Interactive Super Mario Bros Evolution

Patrikk D. Sørensen, Jeppeh M. Olsen and Sebastian Risi
Center for Computer Games Research
IT University of Copenhagen
Copenhagen, Denmark
{pdyr, jmao, sebr}@itu.dk

ABSTRACT

Creating controllers for NPCs in video games is traditionally a challenging and time consuming task. Automated learning methods such as *neuroevolution* (i.e. evolving artificial neural networks) have shown promise in this context but they often require carefully designed fitness functions to encourage the evolution of desired behaviors. In this paper, we show how casual users can create controllers for *Super Mario Bros* through an *interactive evolutionary computation* (IEC) approach, without prior domain or programming knowledge. By iteratively selecting Super Mario behaviors from a set of candidates, users are able to guide evolution towards a variety of different behaviors, which would be difficult with an automated approach. Additionally, the user-evolved controllers perform similarly well as controllers evolved with a traditional fitness-based approach when comparing distance traveled. The results suggest that IEC is a viable alternative in designing complex controllers for video games that could be extended to other games in the future.

Keywords

Interactive Evolution, Neuroevolution, Game AI

1. INTRODUCTION

In recent years it has become more and more popular for video games to enable users to create and share game content. The user-created content most often comes in the form of levels; very few games let the user create or modify the behavior or underlying structure of the Non-Player-Characters (NPC). Usually, the NPC behaviors are constructed by programmers and function in predetermined and static ways.

In this paper we show that casual users are able to create sophisticated behaviors for the *Super Mario Bros* video game by using a simple interface to evaluate several candidate behaviors that is reiterated upon. The approach is based on *interactive evolutionary computation* (IEC) [3] and requires no prior knowledge of neither AI methods nor programming. The developed IEC framework presents users with a selection of animated GIFs of short level playthroughs, from which they choose a parent for the next generation.

GECCO '16 July 20-24, 2016, Denver, CO, USA

ACM ISBN 978-1-4503-4323-7/16/07...$15.00
DOI: http://dx.doi.org/10.1145/2908961.2908997

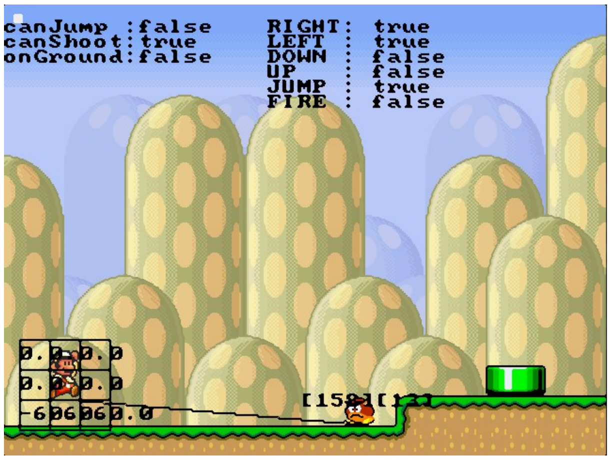

Figure 1: Mario Representation. The controlling ANN receives a 3 × 3 grid as input together with information about the distance and angles to enemies, and the conditional domain variables *canJump* and *onGround*. The ANN outputs (shown at the top-right corner) determine the action that Mario performs each tick.

The results in this paper show that casual users are able to not only interactively evolve a variety of interesting and unique behaviors but also behaviors that perform competitively in comparison to automated searches. Additionally, and potentially more important, users reported that they (1) had fun while evolving Mario behaviors, and (2) felt that they had an impact on evolution. These results indicate that IEC could be a viable and entertaining approach to empowering players to create their own NPC behaviors.

2. METHODS AND REPRESENTATION

Neural Network Setup. The neural networks controlling Mario are evolved with the NEAT algorithm [2]. Similarly to the setup in Togelius et al. [4], the ANN receives a 3 ×3 grid of cells centered around Mario as input, where different cell values represent different terrain types (Figure 1). Additionally, the ANN receives the distance and angle to the two nearest enemies. The ANN has six outputs: right, left, up, down, jump, and fire. If the output value for an action is higher than 0.5, the particular button is pressed.

The IEC Interface. While the user is watching, a total of nine controllers are playing – one after the other – through a small part of a Mario level. During these playthrougs, GIFs are recorded for each of the different controllers that show Mario in action. Once the whole population has been played and recorded, a window with all nine recorded GIFs is shown to the user. The user is then able to evaluate and compare each individual in the population and select one preferred controller by simply clicking on the particular GIF. Based on the user's selection, the next generation of controllers is created through mutating the selected individual and the process starts again. That way users can guide the evolutionary search towards Mario behaviors they find interesting.

