

Multi-Faceted Evolution Of Simple Arcade Games

Michael Cook and Simon Colton,
Computational Creativity Group, Department of Computing,
Imperial College, London, UK
ccg.doc.ic.ac.uk

Abstract— We present a system for generating complete game designs by evolving rulesets, character layouts and terrain maps in an orchestrated way. In contrast to existing approaches to generate such game components in isolation, our ANGELINA system develops game components in unison with an appreciation for their interrelatedness. We describe this multi-faceted evolutionary approach, and give some results from a first round of experimentation.

I. INTRODUCTION

Computational intelligence research has done much to help the games industry attack the content generation problem in recent years. As consumers expect a greater breadth and depth of content in videogames, procedural generation techniques have been used to alleviate the burden on developers and give way to more dynamic and interesting games. This research has contributed work towards the generation of level designs [1], [2], [3], [4], [5], visual content [6], [7], controllers for game entities [8] and game rulesets [9], [10], [11]. Much of the material written about this work mentions a move towards *automated game design*, where an intelligent system is responsible for the overall development of a game. However, we argue that in its current state, such research is little more than an assistant in the game design process, not a designer itself.

For a system to take on the role of a game designer, it is not sufficient for it merely to generate all of a game's components by itself. It must also understand the relationship those components have with one another, and how a change in one component affects the final game as a whole. We have developed a system, called ANGELINA¹, that designs games from scratch by evolving the constituent game components with respect to one another. ANGELINA is composed of separate evolutionary processes which design a game's ruleset, its level and the layout of the game's non-player characters (NPCs), combined with an understanding of the dependencies between these processes that allow it to influence the end result accordingly. The types of games that ANGELINA can currently generate are limited to simple arcade games, as portrayed in Fig. 1, but we intend to extend this to more sophisticated games in the future.

The remainder of this paper is organised as follows: in section II, we introduce ANGELINA and outline the process by which it evolves games. In section III, we present some sample games generated by ANGELINA and evaluate the quality of the output, showing that ANGELINA is capable of

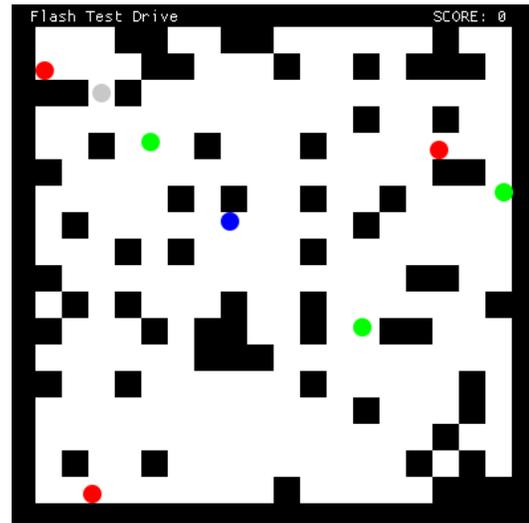


Fig. 1: A screen from a game made by ANGELINA. The player is grey, and the NPCs are coloured red, blue and green.

working alone or assisted by a human, and demonstrating the importance of the system's awareness of the finished game. In section IV, we look at existing research in automated game design, and discuss how ANGELINA fits into this context. Finally, in section V, we discuss the future for ANGELINA and automated design beyond this project.

II. THE ANGELINA SYSTEM

A. The Anatomy of a Game

We represent a game as three flexible components - a *map* which defines passable and impassable areas in a two-dimensional grid of square tiles, where a tile is a small square of space in the game world; a *layout* which describes the arrangement of player- and non-player-characters on a similar two-dimensional grid; and a *ruleset* which describes the effects of collisions between types of game entities, as well as movement types for the NPCs, time limit and score limit for the game.

A representation of an example map and layout can be seen in the game in Fig 2. The state space for these two components is the space of all possible configurations of the 2D arrays representing them, and their representation inside the system is simply a list of binary (in the case of maps) and integer (in the case of layouts) values representing the 2D grid. For a map, each square has two possible values.

