Darwin's Avatars: a Novel Combination of Gameplay and Procedural Content Generation

Dan Lessin ITU Copenhagen Copenhagen, Denmark dles@itu.dk

ABSTRACT

The co-evolution of morphology and control for virtual creatures enables the creation of a novel form of gameplay and procedural content generation. Starting with a creature evolved to perform a simple task such as locomotion and removing its brain, the remaining body can be employed in a compelling interactive control problem. Just as we enjoy the challenge and reward of mastering helicopter flight or learning to play a musical instrument, learning to control such a creature through manual activation of its actuators presents an engaging and rewarding puzzle. Importantly, the novelty of this challenge is inexhaustible, since the evolution of virtual creatures provides a way to procedurally generate content for such a game. An endless series of creatures can be evolved for a task, then have their brains removed to become the game's next human-control challenge. To demonstrate this new form of gameplay and content generation, a proof-of-concept game—tentatively titled Darwin's Avatars—was implemented using evolved creature content, and user tested. This implementation also provided a unique opportunity to compare human and evolved control of evolved virtual creatures, both qualitatively and quantitatively, with interesting implications for improvements and future work.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning—connectionism and neural nets; I.6.8 [Simulation and Modeling]: Types of Simulation—animation; I.2.1 [Artificial Intelligence]: Applications and Expert Systems—games; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—animation

Keywords

evolved virtual creatures; artificial life; muscle drives; physicsbased character animation; procedural content generation

GECCO '15, July 11 - 16, 2015, Madrid, Spain

© 2015 ACM. ISBN 978-1-4503-3472-3/15/07...\$15.00

DOI: http://dx.doi.org/10.1145/2739480.2754749

Sebastian Risi ITU Copenhagen Copenhagen, Denmark sebr@itu.dk



(a) User view of the game.



(b) Volunteers playtesting the human vs. human version of the game.

Figure 1: Proof-of-concept game implementation. See video at http://youtu.be/p42wPQ3FYLU .

1. INTRODUCTION

This paper describes a new form of gameplay and a corresponding method of procedural content generation [17, 19, 5] which are made possible by evolutionary computation. The fundamental gameplay mechanism is a complex nonintuitive problem: How to control a physically simulated three-dimensional creature by manually activating its muscles.

The core content of the game—the creature to be controlled is generated procedurally using a muscle-driven evolved virtual creature (EVC) system like the one described in Lessin et al. [6]. Due to this close coupling between the core gameplay mechanics (learning to control an unfamiliar creature) and the associated method for procedural content generation

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

(evolving virtual creatures), the game's entertainment value is inherently vibrant, novel, and open ended. The implementation, mechanics, and procedural generation of content for the game—referred to here as *Darwin's Avatars*—are described in Section 3.

An additional benefit of this game system is that it permits the co-evolved controllers for these creature bodies to be compared to human control techniques. By seeing the ways in which human and evolved controllers differ (both in strategy and ability), we may be able to produce better evolved controllers. In Section 5, the results of such a comparison (both quantitative and qualitative) are presented, using data collected across multiple trials, users, and creatures. Useful implications of these comparisons are discussed in Section 6.

2. BACKGROUND

Before the presentation of the new system, relevant related work–both in terms of implementation and gameplay–will be described in this section.

2.1 EVC System

The EVC method at the heart of this new system of gameplay and content generation is an adaptation of Sims' original work [16], with a few small changes as well as some novel features. That method is described next.

Evolutionary Algorithm

A conventional evolutionary algorithm is used, with elitism, fitness-proportionate selection, and rank selection [8]. In addition, the most challenging tasks employ shaping [20, 3]. Fitness is evaluated in a physically simulated virtual environment implemented with NVIDIA PhysX [10].

Morphology

As in Sims' original work, creature morphology is described by a graph-based genotype, with graph nodes representing body segments, and graph edges representing joints between segments (Figure 2). By starting at the root and traversing the graph's edges, the phenotype is expressed. Reflexive edges as well as multiple edges between the same node pair are allowed, making it possible to define recursive and repeated body substructures easily. In addition, as in Sims' work, body symmetry is made readily available to evolution, with only a single mutation required to produce it. In this implementation, all PhysX primitives are available for use as body segments: boxes, spheres, and capsules. Joints between segments may be of most of the types offered by PhysX, specifically: fixed, revolute, spherical, prismatic, and cylindrical. In contrast to the typical technique of separately evolving explicit joint limits, most limitations on joint movement are provided implicitly by creature structure through natural collisions between adjacent segments.

