Exploiting Environmental Differentiation to Promote Evolvability in Artificial Evolution

Jonata Tyska Carvalho[1,2]
[1]Center for Computational Sciences (C3) - Federal University of Rio Grande (FURG)
Av. Italia, km 8
Rio Grande, Brasil 96203-900
jonatatyska@gmail.com

Stefano Nolfi[2]
[2]Institute of Cognitive Sciences and Technologies(ISTC) - National Research Council(CNR)
Via S. Martino della Battaglia, 44
Roma, Italia 00185
stefano.nolfi@istc.cnr.it

ABSTRACT

In this work we investigate the possibility to exploit environmental differentiation to promote evolvability in artificial evolution. More specifically we propose a new algorithm and demonstrate how agents evolved for the ability to solve the double-pole balancing problem in differentiated environmental conditions among the population outperform agents evolved in homogeneous environmental conditions. The algorithm operates by evolving the agents on multiple environmental niches with randomly varying environmental characteristics and by enabling agents displaying superior performance in other niches to colonize them. Agents evolved through the proposed algorithm outperform agents evolved in homogeneous environments, either on stable or temporally varying environments.

CCS CONCEPTS

•Computing methodologies → Genetic algorithms; Evolutionary robotics; •Theory of computation → Evolutionary algorithms; •Computer systems organization → Neural networks;

KEYWORDS

Evolutionary Robotics, Genetic Algorithms, Evolvability, Niches

ACM Reference format:
Jonata Tyska Carvalho[1,2] and Stefano Nolfi[2]. 2017. Exploiting Environmental Differentiation to Promote Evolvability in Artificial Evolution. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017,* 2 pages.
DOI: http://dx.doi.org/10.1145/3067695.3075994

1 INTRODUCTION

Notably, despite natural evolution clearly occurs in environments that vary over space and time, artificial evolution is normally realized in non-differentiated environments. Even in the cases in which the evolving population is divided into separated islands or groups (as in [2, 4, 5]) the environmental characteristics of the different islands/niches are maintained constant. In these works, the subdivision of the population into islands is introduced to preserve diversity among the population. However, the lack of environmental differentiation among islands prevents the exploitation of another important mechanism that promote population diversity that consists in the need to adapt to varying environmental niches (see [3]).

In this work we propose a new evolutionary algorithm that operates on multiple differentiated environmental niches. The need to adapt to specific niches promotes the development of specialist solutions that are adapted to their niche but that are not necessarily fit in other environmental conditions. On the other hand, the possibility for the individuals displaying good performance in different environmental conditions to colonize new niches promotes the development of more general solutions that operate successfully on multiple niches. The comparison of the results obtained in two control experiments in which the environment is undifferentiated, stable or varying randomly, demonstrates that environmental differentiation favors the evolution of better solutions.

2 METHOD

As experimental set up we chose the so-called double-pole balancing task [6], the problem consists in controlling a mobile cart with two poles of different length attached through passive hinge joints on the top of the cart for the ability to keep the poles balanced. We consider the non-markovian version of the problem proposed by [6] in which the agent controller does not have access to the speed of the cart and of the poles.

We used a neural network as the cart controller with three sensory neurons, 10 internal neurons with recurrent connections and one motor neuron, a sigmoid was used as the activation function for the internal and output neurons. The sensory neurons encoded the position of the cart, and the angular position of the two poles. Each agent was evaluated for a certain number of trials that differ with respect to the initial position of the cart and of the poles. The initial angle of the poles and the position of the cart were selected within [-0.2, 0.2]rad and [-2.4, 2.4]m, respectively. Trials lasted up to 1000 steps and were terminated when the angular position of one of the two poles exceeded the range [-36°, 36°] or when the position of the cart exceeded the range [-2.4, 2.4]m. The fitness function include two components that rate agents for the ability to maintain the poles balanced and to penalize oscillations that could lead to an exploitation of round off errors in the simulation, as proposed in [1]. The equations describing the dynamics of the system are the same used in [6]. The mass of the cart was 1kg, the

GECCO '17 Companion, Berlin, Germany
 978-1-4503-4939-0/17/07…$15.00
DOI: http://dx.doi.org/10.1145/3067695.3075994

long and short pole had a mass of 0.1kg and 0.01kg, and were 0.5m and 0.05m long respectively. The friction was set to 0.

To check whether environmental differentiation promotes the discovery of better solutions we carried out two series of the experiments in which the agents were placed in an undifferentiated environment, i.e. were evaluated from the same initial conditions that could change or not on each generation, and one in which the agents were placed into 20 different niches and were evaluated from different initial conditions in each niche. Since the adaptive problem is influenced by the initial states of cart and poles, we decided to use them to differentiate the ecological niches of the agents.

The evolutionary process for the undifferentiated experiments was done on the basis of the (μ+1) evolutionary strategy that uses μ parents, 20 in our case, enables each parent to produce a single mutated offspring, 4% mutation rate, and use the 20 best individuals among parents and offspring as new parents. The first series of experiment we called homogeneous was composed by a single niche with 200 randomly generated initial conditions that never changed. The second series we called random and was composed by a single niche with 10 conditions that changed on each generation. All individuals were placed in the same niche and so were evaluated on the basis of the same initial conditions.