3. EXPERIMENTS

To test whether users could evolve interesting Mario behaviors, a user study was performed on site at the IT University of Copenhagen with a total of 20 participants. Each participant was asked to evolve controllers for 20 generations through the IEC interface. Additionally, we divided the participants in two groups of ten. The first ten participants were encouraged to evolve whichever behaviors they preferred, while the other ten participants were told to evolve controllers that could travel as far as possible.

After the experiment the players were asked to answer two questions on a scale from 1–6 (where 6 is best): (1) *How much impact do you feel that your choices had on the evolutionary process?* (2) *How fun was it to develop your AI this way?*

Controllers evolved through IEC were compared to controllers evolved with a traditional automated fitness-based search approach. The fitness for the automated approach was the number of cells passed at the end of the simulation. The simulation was terminated if Mario reached the end of the level, he died or the time ran out. Additionally, to create more robust controllers, the starting position of the avatar was moved every four generations to a different position in the same level, for both the automated and IEC approach.

The population size was set to nine for both the automated and IEC approach. The number of generations was 20. Offspring had a weight mutation chance of 0.55, 0.01 chance of node addition, and 0.01 chance of link addition.

4. RESULTS

The IEC results are based on the ten participants of each experiment and the fitness-based results are collected from ten independent evolutionary runs. Figure 2 depicts a general trend for all approaches to improve over generations. Not surprisingly, for all three approaches there are drops in fitness when the starting point of Mario is moved in generations 5, 9, 13 and 17. However, there is a significant increase in performance for all methods comparing first and last generations ($p < 0.05$; Student's t-test). The pair-wise differences between the approaches are not significantly different, which indicates that both automated and IEC approaches are able to evolve similar performing Mario controllers.

Indeed, the participants in our user study were able to evolve controllers with a variety of different behaviors. One was rushing through the level, another was being very careful around enemies and yet another was shooting fireballs at every enemy it came across. The reader is encouraged to take a look at the video accompanying this paper (available

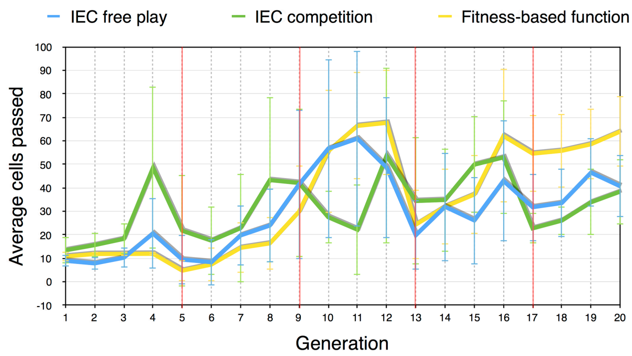

Figure 2: Training Performance Average cells passed during training over generations. The vertical red lines indicate the generations when Mario is set to a new position in the level.

at http://youtu.be/vHVgN_hFN2c), to get a better sense of the types of controllers that were evolved.

The results from the questionnaire show that the participants felt that they had significant impact on evolution with an average rating of 4.29 out of 6 (SD=0.8). Maybe slightly surprising, over 25% gave the maximum score when asked about the level of fun they had breeding Mario controllers, with a score of 4.88 on average (SD=1.0).

5. CONCLUSION

The presented approach allows users, for the first time, to interactively evolve behaviors for Super Mario. The controllers evolved with IEC perform similarly well compared to a fitness-based search in terms of distance traveled, but show more varied strategies and behaviors. Importantly, users of the IEC system reported that they had fun while evaluating and evolving Mario controllers. In the future, this system could be extended to other video games and to allow users to evolve behaviors collaboratively online. To encourage an even greater variety of NPC behaviors, an approach that combines a non-objective search such as *novelty search* [1] with IEC could be promising [5].

6. REFERENCES

[1] J. Lehman and K. O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.

[2] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.

[3] H. Takagi. Interactive evolutionary computation: Fusion of the capabilities of ec optimization and human evaluation. *PROCEEDINGS OF THE IEEE*, 89(9):1275 – 1296, 2001.

[4] J. Togelius, S. Karakovskiy, J. Koutník, and J. Schmidhuber. Super mario evolution. In *Computational Intelligence and Games, 2009. CIG 2009. IEEE Symposium on*, pages 156–161. IEEE, 2009.

[5] B. G. Woolley and K. O. Stanley. A novel human-computer collaboration: combining novelty search with interactive evolution. In *Proceedings of the 2014 conference on Genetic and evolutionary computation*, pages 233–240. ACM, 2014.