In addition to the typical segments and joints, the implementation of the underlying EVC system also evolves muscle drives, described next.

Muscles

In a break with traditional evolved virtual creature systems, which typically use forces exerted directly at joints, the underlying EVC system of this paper uses simulated muscles as actuators [7]. Muscles are rendered as lines (as seen in Figures 4-6), with color from white to red indicating the degree



(a) Simple topology.



(b) Multiple edges for repeated substructures.



(c) Reflexive edge for recursive structure.



(d) Multiple and reflexive edges together.



(e) Multiple reflexive edges.

Figure 2: Hand-designed genotype/phenotype pairs (after those of Sims [16]) demonstrate the encoding power of EVC systems like the one employed in this paper.

of activation, from zero to full, respectively. Each muscle is defined by two attachment points on adjacent segments, along with a maximum strength value. In simulation, the muscle is implemented as a spring, with muscle activation modifying the spring constant. In addition to acting as an effector, each muscle also produces a proprioceptive feedback signal based on its current length. For each muscle, two nodes are added to the brain: one that accepts an input to set the muscle's activation, and another that makes the muscle's proprioceptive output signal available to the rest of the brain.

Control

In a manner which is again very similar to that of Sims, creature control is provided by a brain composed of a set of nodes connected by wires (as in Figure 3). Nodes receive varying numbers of input wires, and use their inputs to compute an output value (always in the range [0,1]) which may be sent to other wires. Signals originate from sensors in the

body as well as certain types of internal brain nodes, travel through the network of internal nodes and wires, and ultimately control the operation of actuators (muscles) in the physically simulated body. For each step of physical simulation, control signals move one step through the brain.



Figure 3: An EVC control network, co-evolved with a body for locomotion.

In addition to the special node types for muscles and their associated proprioceptors (described above), the following node types are allowed: sinusoidal, complement, constant, scale, multiply, divide, sum, difference, derivative, threshold, switch, delay, and absolute difference.

2.2 Related Work

The gameplay presented here, while novel, is not without precedent. The most obvious and closely related example is Bennett Foddy's online game $QWOP^1$ [2]. In QWOP, the user is presented with a two-dimensional physically simulated human runner whose muscles are controlled by pressing Q, W, O, and P on the keyboard. This work is useful in that it demonstrates the viability of gameplay which uses manual keyboard control of physically simulated muscles. The primary difference between the new system and QWOP is that this new control challenge is deeper and open ended, since it takes place in an unlimited supply of novel creatures produced through evolution. In addition, this evolutionary process provides viable automatic control systems for the creatures, making them available as competitors or teaching models.

Also of note, *Incredipede*² demonstrates the rewarding gameplay challenge of non-intuitive creature control (again in 2-D), although in this case, morphology can vary. This variability of body and control are at the heart of that game's entertainment value, serving to demonstrate the worth of the same gameplay challenge in the new system presented here. In Incredipede, however, morphological novelty is added by the player through a manual creature construction process, in contrast to the automated mechanism of procedural content generation described in this paper.

With respect to human vs. evolved control of physically simulated creatures, QWOP was recently the subject of such a comparison [12]. While not yet able to compete with human QWOP champions [11], their system is described by the authors as the first autonomous evolution of successful QWOP gaits.

When applied to games, procedural content generation (PCG) allows game elements (e.g. maps, textures, items, quests, etc.) to be generated algorithmically rather than through direct human design [19, 5]. This approach can reduce design costs and can also benefit players by providing them unique experiences every time they play. In particular, evolutionary computation and other search-based approaches [19] can enable the design of new content outside the scope of a fixed space of rules.

An example of PCG applied to video games is Galactic Arms Race (GAR [4]), in which weapons are evolved automatically based on user behavior. Further examples include Avery et al. [1], who evolved several aspects of a tower defense game, Shaker et al. [15] who evolved levels for the platform game Super Mario Bros., Risi et al. [13], whose game *Petalz* allows players to breed an unlimited variety of virtual flowers, and Togelius et al. [18], who experimented with evolving the rules of the game itself.

While the potential of PCG in competitive gaming is now well established [4, 9, 14], this paper presents the first example of PCG-generated EVCs for game content. The promise of this new approach is that it could supply a novel and unlimited stream of interesting creatures for the player to master.

