Benefits of Lamarckian Evolution for Morphologically Evolving Robots

Milan Jelisavcic
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
m.j.jelisavcic@vu.nl

Rafael Kiesel
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
rafael.kiesel@web.de

Kyrre Glette
University of Oslo
Oslo, Norway
kyrrehg@ifi.uio.no

Evert Haasdijk
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
e.haasdijk@vu.nl

A. E. Eiben
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
a.e.eiben@vu.nl

ABSTRACT

Implementing lifetime learning by means of on-line evolution, we establish an indirect encoding scheme that combines Compositional Pattern Producing Networks (CPPNs) and Central Pattern Generators (CPGs) as a relevant learner and controller for open-loop gait controllers in modular robots which have evolving morphologies. Experimental validation on the morphologically evolved robots shows that a Lamarckian setup with CPPN-CPG provides substantial benefits compared to controllers learned from scratch.

CCS CONCEPTS

•**Computing methodologies** → **Evolutionary robotics;** *Mobile agents;* •**Theory of computation** → *Evolutionary algorithms;*

KEYWORDS

Evolutionary robotics, On-line evolution, Indirect encoding, Lamarckian evolution, Gait learning

ACM Reference format:

Milan Jelisavcic, Rafael Kiesel, Kyrre Glette, Evert Haasdijk, and A. E. Eiben. 2017. Benefits of Lamarckian Evolution for Morphologically Evolving Robots. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017,* 2 pages.
DOI: http://dx.doi.org/10.1145/3067695.3076046

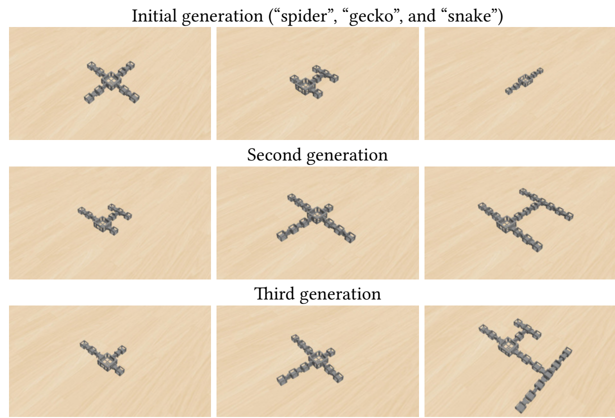

Figure 1: The top row shows the three hand-designed initial morphologies. The following rows of robots result from recombination and mutation of the previous generations.

1 INTRODUCTION

Evolutionary robotics offers a methodology to consider the development and adaptation of robot morphology and control holistically. While there is a relatively large body of research in evolutionary robotics considering evolution and adaptation of robotic control systems, the simultaneous development of robot morphologies and control systems is a difficult task, and we have only seen relatively simple results so far, as noted by Cheney et al. [1]. Some of the difficulty is due to the increased dimensionality of the search, but a more pernicious aspect may be the increased ruggedness of the search space: a small mutation in the morphology can easily offset the performance of the controller-body combination found earlier.

Here, we address this challenge by means of a lifetime learning approach for robot controllers, focussed on scenarios where robots' morphologies evolve. In particular, we show how the use of an *indirect* encoding can enable a Lamarckian evolutionary set-up, allowing learned control knowledge of parent robots to be transferred to offspring. This is achieved through the use of NEAT and Compositional Pattern Producing Networks (CPPNs) [3], in an on-line evolutionary fashion, to learn suitable parameters for Central Pattern Generators (CPGs) in every active joint of modular robots of varying morphology. By using this indirect encoding, the CPPN can easily be transferred from one robot to a robot with a different morphology. Note that we use on-line evolution as a means of achieving lifetime learning, and do not consider the long-term effects on the evolution of morphologies in this case.

2 EXPERIMENTS AND RESULTS

We employ a framework for modular robots similar to the one described in [2], and the morphologies studied are shown in Fig.1. The control system consists of CPGs in every active joint, with connections linking neighboring CPGs.