¹A Novel Game-Evolving Labrat I've Named ANGELINA

```

<scorelimit>44</scorelimit>
<timelimit>72</timelimit>
<redmovement>CLOCKWISE</redmovement>
<greenmovement>RANDLONG</greenmovement>
<bluemovement>RANDLONG</bluemovement>
<rules>
  <rule>PLAYER, RED, DEATH, DEATH, 1</rule>
  <rule>GREEN, BLUE, TELEPORT, DEATH, 0</rule>
  <rule>OBSTACLE, RED, NOTHING, TELEPORT, 0</rule>
</rules>

```

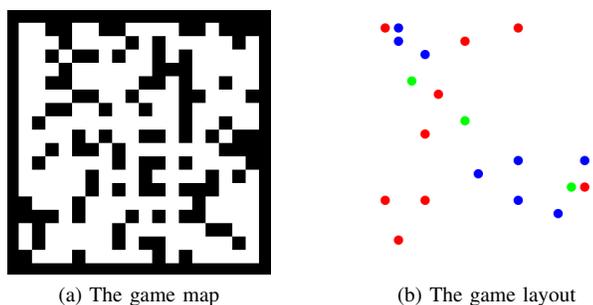


Fig. 2: A sample game output showing the structure of a game produced by ANGELINA.

For a layout, each square has five possible values (empty, player, and three types of NPC - red, green and blue). For an average map of twenty by twenty tiles, the state space contains 1.024×10^{13} layouts and 1.6×10^5 maps.

The state space described by the rulesets is derived from those defined in [10]. When two objects collide, rules in the ruleset may fire. The XML code in Fig 2 shows a sample ruleset. Rules have conditions, which are the two entities whose collisions they are triggered by (possibly including the player), and the effects applied to the two colliding objects. These effects are one of either *death*, i.e., removing the object from play; or *teleport*, which moves the object to a randomly-chosen, non-obstructed point on the map; or simply *nothing*. A rule may only apply an effect to the two objects that triggered the rule. Each rule also has an associated score value which may be positive, negative or zero.

Overall, a rule can be represented as a 5-tuple of integers coding for the various components; two entity types (chosen from the player, the three NPC colours, and walls) and two effect types (chosen from *death*, *teleport* and *nothing*), and a score change of one of $+1$, 0 or -1 . The state space for rulesets varies, because a ruleset can contain any number of rules. Each rule has a state space of $5^2 \times 3^3$, so an average-sized set of three rules has a state space of 675^3 . In addition, a ruleset in ANGELINA also contains a score and time limit, both of which are integers valued between 0 and 100 inclusive, as well as a movement class for each NPC colour, which is one of three behaviours:

- Random Path - the NPC makes a random 90° turn after a fixed distance of movement.

- Wall Hug - the NPC moves in straight lines until it reaches an obstacle, whereupon it turns either left or right 90° .
- Random Walk - the NPC makes 90° turn at random intervals.

A game created out of a combination of a map, ruleset and layout can be converted into an intermediate representation by ANGELINA for compilation to an executable.

B. Multi-Faceted Evolution

The high-level structure that defines how ANGELINA evolves its games is based on an extension of computational evolution we have called *multi-faceted evolution*. We have developed a new framework for artefact generation that uses individual evolutionary strands in a compositional manner, to share information during the evolutionary process and use the fitness of individual components to increase the overall fitness of the finished artefact.

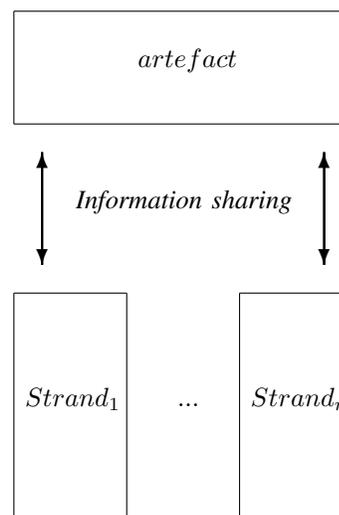


Fig. 3: The structure of a multifaceted evolutionary process.

Fig. 3 shows a basic, high-level view of a multi-faceted evolutionary process. At the top is an *artefact* object, an abstract representation of the objects that the system is intending to produce. In the case of ANGELINA, this object is a Game. In MFE, artefacts are co-ordinating objects that share information between two or more *strands*. A strand is an autonomous evolutionary process that has crossover methods, fitness functions and a population of solutions.

In the case of ANGELINA, we have three strands - a Ruleset strand, a Layout strand and a Map strand. Each strand is responsible for evolving its namesake subcomponent described in the previous section. All strands have populations containing 100 members, and are typically iterated for between 100 and 250 generations. The key difference between a strand and a standard evolutionary process is that in MFE a strand has two fitness functions - an *external* function and an *internal* one.

The *internal* fitness function is similar to what an ordinary evolutionary process might contain. It evaluates a candidate solution according to metrics appropriate to the problem, and does not appeal to any information outside the strand. In ANGELINA, an internal fitness function for a map considers how ‘interesting’ the map is, according to several metrics that we discuss later. It does not consider the layout or rules in order to do this. In contrast, the *external* fitness function allows a strand to evaluate a member of its population in the context of the larger problem. In ANGELINA, a technique of *playouts* are used to provide external fitness data to each strand. Playouts will be explored in a subsequent subsection.