3. THE GAME: DARWIN'S AVATARS

This paper's primary contribution is a new and powerful combination of gameplay and content generation. The fundamental challenge of the game is to manually control the muscles of a novel creature to achieve a certain task. In this proof-of-concept implementation, the challenge is a race, either against another human user, or against an evolved controller.

In this section, this system is presented in detail. First, the game's method of procedural content generation is described. Second, the mechanics of the game implementation are presented. Finally, the results of a human vs. human playtesting trial of the game are given.

3.1 Procedural Generation of Game Content

The game described above derives its entertainment value from the user's mastery of an unfamiliar control problem. In other games (e.g., QWOP and Incredipede) this challenging unfamiliarity is either fixed or the result of a user-controlled construction process. In contrast, this new game's corresponding procedural content generation is both open-ended and automatic, providing the user with an essentially inexhaustible source for the challenge at the heart of its gameplay.

To produce novel creatures for the user to control, an EVC system (Section 2.1) is run, with fitness corresponding to what will be the user's goal during gameplay. In the game described in this paper, for example, the creatures are the result of evolution for a locomotion task. Typically, ten runs might be started in parallel (with different random seeds), each with population size 100. These might evolve for 500 generations, with the winner from each run examined for

¹http://www.foddy.net/Athletics.html

²http://www.incredipede.com

both fitness success and aesthetic suitability. Approximately half of these runs might be expected to yield useful results.

Because the body is co-evolved with a successful control mechanism, two important aims are achieved: First, it is demonstrated that useful control of the evolved body is possible. If the body were merely evolved on its own, creatures might result which had no possibility of successful control. Second, the existence of an evolved control system for each creature makes a non-player character available, enabling a single user to compete against—or learn by example from the evolved brain.

3.2 Game Mechanics

Drawing upon the bodies evolved as described in Section 3.1, users are presented with three creatures, each with a different body plan, muscle structure, and likely locomotion style. The first creature (Figure 4) has three muscles, and typically moves by raising and lowering its heavy front limb to generate a fast jumping gait. The second creature (Figure 5) has four muscles controlling sphere limbs in a body which is symmetrical from front to back. This creature is more challenging to control, as its body makes it as easy to move backward as it is to move forward. It is also possible to tip over with this creature's morphology, although recovery seems always to be possible. The third creature (Figure 6) has the most complex body, with seven segments and six muscles, which are typically used in a familiar (and slower-paced) quadrupedal gait.



Figure 4: Creature type A: two segments, three muscles.



Figure 5: Creature type B: three segments, four muscles.

Once the creature is selected, the user is presented with the set of keys which will control their character. These keys are selected in a fixed yet essentially arbitrary order, with



Figure 6: Creature type C: seven segments, six muscles.

multiple players having a matching pattern of keys in the right and left halves of the keyboard. This non-intentional ordering of muscles and keys is appropriate to the fundamental gameplay challenge of learning a non-intuitive interface to creature control, and it may also be required if the game is to take advantage of the open-ended generation of content provided by the EVC system. Keys control muscle activation directly, with a key being up or down at any point in time corresponding to zero or full activation of a muscle, respectively. Keys controlling muscles function independently, allowing them to be activated in arbitrary overlapping patterns.

Two copies of the chosen creature are presented in a sideby-side pair of tracks (Figure 1a), with one creature controlled by player 1, and the other controlled by either player 2 or the creature's original evolved brain, as selected by the user. In this implementation, the tracks were straight, and creatures were prevented from leaving them by invisible physically simulated walls.

At a key press, the start timer counts down to zero, and the race begins. Time and distance are tracked for each creature, until both creatures have reached the end of the tracks or the time limit is reached. The time limits (10, 15, and 20 seconds for creatures A, B, and C, respectively) are based on the evolved controller's known speed for the current creature, with approximately double that time used as a limit in this implementation. When the race ends, each player's speed is presented. Speed is computed as positive or negative distance covered from the start point, divided by time to reach the end of the track (or the limit time, if the end of the track was not reached).

4. HUMAN VS. HUMAN USER EXPERIENCE

In preliminary human-vs.-human playtesting (four users, five matches with each of the three creatures, in the order A, B, C), the value of the proposed gameplay mechanics appear to have been validated. While some users may find direct competition against an evolved brain discouraging due to the mismatch in speed (see Section 5.2), most appear to enjoy the human-versus-human version of the game. As would be expected, such a one-on-one competition tends to reduce frustration, providing real hope for success as long as both players are at a similar level. For example, in the sequence of contests shown in Figure 7, the players are close enough in skill that user 012, while worse on average, is still within reach of user 011's speed, even surpassing it in the fourth run. During testing, many players displayed great curiosity about their scores and relative rankings, suggesting an interest in performance, accomplishment, and competition that bodes well for a motivated and rewarding gameplay experience.



