Engineering Online Evolution of Robot Behaviour

(Doctoral Consortium)

Fernando Silva
Supervisors: Anders Lyhne Christensen and Luís Correia
BioMachines Lab & Instituto de Telecomunicações & BioISI
Lisboa, Portugal
fsilva@di.fc.ul.pt

ABSTRACT

Evolutionary computation techniques have been widely studied to automate the synthesis of behavioural control for robots. In online evolution, an evolutionary algorithm is executed on the robots themselves during task execution so as to continuously optimise the robot controllers. Online evolution provides numerous potential benefits, including enabling robots to modify their behaviour in response to changes in the task and in the environment. Current approaches to online evolution on physical robots, however, often require a prohibitively long evolution time and a substantial amount of human experimentation, and have not yet scaled to complex real-world tasks. In this research, we study how to accelerate and scale online evolution to more complex tasks while minimising the amount of human intervention. Our ultimate objective is to enable the realisation of real-world multirobot systems that can effectively learn new behaviours and adapt online to take on dynamic tasks in a timely manner.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning—*Connectionism and neural nets, concept learning*

General Terms

Algorithms

Keywords

Adaptation and learning; artificial neural network; evolutionary robotics; online evolution; robot controller

1. INTRODUCTION

A long-standing goal in artificial intelligence is synthesising agents that can effectively learn and adapt *online* throughout their lifetime [4]. When considering embodied agents such as robots, the benefits of online learning and adaptation are numerous. Online learning refers to a robot's capacity to modify its behaviour in response to unforeseen circumstances, such as changes in the task and in the environment. Online adaptation is related with behavioural robustness and refers to a robot's ability to maintain qualitatively similar behaviours despite environmental or task-related changes.

One of the most radical and open-ended approaches to adaptation and learning in robotic systems is *online evolution* [10], which falls under the umbrella of evolutionary robotics. In online evolution approaches, an evolutionary algorithm (EA) is used as a means to continuously optimise robot controllers. The EA is executed on the robots themselves during task execution. The main components of the EA (evaluation, selection, and reproduction) are performed by the robots without any external supervision. As a result, robots may be capable of continuous self-adaptation in a completely autonomous manner.

The first study on online evolution in a real robot was carried out by Floreano and Mondada [1]. Afterwards, Watson *et al.* [11] developed *embodied evolution*. In this approach, the use of multirobot systems was motivated by the potential speed-up of evolution due to groups of robots that evolve in parallel and exchange genetic information. Over the past decade, different approaches to online evolution have been introduced, see [10] for examples. However, there are central issues that must be addressed before online evolution becomes a feasible approach to adaptation and learning in real robots. The major impediments to large-scale adoption are: (i) the prohibitively long time that online evolution typically requires (several hours or days) [6], and (ii) the fact that evolutionary robotics has not yet scaled to complex real-world tasks conducted outside controlled laboratory conditions [8]. Furthermore, current online evolution approaches often require a substantial amount of experimentation and human knowledge to determine, for instance, a suitable structure for controllers such as artificial neural networks [10]. If an inappropriate controller structure is chosen, the evolutionary process and thus the potential for adaptation and learning are affected.

2. OBJECTIVES AND METHODOLOGY

The ultimate goal of our research is to enable the realisation of real-world multirobot systems that can effectively learn and adapt online in a timely manner, and that are able to solve tasks with a level of complexity that goes beyond the current state of the art. Specifically, we focus on studying and developing novel methods that can accelerate and scale online evolution to more complex tasks while minimising the amount of intervention from the system designer. Although our research is focused on robot systems, we believe that

Appears in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2015), Bordini, Elkind, Weiss, Yolum (eds.), May 4–8, 2015, Istanbul, Turkey.*

our results generalise to other multiagent systems in which online adaptation and learning are essential.

2.1 Automating the Evolution of Controllers

In online evolution, robot controllers are often based on artificial neural networks (ANNs) because of their ability to tolerate noise in sensors. Previous approaches to online evolution have typically been limited to the evolution of weights in ANNs with a fixed-topology [10]. Such methods require the system designer to decide on a suitable topology for a given task, which usually involves intensive experimentation. We have recently introduced an algorithm called odNEAT [10] for decentralised online evolution of ANN-based controllers in multirobot systems. In odNEAT, both weights and topology of ANNs are under evolutionary control. The algorithm starts with minimal networks with no hidden neurons. ANN topologies are gradually complexified by adding new neurons and new connections through mutation, thus allowing odNEAT to bypass the inherent limitations of fixed-topology algorithms and automatically find an appropriate degree of complexity to solve the current task.