An MFE process is carefully designed to encourage separation between its strands. Therefore, there is no need for the external fitness function to be applied to every generation of every strand. A strand may only use external evaluation intermittently to check how the current population is evolving in the wider context. For ANGELINA, both internal and external fitness functions are applied at each generation, and their fitness data combined to select candidate solutions for crossover.

```

1 gen = 0
2 while fitness level not met:
3   gen++
4   foreach strand in artefact:
5     if gen % external_rate == 0:
6       evaluate population externally
7     evaluate population internally
8     select solutions for crossover
9     produce new population
10    if gen % update_rate == 0:
11      update exemplars in artefact

```

Fig. 4: A pseudocode overview of an MFE process.

In MFE processes, external fitness functions do not contact other strands directly for information. Instead, they appeal to the artefact object only for the information they need, as we see in the abstract diagram in Fig. 3. On line 10 of Fig. 4 we see that each strand checks to see if it should update the information held about it in the main artefact. If this is the case, then on line 11 it sends new exemplar data to the artefact. This data is then used on subsequent iterations on line 6, where the external fitness evaluations take place.

For ANGELINA, the exemplars used in the external fitness functions represent the highest-fitness solutions from the subpopulation. This ensures that external evaluations are done using the best examples from each strand as context. However, other applications or approaches to this problem could use a randomly-selected set of exemplars, or compare the most and least fit solutions in a given population.

C. Crossover and Mutation

In order to recombine the various elements of the games we implement a range of crossover styles, obtained through experimentation with each component type. For maps, crossover of two maps contributes eight maps in total to the new population, including the original parents themselves.

Two further children are produced by combining the two parent maps using binary OR and AND operators, where a tile in the child is blocked if it is blocked in either parent, or both parents respectively. The remaining four children are produced through scattered inheritance (motivated by initial experimentation), where an arbitrary map tile in the child map map_{child} is defined as:

$$map_{child}[i][j] = \begin{cases} map_1[i][j] & \text{if } rnd() < 0.5 \\ map_2[i][j] & \text{if } rnd() \geq 0.5 \end{cases}$$

For NPC layouts, we use some similar approaches, employing binary operators AND and OR to produce children, and including the two parents in the set of children. We also transpose sets of coloured NPCs between the parents; for example, one child may be a copy of one of its parents, but have its set of blue NPCs switched out and replaced with the set of blue NPCs from its other parent. Which sets are crossed over is chosen randomly, and each pair of parent layouts generates four children through this method, meaning that two parents create eight layouts for the next generation overall.

Ruleset evolution includes some single-parent mutation as well as dual-parent recombination. Single parents generate mutated offspring by choosing a random rule and randomly changing one of its parameters, such as the effect of the collision, or the coloured NPCs it is triggered by. Crossover works by swapping rules over between parents, so that the resulting child has a mixture of rules from its parents. For every eight children generated this way, we also include two randomly generated ‘fresh’ rulesets to avoid stagnation in the evolutionary search, which was noted in early experiments.

These evolutionary settings have been determined by some preliminary experimentation, but in future, we plan to perform systematic testing of the evolutionary parameters, to optimise the system

D. Internal Fitness Functions For Game Components

1) *Maps*: Map designs are scored using two metrics which we have termed *fragmentation* and *domination*. A map’s fragmentation score is equal to the number of *islands* present in the map, where an island is a set of blocks that are adjacent to one another and disconnected from both the wall and other islands. A map’s domination score represents the number of tiles which *dominate* sections of the map. A tile is said to dominate two other tiles if all paths between those two tiles must pass through the dominated tile. A similar definition can be found in [5]. In the vocabulary of game design, a dominating tile represents a doorway or a corridor; a tile separating two regions of the map.

Calculating which tiles dominate others is costly to do, given that it must be performed on every tile in every map in the population at every generation. We check if a tile is dominating by calculating how many tiles are reachable from that tile in the map. We call this set of reachable tiles a *flood plain*. For a tile, t , the flood plain is defined as:

$$floodplain(t, map) = |\{t_2 \text{ such that } rchable(t, t_2, map)\}|$$

where

$$rchable(p, q, m) \iff \exists \text{ a path from } p \text{ to } q \text{ in map } m$$

We then modify the map by blocking out t and making it inaccessible. We then perform the same calculation for a neighbour of t , t_{nb} , in the modified map map_{br}

$$map_{br} = map[t = \text{obstacle}]$$

where $map[t = \text{obstacle}]$ represents a remapping of map such that tile t is now inaccessible. Therefore, tile t is a dominant tile for map if the modified floodplain is more than one tile smaller than the unmodified floodplain.

$$floodplain(t, map) < floodplain(t_{nb}, map_{br}) + 1$$

Since we have only blocked off one tile in map_{br} , if the floodplain is more than one tile smaller then blocking t must also have blocked access to other tiles. In other words, t dominates at least one tile. In general, fragmentation controls the number of obstacles in a map, while domination controls the openness of the play areas. These map-generation metrics are simple, but include a range of arcade-style archetype maps varying from open arena-based games, such as *Pong* or variants of *Snake* with minimal obstacles, to the more labyrinthine maps with maze-like properties such as those used by *PacMan* or *Sokoban*.