DOI: http://dx.doi.org/10.1145/3067695.3076046

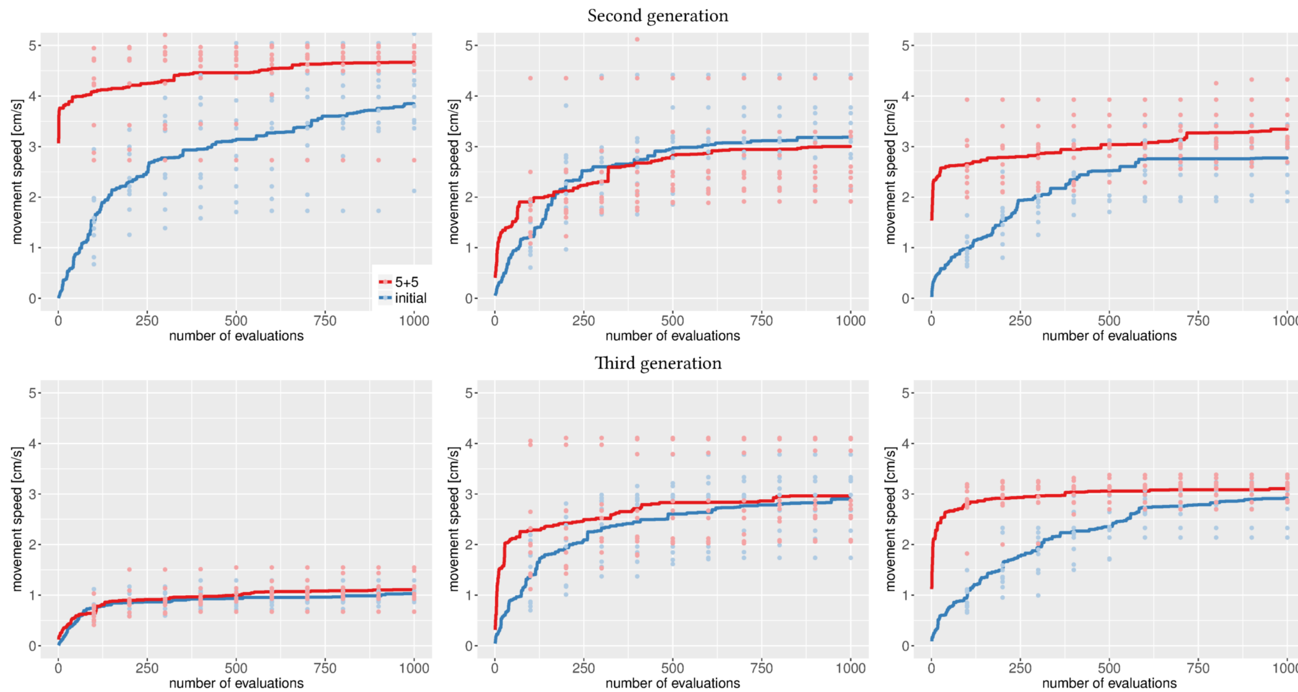

Figure 2: Lifetime learning performance on morphological offspring, using either re-initialized CPPNs (initial), or a Lamarckian set-up where 5 CPPNs are inherited from each parent (5+5). All learning methods are run 10 times, lines show the mean and dots show the individual run results. The layout of the plots corresponds to the two bottom rows of morphologies in Fig. 1.

Starting with a set of three hand-designed morphologies, we apply recombination and mutation on the tree structures describing these robot morphologies to produce a second generation of robots and iterate to produce a third generation. We then run 1000 evaluations of NEAT, with a population size of 10, to learn efficient forward locomotion performance for the hand-designed robots in our simulation environment. We then perform a comparison of the learning performance on the offspring robots, running learning either from 10 randomly initialized CPPNs or from 5 CPPNs from the learning process in each of the morphological parents.

Fig. 2 shows the results of lifetime learning using randomly initialised CPPNs and using the Lamarckian set-up with inherited CPPNs on the offspring morphologies. We observe that at the end of the evaluation budget (1000 evaluations), the approaches have reached mostly similar locomotion performances, although the Lamarckian approach seems to have a slight advantage. This may indicate that, given enough time, both approaches could reach the practical limits for a specific morphology. However, the Lamarckian approach clearly has an advantage in the earlier stages of the learning process, reaching higher locomotive performance in the first 100 evaluations. This characteristic of quick convergence to a high locomotive performance is particularly interesting for a scenario where the learning is performed on real robots. In such cases, evaluations are both time-consuming and may cause considerable wear and tear on the robots, and thus an early termination of the learning phase would be a considerable advantage. We also observe from the bottom left plot in Fig.2 that both approaches struggle to achieve good performance for this particular morphology, however, we speculate that this could be related to the morphology being small and less suited for a fast locomotion strategy.

3 CONCLUSION

This paper established that CPPNs combined with CPGs can enable a Lamarckian scheme for morphologically evolving modular robots. We have demonstrated that our approach for transferring control knowledge from parent robots can be beneficial for faster convergence of good locomotion performance in a lifetime learning process.

Future work will focus on analyzing this effect, e.g. investigating if insertion of fresh genetic material in the learning process would be beneficial, and how the genetic material is related to the performance increase. Moreover, it would be beneficial to study evolutionary timescale effects (e.g., of different selection schemes) on lifetime learning, and to investigate both Lamarckian and Baldwinian approaches for different scenarios.

REFERENCES

[1] Nick Cheney, Josh Bongard, Vytas Sunspiral, and Hod Lipson. 2016. On the Difficulty of Co-Optimizing Morphology and Control in Evolved Virtual Creatures. In *Proceedings of the Artificial Life Conference (ALIFE XV)*. MIT Press, 226–234.

[2] Milan Jelisavcic, Matteo De Carlo, Evert Haasdijk, and A. E. Eiben. 2016. Improving RL Power for On-Line Evolution of Gaits in Modular Robots. In *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE.

[3] Kenneth O Stanley, David B D'Ambrosio, and Jason Gauci. 2009. A hypercube-based encoding for evolving large-scale neural networks. *Artificial Life* 15, 2 (2009), 185–212.