In previous simulation-based studies, we have shown that: (i) odNEAT effectively evolves controllers for robots that operate in dynamic environments with changing task parameters [9], (ii) the controllers evolved are robust and can often adapt to changes in environmental conditions without further evolution [10, 5], and (iii) that robots executing odNEAT can display a high degree of fault tolerance: they are able to adapt and learn new behaviours in the presence of faults in the sensors [10].

In order to effectively evolve complex and large-scale controllers, we have also introduced a novel hybrid-encoding approach called R-HybrID [7]. Traditional evolutionary approaches use *direct encodings*, which optimise each ANN parameter independently and are therefore limited in their ability to evolve large-scale ANNs. *Indirect encodings*, on the other hand, facilitate scalability because each gene can be reused multiple times to construct the ANN. However, indirect encodings are biased towards regularity and often become ineffective when irregularity is required (e.g. fine-tuning a single ANN weight). In R-HybrID, controllers have both indirectly encoded and directly encoded structure. Because the portion of structure following a certain encoding is under evolutionary control, R-HybrID can *automatically* find a suitable encoding combination for a given task.

3. ONGOING AND FUTURE WORK

We are currently studying the application of odNEAT to different real-world scenarios. We are analysing how to effectively: (i) reduce the currently prohibitive long time required to evolve controllers directly on real robots, and (ii) scale online evolution to more complex tasks. To that end, we have recently extended odNEAT to evolve general behavioural building blocks that are prespecified in the ANN as *macro-neurons* [6]. Macro-neurons can be either evolved or preprogrammed. During online evolution, the structure and the parameters of macro-neurons are optimised together with the ANN's weights and topology in a unified manner. In simulation-based experiments [6], we have shown that macro-neurons enable a substantial speed up of the evolution time, and the synthesis of high performing solutions.

The results we have obtained so far are promising, and suggest that our proposed algorithms are a potentially viable approach to efficient online evolution. The aim of our ongoing work is to evaluate the performance of odNEAT in real multirobot systems. In this respect, we also intend to investigate the effects of behavioural diversity-based methods [3, 2] in online evolution. In tasks where the environmental conditions or task requirements constantly change, explicitly encouraging behavioural novelty may enable evolution to more easily produce a large variety of effective solutions, thereby opening a new path towards more evolvable robot systems.

Acknowledgments

This work was partially supported by FCT under grants SFRH/BD/89573/2012, UID/EEA/50008/2013, UID/Multi/04046/2013, and EXPL/EEI-AUT/0329/2013.

REFERENCES

[1] D. Floreano and F. Mondada. Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In *3rd International Conference on Simulation of Adaptive Behavior*, pages 421–430. MIT Press, 1994.

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

[3] J. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary Computation*, 20(1):91–133, 2012.

[4] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.

[5] F. Silva, L. Correia, and A. L. Christensen. Dynamics of neuronal models in online neuroevolution of robotic controllers. In *16th Portuguese Conference on Artificial Intelligence*, pages 90–101. Springer, 2013.

[6] F. Silva, L. Correia, and A. L. Christensen. Speeding up online evolution of robotic controllers with macro-neurons. In *16th European Conference on the Applications of Evolutionary Computation*, pages 765–776. Springer, 2014.

[7] F. Silva, L. Correia, and A. L. Christensen. R-HybrID: Evolution of agent controllers with a hybridisation of indirect and direct encodings. In *International Conference on Autonomous Agents and Multiagent Systems*. IFAAMAS, 2015. In press.

[8] F. Silva, M. Duarte, S. M. Oliveira, L. Correia, and A. L. Christensen. The case for engineering the evolution of robot controllers. In *14th International Conference on the Synthesis and Simulation of Living Systems*, pages 703–710. MIT Press, 2014.

[9] F. Silva, P. Urbano, and A. L. Christensen. Online evolution of adaptive robot behaviour. *International Journal of Natural Computing Research*, 4(2):59–77, 2014.

[10] F. Silva, P. Urbano, L. Correia, and A. L. Christensen. odNEAT: An algorithm for decentralised online evolution of robotic controllers. *Evolutionary Computation*, 2015. In press.

[11] R. Watson, S. Ficici, and J. Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18, 2002.