2) *Layouts*: Layouts are scored using sparseness and volume as metrics. This space is defined thus: a layout is sparse if the average distance between two NPCs is high. For a layout with n items in it (of any colour) the sparseness of a layout is measured as:

$$2 \times \left(\sum_{i=1}^n \sum_{j=1}^n dist(i, j) \right) \times \frac{1}{(n \times n - 1)}$$

which assumes the existence of a *dist* function, measuring the Euclidean distance between two items. Volume is the measure of how many NPCs there are on-screen. For a map m , its volume is defined as:

$$vol(m) = |reds(m)| + |blues(m)| + |greens(m)|$$

where $colour(m)$ is the list of items in map m of that colour. The layouts also compare themselves with current map instantiations to ensure that they are producing legal, playable layouts - that is, layouts which do not place NPCs over obstacles such as walls. In this way, the layouts and the maps shape each others' development.

3) *Rulesets*: We score rulesets using a variety of heuristics and checks. First, we run a series of checks to filter out the most unplayable or useless rulesets. We do this by examining the rulesets directly and checking for pathological combinations of rules. For instance, a ruleset is less fit if ways of increasing score are linked to player death, or if

rules overlap in ways that would cause a player to lose score; similarly, a game should never specify a rule for a NPC type if that type is entirely absent in the game layout.

In deciding on what form these checks should take, we took inspiration from McGonigal's definition [12] of a game's components as a *goal*, a set of *obstacles* and a *feedback mechanism*. In particular, we ensure the games are directed towards a goal using score gain and provide an obstacle to the player through the loss of score or death. These guidelines shrink the state space for rulesets, whilst retaining a core space of 'interesting' rulesets that can be explored through external playouts.

E. External Fitness For Strands Via Playouts

A *playout* is a simulated execution of a game with an algorithm taking the place of the human player. The external fitness functions for all three strands in ANGELINA use playouts as their source of data, as they provide a perfect model for how the three strands interact and perform together.

We characterise several types of playout based on the behaviour of the algorithm controlling the human player character. The types of playout are:

- *Empty*, where the player character is removed from the game and the game is executed until the time limit. This gives useful information on the game's natural entropy. For a game like *Pong*, this would result in a large negative score, while a game like *Pacman* would return a zero score.
- *Static*, where the player is included in the game but does not move. This might be considered a step up from the Empty playout, and provides a similar amount of 'null state' information.
- *Random Walker*, where the player mimics the NPC behaviour type of random walk. The player moves through the environment, turning at randomised intervals.
- *Pre-Collision*, which are identical to Random Walk playouts, except that the player is collided with all NPCs of one or more types before the play begins. These playouts are useful to compare to Random Walks as they show what the extremes of player involvement are. A game like *Pacman* yields very useful information from these playouts, as they include a perfect run (colliding with all the pills before the game begins completes the level) and a worst-case (colliding with the ghosts before the game begins, which results in the game over screen).
- *Long Play*, which are also identical to Random Walk playouts, but they are not limited by time.
- *True Player*, which uses A* search and the game's ruleset to seek out ways of increasing score and avoid losing score.

Playouts record information such as the player score, time taken to complete the level (or if the level was not completed), where and when the player died and what rules were fired. This information is used in individual ways by each strand to evaluate how members of their population interact with the larger game.

For example, the Ruleset strand will compare the score and survival data of the more ‘intelligent’ playouts such as True Player with the simpler Empty and Static playouts. We expect to see an improvement in the player’s performance as the playout type becomes more sophisticated, and evidence of this improvement results in a higher fitness for the ruleset. By contrast, rulesets for which the playout player struggles under all playout types may be considered too hard, and are therefore penalised.

III. EXPERIMENTATION AND RESULTS

We conducted several experiments to test the limits of ANGELINA’s ability to design games both assisted and alone, and in this section we introduce each of these and examine the results arising.