Figure 7: A competitive sequence of five runs of human vs. human competition using the A-type creature.

Also of note, there is evidence for an increase in skill with practice, even over as little as the five runs allowed for each creature type during playtesting. Table 1 illustrates the change in mean speed between the first and fifth run for each creature, using data collected as described in Section 5. While only creature type C's increase is statistically significant (p < 0.01, using two-tailed, paired *t*-test), all three types show an increase in the mean that is at least suggestive.

Creature	Run 1	Run 1	Run 5	Run 5	р
Type	Mean	SD	Mean	SD	
A	0.4121	0.1651	0.6023	0.2986	0.1591
В	0.1988	0.1293	0.2424	0.0711	0.3930
С	0.1818	0.0543	0.3049	0.1265	0.0088

Table 1: Change in mean speed for each creature type between first and last run of playtesting, along with standard deviations and p (computed using a two-tailed, paired ttest).

This type of learning and accomplishment provides the player with satisfaction, increasing the player's enjoyment and the game's viability. And note that this core aspect of the game's entertainment value—the challenge and satisfaction of improving at a novel control task—is exactly what is extended by the game's corresponding procedural content generation solution.

5. HUMAN VS. EVOLVED-CONTROL RESULTS

In addition to producing the coupled game mechanics and content generation described above, the manual control of muscles in evolved virtual creatures provides a new way to compare human and evolved control of EVCs. In this section, both qualitative differences in control strategy and quantitative differences in resulting speed are presented.

These results are derived from single-player user tests by seven volunteers. In each user test, five runs were performed for each of the three creature types (first A, then B, then C), with both final speed and muscle-activation patterns recorded for each run.

5.1 Qualitative Comparison

For qualitative comparison of human and evolved controllers, patterns of recorded muscle activation are compared. A useful way to present such patterns is as a variation on a spectrogram, with each row corresponding to a single muscle's activation over time. In Figures 8-10, a spectrogram for the final user run of each creature type is presented (in green), along with a spectrogram from that creature's evolved controller (in red).



Figure 8: Creature type A: muscle-activation spectrograms, with corresponding speeds.

For the A-type creature (Figure 8), human control strategies are similar to that of the evolved controller. Users tend to employ the same two muscles as the evolved controller, and many also replicate its simple alternating pattern, as well. The user pattern most similar to the evolved solution (Figure 8d) was also the fastest human example for that run (1.1146 m/s), although still far from the evolved controller's mean for that creature (2.3281 m/s). Beyond differences in control technique, another reason for lower human scores which is indicated in these spectrograms is that many users do not begin activating muscles until a fraction of a second after the run begins, while the evolved controllers start without any delay.



Figure 9: Creature type B: muscle-activation spectrograms, with corresponding speeds.

For the B-type creature (Figure 9), most human activation patterns differ significantly, both from each other and from the evolved controller, and none are particularly successful compared to the evolved solution. Although its importance is not proven, this evolved controller makes the most obvious use of continuous activation values—something not available to a human user through the keyboard interface (as discussed in Section 6). This may be one reason why human users find this creature particularly challenging. Another difficult aspect of controlling this creature originates in its morphology: Because its body is symmetrical from front to back, it just as easy for it to move backwards as forwards, with some users even achieving a negative speed during evaluation runs.

For the C-type creature (Figure 10), there is again great variety in timing and muscles used, both among human and evolved results. In this creature, however, human control has finally begun to compete with evolved control. Interestingly for this creature, the human control pattern most visually similar to the evolved one (Figure 10b, 0.2983 m/s) is not nearly the fastest. Instead, it is the pattern of Figure 10c, using very different muscles and timing from the evolved



Figure 10: Creature type C: muscle-activation spectrograms, with corresponding speeds.

solution, which scores best, reaching 0.5606 m/s, beating the evolved controller's mean of 0.5515 m/s.