For the heterogeneous environmental conditions, instead, the evolutionary process was carried out on the basis of a new algorithm that we developed, in which each of the 20 individuals forming the population was placed in a different corresponding niche. Each niche was characterized by 10 different randomly generated initial conditions that differed among niches, consequently each individual was evaluated in different environmental conditions. Individuals adapted in the basis of a (1+1) evolutionary strategy for 50 iterations, i.e. they generated an offspring, 4% mutation rate, that was used to replace the parent individual depending on whether it's fitness was equal or greater than that of the parent. The best mutated agent obtained during the 50 iterations was then evaluated in all other niches and was used to colonize the niche in which it displayed the greater improvement, with respect to the individual currently occupying that niche.

For each condition we ran 30 replications of the experiment and for each replication were done 20 millions evaluations (trials) in order to maintain the computational cost constant. The performance obtained in the different experimental conditions was measured by post-evaluating the evolved agents for 625 trials during which the initial position of the cart, the angle and velocity of the long pole varied uniformly as proposed in [1].

3 RESULTS

The measured performance refer to the number of trials in which the agents managed to balance the pole up to the end of the trial during the post-evaluation procedure described in the previous section. Notice that the optimal performance level is unknown but it is probably much lower than the theoretical maximum given that balancing the poles from certain initial states is likely impossible. As can be seen in the Figure 1 the agents evolved in the heterogeneous environment display better performance than those evolved in the undifferentiated environments, homogeneous and random.

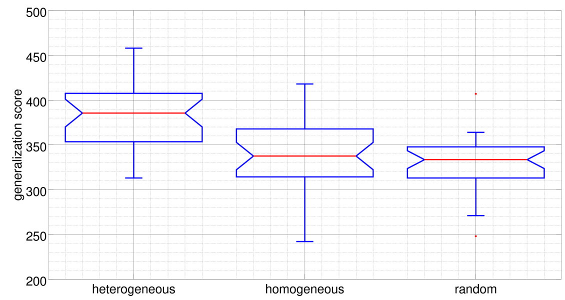

Figure 1: Performance in the three experimental conditions.

The differences between the first and the other two conditions are statistically significant (student's t-test, $p < 0.001$).

4 DISCUSSION

We investigated whether the presence of differentiated environmental conditions promotes the synthesis of better solutions in agents evolved for the ability to solve the double-pole balancing problem, a hard problem commonly used as a benchmark for the evolution of embodied agents. The experiments in the heterogeneous environment were carried out through a new evolutionary algorithm that we designed to explore the advantage of environmental differentiation. The algorithm operates by: (i) allocating each individual of the population in a corresponding niche that varies randomly with respect to other niches, (ii) periodically checking the presence of general solutions (i.e. agents that are able to display superior performance in niches that differ from the one in which they adapt), and (iii) enabling generalist individuals to colonize other niches. In other words the algorithm has been designed so to facilitate the occurrence of the incremental evolutionary process described above in which evolution proceeds by first generating specialist high-quality solutions and then progressively more generalist high-quality solutions.

ACKNOWLEDGMENTS

Partially funded by CAPES/Science without borders/Brazil.

REFERENCES

[1] Frederic Gruau, Darrell Whitley, and Larry Pyeatt. 1996. A Comparison Between Cellular Encoding and Direct Encoding for Genetic Neural Networks. In *Proceedings of the 1st Annual Conference on Genetic Programming*. MIT Press, Cambridge, MA, USA, 81–89. http://dl.acm.org/citation.cfm?id=1595536.1595547

[2] Gregory S. Hornby. 2006. ALPS: The Age-layered Population Structure for Reducing the Problem of Premature Convergence. In *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation (GECCO '06)*. ACM, New York, NY, USA, 815–822. DOI: http://dx.doi.org/10.1145/1143997.1144142

[3] R. Kassen. 2002. The experimental evolution of specialists, generalists, and the maintenance of diversity. *Journal of Evolutionary Biology* 15, 2 (March 2002), 173–190. DOI: http://dx.doi.org/10.1046/j.1420-9101.2002.00377.x

[4] D. Simon. 2008. Biogeography-Based Optimization. *IEEE Transactions on Evolutionary Computation* 12, 6 (Dec. 2008), 702–713. DOI: http://dx.doi.org/10.1109/TEVC.2008.919004

[5] Darrell Whitley, Soraya Rana, and Robert B. Heckendorn. 1999. The Island Model Genetic Algorithm: On Separability, Population Size and Convergence. *CIT. Journal of Computing and Information Technology* 7, 1 (March 1999), 33–47. http://hrcak.srce.hr/index.php?show=clanak&id_clanak_jezik=221148

[6] A. P. Wieland. 1991. Evolving neural network controllers for unstable systems. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*, Vol. ii. 667–673 vol.2. DOI: http://dx.doi.org/10.1109/IJCNN.1991.155416