A. Density Of The Game Design Space

First, we introduce a preliminary experiment to help understand the scale of the state space. We produced an initial population of 100 randomly created games. If a game was playable, with a consistent map and layout and a goal-oriented ruleset, we classified them as Playable. If the games were not playable, they were classified with further detail on what factors were lacking. The results are given in Fig. 5.

Playable	3
Illegal Map/Layout	100
Broken Rules	74
No Goal	69
No Obstacles	54

Fig. 5: Classification of 100 random games.

We define ‘playable’ through the same judgements we apply in our fitness functions, along with further evaluation through human playouts. Each game was compiled and played by hand, and games which had clear goals and obstacles, and in which the player had some purpose or way to affect the system, were deemed playable. All of the games, however, had conflicting layouts and maps. Every game generated had NPC elements that were either embedded into walls or cut off from the rest of the map. We allowed a game to be declared playable if the player character was legally placed, although other game items may be placed illegally. As long as the goal remained achievable - that is, one or more of the goal items are not placed illegally - then the game as a whole is technically playable.

74 of the games examined had broken rules. These included rules that can never fire, such as a rule for a collision between two wall entities, or a rule that makes no sense given the current game design, such as a rule for green NPCs if the layout does not contain any. Further to this, 69 of the games had no way for the player to gain score, or avoid losing score, and 54 of them had no challenge for the player, either through death or the loss of score. Often, failing one of these checks was enough to stop a game being playable. However, most games failed three or more checks.

B. Independent Game Design

To investigate whether the whole is more than a sum of its parts, we first present games developed by ANGELINA with the system running without interaction or awareness between the individual evolutionary strands. What results are games that are composed of uniquely fit subcomponents, but that are not considered fit as overall games. Fig. 2 shows a sample game from these runs. The game’s rules, roughly, challenge the player to reach a red object. Touching the red object kills the player, and is the only way to gain score.

We see that each component complies with its own fitness constraints. For instance, the map in Fig. 3 exhibits a high domination score, creating a maze-like layout. The rulesets also conform to the loose common-sense fitness function devised for them. Despite this, the game in Figure 2 is unplayable, because, the map and layout conflict with each other by overlaying NPCs with wall tiles. This means those characters are unable to move or may be inaccessible to the player. The ruleset also includes two problems in particular - one of the rules, for collisions between green and blue, will never be used because the layout does not include any green NPCs. Another rule, the only one which allows the player to gain score, causes the player to die. This means that there is no way for the player to achieve the game’s goal. Recalling our definition of a game as a goal, a series of obstacles, and a feedback mechanism, a broken ruleset such as this provides incorrect obstacles, and does not have any path to the goal for the player to pursue.

C. Flexibility And Responsiveness

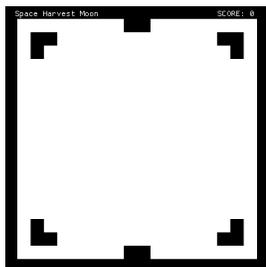
Fig. 6 shows a game designed with human input. We presented ANGELINA with a human-designed map and layout and asked it to derive a ruleset. The top images show an open map with a few obstacles that impede movement. The game ANGELINA designed exhibits rules relating to the placement of Blue non-player characters, fixes goals and obstacles for the player, and chooses movement behaviours for the various non-player character types. The player must avoid the moving blue NPCs and touch the walls to score.

We then changed the map and layout to that shown in the lower images in Fig. 6. This map and layout is less open than the previous one with a similar NPC arrangement. The fittest ruleset derived is much different to the previous example; the player must stay around blue objects to gain score, and avoid the fast-moving and unpredictable red objects. These objectives are more about awareness and navigation than agility, which fits with a more confined space and unpredictable opponents, since the red NPCs now teleport erratically. The module responsible for ruleset design knows, through playouts and analysis of NPC populations, that the space of ‘good’ games has changed because of the new level design, and the rulesets produced change appropriately. In particular, the second game designed does not use collisions with obstacles as a rule, because the higher number of obstacles would make the game too difficult (or too easy, depending on the effect of colliding with the obstacles). In-

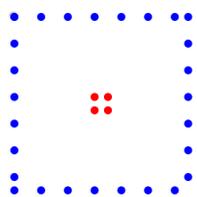
```

<scorelimit>55</scorelimit>
<timelimit>91</timelimit>
<redmovement>CLOCKWISE</redmovement>
<greenmovement>RANDSHORT</greenmovement>
<bluemovement>RANDLONG</bluemovement>
<rules>
<rule>PLAYER, OBSTACLE, TELEPORT, NOTHING, 1</rule>
<rule>BLUE, PLAYER, TELEPORT, TELEPORT, -1</rule>
</rules>

