cameron Browne. Search-based procedural content generation. In *Applications of Evolutionary Computation*, volume 6024 of *Lecture Notes in Computer Science*, pages 141–150. Springer Berlin / Heidelberg, 2010.
- [34] J. H. van Lint and R. M. Wilson. *A Course in Combinatorics, second edition*. Cambridge University Press, New York, NY, 2001.