5.2 Quantitative Comparison

For quantitative comparison of human and evolved control of EVCs, the same human test data is examined with respect to locomotive speed. Figures 11-13 illustrate humancontrolled speed vs. evolved-controller speed for each of the three creature types. For each creature, data for seven test subjects over five runs are plotted, with mean, standard deviation, and best result indicated for each run. For comparison, seven runs of the evolved creature are similarly recorded and graphed. For each of the three creature types, evolved controllers score better in mean and best result during every run, but there are some important differences worth noting.

For creature types A and B, humans were not competitive with the evolved solution throughout the recorded trials. The best human score remained well below the evolvedcontroller mean for all such runs. Given the limited number of runs, however, and real-world examples of human controltask learning (consider learning to play the piano, or compete in Olympic gymnastics), it is not unreasonable to think that a human player might be competitive with an evolved controller, if given enough time to practice.

With creature type C, however, human performance was much closer to evolved performance throughout, with the best human controller once approaching the evolved-controller



Figure 11: Comparison of evolved and human controller speeds for creature type A.



Figure 12: Comparison of evolved and human controller speeds for creature type B.

mean, and once even exceeding it. In fact, the last run for this type showed both the highest mean and highest best score for human controllers, with a statistically-significant indication of increasing score between the first and last runs for this creature (as described in Section 4 and Table 1). With human results approaching evolved results and shown to be still increasing, C-type creature results seem to offer humanity's best hope for matching or surpassing evolved controllers for the virtual creatures examined here.

6. DISCUSSION AND FUTURE WORK

The results of Section 5's comparison between human and evolved controllers have some interesting implications, and suggest potentially valuable future work. One important observation is that humans may be able to produce usable (although likely slower) control strategies in a significantly shorter wall-clock time than evolution—on the order of minutes for humans versus hours or more for evolution. Another



Figure 13: Comparison of evolved and human controller speeds for creature type C.

important observation is that human and evolved controllers often use very different techniques, with even the most competitive human controller differing greatly in strategy from the evolved controller for the same creature. This suggests the potential for a productive transfer of knowledge from human user to evolution. For example, evolved solutions could be rewarded for their similarity to successful human play traces, just as a teacher might instruct a student on technique when learning a physical skill. The proper combination of human and evolved skill could potentially lead to better solutions than either might have achieved alone, and seems worthy of future study.

Another important topic worthy of discussion is the difference in possible activation patterns between human and evolved controllers in the system this paper describes. As in typical EVC systems, the evolved controllers studied here generate continuous activation signals, while human users are limited to a binary signal for each muscle, owing to the keyboard interface. For human-vs.-human competition, this limitation may be acceptable, since useful control was shown to be possible with such discrete input, and because both players have this constraint in common. For humanvs.-evolved competitions or comparisons, however, it is unknown if this limitation contributed to the observed difference in achievement between the two types of control. To resolve this issue, future evolved controllers could be forced to produce binary muscle activation signals by applying a threshold filter to their outputs. Then, identical human and evolved activation patterns would be possible, placing human and machine on equal footing in both gameplay and research comparisons.

Another interesting difference between human and evolved controllers is the input that each receives. Although not all EVCs in this implementation use it, they have the potential to employ proprioceptive signals based on the current length of their muscles. This information is not directly available to the human user. Similarly, human users have access to a visual representation of the creature in simulation, providing them with information not available to the evolved controller (e.g.: the orientation, position, and velocity of many body segments; or which parts of the creature are in contact with the ground). While it is not clear which of these might confer a larger advantage, or whether the differences are significant at all, it is at least worth noting the difference.

On the topic of future work, one important step would be more deeply exploring the potential for evolution of new creatures and online interaction. In a full version of the game, great value might be added by allowing users to evolve their own creatures, trade them with each other, and compete against each other online.

Another potentially valuable extension to this initial experiment would be the addition of different tasks. Any action that EVCs can be evolved to perform (jumping, climbing, swimming, fighting, etc.) might serve as a rewarding new control challenge for the game.

7. CONCLUSION

This paper has presented Darwin's Avatars—a novel combination of gameplay and procedural content generation made possible by evolutionary computation. Although similar physical control challenges have already proven successful in previous games, this new game adds the novel elements of threedimensional creatures and creatures which are unfamiliar without being constructed by the user. Preliminary playtesting showed that these gameplay mechanics have the potential to be entertaining and challenging, offering users the opportunity for learning and close competition. And importantly, this system's inherent ability to generate the core ingredient of its gameplay-novel control challenges, in its unlimited supply of newly evolved creatures-makes for a unique and compelling new combination. In addition, the game's core mechanic of human control of muscles in EVCs made it possible to compare human and evolved controllers, providing some initial insight into how these two types of controller compare numerically, how they may differ in terms of technique, and potential improvements and future work implied by these differences.