```



(a) The game map

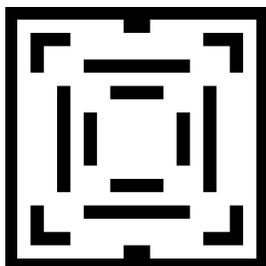


(b) The game layout

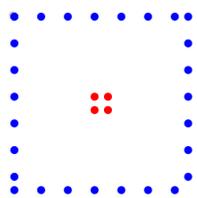
```

<scorelimit>89</scorelimit>
<timelimit>21</timelimit>
<redmovement>RANDLONG</redmovement>
<greenmovement>CLOCKWISE</greenmovement>
<bluemovement>CLOCKWISE</bluemovement>
<rules>
<rule>PLAYER, BLUE, NOTHING, NOTHING, 1</rule>
<rule>PLAYER, RED, DEATH, DEATH, 0</rule>
<rule>PLAYER, NONE, DEATH, TELEPORT, -1</rule>
<rule>OBSTACLE, RED, NOTHING, TELEPORT, -1</rule>
</rules>

```



(c) The game map



(d) The game layout

Fig. 6: Two games designed by ANGELINA, showing a response to the change in level design.

stead, the proposed game takes advantage of the more maze-like approach by developing a game with an emphasis on chase and evasion. ANGELINA's responsiveness to change in one element here has resulted in a dramatically different result that shows an awareness of the game's design as a whole being exhibited by our system.

D. Free Design

Fig. 7 shows two games designed by ANGELINA without constraints. The first game is a close relative of *Pacman* - the player is encouraged to collide with blue objects (which are destroyed, and give the player score) and avoid red objects (contact with them causes death). The map designed is not as symmetrical as *Pacman*'s, but is open enough for the player to manoeuvre around. This is a common arcade archetype that ANGELINA discovered within the state space independently.

The second game, shown in Fig. 6, is an arcade game equivalent of the 'steady hand game', where the player must avoid touching walls and pick up all NPCs of a certain colour. This map is designed as more confined, with less free space to manoeuvre. Its layout is more spread out and contains only one type of NPC. These two games show how ANGELINA explores and refines the search space open to it, and creates games that make sense *as a whole*, with the various subcomponents complementing each other. The Steady Hand Game is an interesting contrast to the *Pacman* clone, as it is not a rediscovery of a famous arcade game, yet is recognisable as a coherent game.

IV. RELATED WORK

A. Automated Game Design

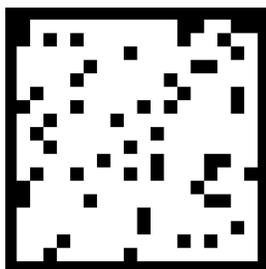
There is relatively little research which considers the problem of automated game generation as being any more than a series of linear content generation problems. Togelius et al. [10] explored the possibility of ruleset generation, using neural networks to evaluate their rulesets according to a learning-based model of fun similar to that described in *A Theory Of Fun* [13]. The rulesets generated by this system define two-dimensional arcade games where coloured shapes are controlled around a map. Although the work made some useful conclusions about its choice of model for game quality, the scope of the project was limited to ruleset design and did not attempt to generate other game components, such as levels or control schemes. However, Togelius' use of an independent layout module to assess rulesets can be seen as a move towards a more aware game designer, as it allowed the system to reflect on its own work as it progressed.

Smith and Mateas discuss automated game design in their presentation of *Variations Forever* [9], a research project into the procedural generation of 'minigames'. The authors use *answer set programming* ([14], [15]), a form of logic programming, to represent a large state space of possible rulesets, and then constrain the space to explore interesting subspaces of rules. This work only generated sets of rules, but is particularly notable because of the discussion of automated game design found within. In particular, the authors propose further work to develop their system with an automated playtester that could evaluate rulesets and use a knowledge base of player behaviour to suggest alterations to game designs. However, the discussion does not extend beyond the realm of designing simple rulesets, in a 'generate and test' paradigm. In contrast, we aim to produce a system

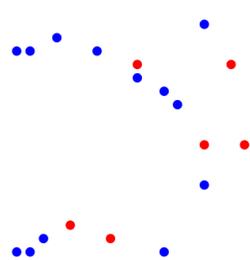
```

<scorelimit>56</scorelimit>
<timelimit>28</timelimit>
<redmovement>CLOCKWISE</redmovement>
<greenmovement>STATIC</greenmovement>
<bluemovement>RANDLONG</bluemovement>
<rules>
  <rule>PLAYER, RED, DEATH, DEATH, -1</rule>
  <rule>PLAYER, BLUE, NOTHING, DEATH, 1</rule>
</rules>

```



(a) The game map



(b) The game layout

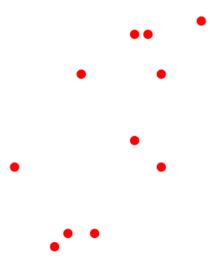
```

<scorelimit>24</scorelimit>
<timelimit>96</timelimit>
<redmovement>STATIC</redmovement>
<greenmovement>RANDLONG</greenmovement>
<bluemovement>RANDSHORT</bluemovement>
<rules>
  <rule>OBSTACLE, PLAYER, NOTHING, DEATH, 0</rule>
  <rule>PLAYER, RED, NOTHING, TELEPORT, 1</rule>
  <rule>RED, OBSTACLE, TELEPORT, NOTHING, 0</rule>
</rules>

```



(c) The game map



(d) The game layout

Fig. 7: Top: A Pacman-style game where the player must pick up green objects while avoiding red NPCs. Bottom: A 'steady-hand' game designed in an unconstrained mode. This game was not in the archetypal arcade game design space, but is an interesting invention by the system.

capable of understanding the connectedness of *multiple* game components, which means moving beyond rulesets alone.

Browne and Maire [16] explore the problem of automatic game design for the world of boardgames, a medium with many analogues to arcade videogames. The work here closely tackles the problem of automated game design by developing not only rulesets, but board shapes and starting layouts also. The work led to the design of many boardgames, some of which achieved commercial success. Although similar in structure, videogames have many components that boardgames lack. Most notably, the majority of videogames are played in real-time, which makes the task of playing out, and therefore evaluating, candidate games much more difficult due to a larger state space of game executions. Similarly, videogames typically derive their challenge from the presence of computer-controlled non-player characters. Assessing the challenge represented by these characters is harder than considering the challenge represented by binary rules, which in turn makes games harder to evaluate.

B. Evolutionary Computation

Evolutionary computation has been used in many areas of content design for videogames both in research and commercial videogame projects. Often, these projects employ a straightforward, linear Darwinian evolution method where a population of candidate artefacts are evaluated against a single fitness function and then high-scoring artefacts are crossed over to generate a new, fitter population.

Two relevant techniques which have emerged out of research into evolutionary computation are *multi-objective evolution* (MOE) and *co-evolution*. MOE differs from standard, linear evolution by attempting to optimise more than one fitness function for a single evolutionary strand. These fitness functions typically conflict in some way, and forms of Pareto optimality are often employed to find maximally-satisfying artefacts through evolutionary search. An example of such a system, along with a succinct explanation of its operation, can be found in [17].

Co-evolution refers to a naturally-occurring evolutionary system where two organisms evolve in response to changes in each other. This is often found in relationships between insects and plants or larger animals, comparable to symbiotic relationships. Programmatically, co-evolution has been used to develop two or more artefacts together with one another [18], [19]. Two standard evolutionary strands work alone, and in order to assess their fitness, they are composed into a single artefact and evaluated.

Standard evolution and other genetic programming techniques are well-applied to the world of content generation ([3], [8]). However, these techniques work largely in isolation and have no way of interacting with larger systems, and there is no sense of a guided design process.

Co-evolution is the technique closest to ANGELINA's underlying methods, but while co-evolution is primarily concerned with change in one object *causing* change in another, we are more interested in the evolutionary pressure applied by changes in multiple components. To continue

the biological analogy, while co-evolution is concerned with close relationships between two species, we are more interested in the ecosystem as a whole, and how the evolution of one strand can change the objectives of other strands.

V. FURTHER WORK

While the current space of possible games is broadly inherited from [10], it is not expressive enough to further test and demonstrate the capabilities of an autonomous game designer. The core of ANGELINA will be expanded to include a modular rule system and a search space which contains a larger variety of game types. Many arcade games are developed using a very small cache of rules. For instance, *Frogger* is a visually distinct game from *Pacman*, but only includes one new rule (a win-condition for being within a defined area of the map). A larger state space would improve ANGELINA's chances of finding new games. The inclusion of unusual rules, such as the ability for NPCs to change their type from one colour to another, would hopefully allow exploration of a less well-known state space and uncover arcade designs that are not so well known, yet enjoyable.

We believe that our approach to the evolutionary design of multi-faceted artefacts, through self-aware evolutionary strands that interact with each other, is applicable to many complex design tasks besides the design of videogames. We are interested in generalising the techniques we have begun to explore here, and developing a framework for creative systems which concern themselves with creating complex, interconnected artefacts. We are interested in applying these techniques to tasks such as the editorial design of magazines and newspapers, where space restrictions and aesthetic considerations conflict in interesting ways, and the construction of narratives, where there is conflict between the strands of narrative, the histories of characters, and the simulation of the world in which the plot is set. We hope that multi-faceted evolution's approach to complex artefact generation will provide a new way of considering these problems.