8. **REFERENCES**

- P. Avery, J. Togelius, E. Alistar, and R. van Leeuwen. Computational intelligence and tower defence games. In Evolutionary Computation (CEC), 2011 IEEE Congress on, pages 1084–1091. IEEE, 2011.
- B. Foddy. Qwop [flash game](2008). Technical report, Retrieved 11-6-13 from http://www. foddy. net/Athletics. html, 2013.
- [3] F. Gomez and R. Miikkulainen. Incremental evolution of complex general behavior. *Adaptive Behavior*, 5:317–342, 1997.
- [4] E. J. Hastings, R. K. Guha, and K. O. Stanley. Automatic content generation in the galactic arms race video game. *IEEE Transactions on Computational Intelligence and AI in Games*, 1(4):245–263, 2010.
- [5] M. Hendrikx, S. Meijer, J. Van Der Velden, and A. Iosup. Procedural content generation for games: A survey. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCCAP), 9(1):1, 2013.
- [6] D. Lessin, D. Fussell, and R. Miikkulainen. Open-ended behavioral complexity for evolved virtual creatures. In *Proceeding of the Fifteenth Annual Conference on Genetic and Evolutionary Computation*

Conference, GECCO '13, pages 335–342, New York, NY, USA, 2013. ACM.

- [7] D. Lessin, D. Fussell, and R. Miikkulainen. Trading control intelligence for physical intelligence: Muscle drives in evolved virtual creatures. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO) 2014*, 2014.
- [8] M. Mitchell. An Introduction to Genetic Algorithms. MIT Press, Cambridge, MA, USA, 1998.
- [9] C. Pedersen, J. Togelius, and G. N. Yannakakis. Modeling player experience for content creation. *IEEE Transactions on Computational Intelligence and AI in Games*, 2(1):54–67, 2010.
- [10] A. PhysX. Ageia, 82 pioneer way, suite 118, mountain view, ca 94041, 2008.
- [11] R. Ramachandra. Fastest 100m run, qwop (flash game). Guinnes World Records, 2013.
- [12] S. Ray, V. S. Gordon, and L. Vaucher. Evolving qwop gaits. In Proceedings of the 2014 conference on Genetic and evolutionary computation, pages 823–830. ACM, 2014.
- [13] S. Risi, J. Lehman, D. D'Ambrosio, R. Hall, and K. Stanley. Petalz: Search-based procedural content generation for the casual gamer. *Computational Intelligence and AI in Games, IEEE Transactions on*, PP(99):1–1, 2015.
- [14] N. Shaker, G. N. Yannakakis, and J. Togelius. Towards automatic personalized content generation for platform games. In *Proceedings of Artificial Intelligence and Interactive Digital Entertainment* (AIIDE-10). AAAI Press, 2010.
- [15] N. Shaker, G. N. Yannakakis, J. Togelius, M. Nicolau, and M. O'Neill. Evolving personalized content for Super Mario Bros using grammatical evolution. In Proceedings of the Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE 2012), Menlo Park, CA, 2012. AAAI Press.
- [16] K. Sims. Evolving virtual creatures. In Proceedings of the 21st annual conference on Computer graphics and interactive techniques, SIGGRAPH '94, pages 15–22, New York, NY, USA, 1994. ACM.
- [17] J. Togelius, A. J. Champandard, P. L. Lanzi, M. Mateas, A. Paiva, M. Preuss, and K. O. Stanley. Procedural Content Generation: Goals, Challenges and Actionable Steps. In S. M. Lucas, M. Mateas, M. Preuss, P. Spronck, and J. Togelius, editors, *Artificial and Computational Intelligence in Games*, volume 6 of *Dagstuhl Follow-Ups*, pages 61–75. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2013.
- [18] J. Togelius and J. Schmidhuber. An experiment in automatic game design. In Computational Intelligence and Games, 2008. CIG'08. IEEE Symposium On, pages 111–118. IEEE, 2008.
- [19] 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*, 3(3):172–186, 2011.
- [20] J. Urzelai, D. Floreano, M. Dorigo, and M. Colombetti. Incremental robot shaping. *Connection Science*, 10:341–360, 1998.