One advantage of generalising such a framework would be the simpler construction of tools that use evolutionary techniques to assist humans in creative tasks. As we demonstrated above, ANGELINA can develop entire games from scratch, but can also work from partially-specified starting points. We can envisage a computer-aided design tool that uses multi-faceted evolution to fill in the blanks in a creative design that a person has begun, be it a videogame design or otherwise.

VI. CONCLUSIONS

We have presented a new approach to generating arcade-style game designs through evolution, by maintaining individual evolutionary processes that can inspect and affect the fitness functions of other strands as they evolve. We have demonstrated this approach through the ANGELINA system, a game designer that produces simple arcade games from scratch by combining rulesets, maps and object layouts that it has designed itself. We demonstrated how this forms a single process, through multi-faceted evolution.

We showed that individually fit game components cannot be blindly composed together, and showed how ANGELINA's method for evolving these components in tandem with one another is effective and flexible. We discussed examples of these games and showed how ANGELINA was able to rediscover old arcade game archetypes, as well as reveal game designs that were novel and unexpected. Finally, we outlined our plans for a generalisation of the techniques in this paper to design better games, and other creative artefacts.

Some of the games compiled by ANGELINA have been put online, including some outlined in this paper². We hand selected these from a dozen runs of ANGELINA as being fun or offering an interesting challenge. We hope to expand these in the future, as ANGELINA's capabilities broaden.

VII. ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their commentary on the initial paper submission, which led us to add further detail to this paper.

REFERENCES

- [1] J. Togelius, M. Preuss, N. Beume, S. Wessing, J. Hagelbäck, and G. N. Yannakakis, "Multiobjective exploration of the starcraft map space," in *IEEE Conference on Computational Intelligence and Games*, 2010.
- [2] N. Sorenson and P. Pasquier, "Towards a generic framework for automated video game level creation," in *EvoGAMES*, 2010.
- [3] D. Ashlock, "Automatic generation of game elements via evolution," in *IEEE Conference on Computational Intelligence and Games*, 2010.
- [4] G. Smith, M. Treanor, J. Whitehead, and M. Mateas, "Rhythm-based level generation for 2d platformers," in *Proceedings of the 4th International Conference on Foundations of Digital Games*, 2009.
- [5] 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*, 2011.
- [6] A. Martin, A. Lim, S. Colton, and C. Browne, "Evolving 3d buildings for the prototype video game subversion," in *EvoGAMES*, 2010.
- [7] E. J. Hastings, R. K. Guha, and K. O. Stanley, "Evolving content in the galactic arms race video game," in *Proceedings of the 5th international conference on Computational Intelligence and Games*, 2009.
- [8] C. Lim, R. Baumgarten, and S. Colton, "Evolving behaviour trees for the commercial game DEFCON," in *EvoGAMES*, 2010.
- [9] A. M. Smith and M. Mateas, "Variations forever: Flexibly generating rulesets from a sculptable design space of mini-games," in *IEEE Conference on Computational Intelligence and Games*, 2010.
- [10] J. Togelius and J. Schmidhuber, "An experiment in automatic game design," in *IEEE Conference on Computational Intelligence and Games*, 2008.
- [11] M. J. Nelson and M. Mateas, "Towards automated game design," in *Artificial Intelligence and Human-Oriented Computing*. Springer, 2007, pp. 626–637.
- [12] J. McGonigal, *Reality is Broken*. Jonathan Cape, 2010.
- [13] R. Koster, *A Theory Of Fun For Game Design*, 2004.
- [14] Y. Babovich, E. Erdem, and V. Lifschitz, "Fages' theorem and answer set programming," in *In Proc. NMR-2000*, 2000.
- [15] F. Fages, "Consistency of clark's completion and existence of stable models," *Methods of Logic in Computer Science*, vol. 1, 1994.
- [16] C. Browne and F. Maire, "Evolutionary game design," in *IEEE Transactions on Computational Intelligence in AI and Games*, 2010.
- [17] P. Saetrom and M. Hetland, "Multiobjective evolution of temporal rules," in *8th Scandinavian Conference on Artificial Intelligence*, 2003.
- [18] C. Nachenberg, "Computer virus-antivirus coevolution," *Commun. ACM*, vol. 40, January 1997.
- [19] J. R. Koza, "Evolution and co-evolution of computer programs to control independently-acting agents," in *Proceedings of the first international conference on simulation of adaptive behavior on From animals to animats*. Cambridge, MA, USA: MIT Press, 1990.

²<http://www.gamesbyangelina.org>