Validation of a Learning and Evolving Robot Swarm

Rasmus Munk
School of Computing
Edinburgh Napier University
Edinburgh, Scotland, UK
munk1@live.dk

Emma Hart
School of Computing
Edinburgh Napier University
Edinburgh, Scotland, UK
e.hart@napier.ac.uk

Ben Paechter
School of Computing
Edinburgh Napier University
Edinburgh, Scotland, UK
b.paechter@napier.ac.uk

ABSTRACT

Recent research in small populations of Thymio II robots illustrated the relative benefits of populations distinguishing heritable and learning features in robots for a simple obstacle avoidance task. Here we scientifically validate these results by repeating them using a simulation. An additional benefit of this work is to provide confidence in the simulation model. This is important because evolutionary swarm robotics experiments can be very time consuming to run in real robots. Having a reliable simulation allows many more experiments to be run in simulation with only the most interesting results needing to be verified with real robots. We describe the development of a simulation using RoboRobo that's using the same three-tier learning framework that was demonstrated in the real-world. The simulation is shown to replicate the real-world results in terms of illustrating the relative benefits of each type of learning, and if anything, indicates that social learning can be more powerful than originally thought.

CCS CONCEPTS

•**Computing methodologies** → **Artificial intelligence;** *Simulation evaluation; Control methods;*

KEYWORDS

Evolutionary Robotics, Simulated Robotics, Social Learning

ACM Reference format:
Rasmus Munk, Emma Hart, and Ben Paechter. 2017. Validation of a Learning and Evolving Robot Swarm. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017,* 2 pages.
DOI: http://dx.doi.org/10.1145/3067695.3082550

1 INTRODUCTION

Heinerman et al. proposed a three-tiered adaptation engine for a swarm of physical robots (Thymio II), that distinguished between *heritable* and *learnable* features in the robots. An evolutionary mechanism evolved a genome that controlled which of the proximity sensors of the robot were switched on in a given generation. This information is *heritable* from robot to robot. An *individual* learning mechanism was used to learn the weights of a feed forward neural network controller, using active sensors as input and outputs the speed of each of the two motors. The weights are encoded by a *memome* and an Evolution Strategy is used to adapt weights during a generation. Finally, *social learning* is implemented by a robot broadcasting its memome to other robots, who store collected memomes and use this information in conjunction with their current memome to produce an adapted controller. The framework was evaluation using a population of 6 Thymio II robots and demonstrated in particular the benefits of social learning.

Our work repeats these experiments using a simulation rather than real robots. The purpose of this is to provide scientific validation of the previous work by approaching the same research question using a different experimental technique and a more generic robot model.

In addition to this we were motivated by a desire to have further confidence in the simulator to accurately model behaviours which encompassed complex interactions of learning types. When investigating questions in environments that model complex interactions with both other group members and the environment, simulations can provide a straightforward means of testing many hypotheses [2] in a manner that uses far less elapsed time. This allows for extensive experimentation to be performed, with only the most interesting results tested in real robots.

We describe the development of a simple simulation that loosely models the Thymio II robots and repeats the experiments described in [4]. Although commercial simulations over Thymios are available (e.g in Webots[1]) we felt it was important to use a simple model of a generic robot in order to demonstrate that the results are not specific to Thymio hardware or precise morphology. The goal was to demonstrate that the three-tiered architecture provided results with the same features as those shown by Heinerman *et al*, i.e. that social learning increases learning speeds and results in better controllers, before moving on to a deeper investigation of factors influencing social learning, informed from biological systems. The intention is *not* to demonstrate that controllers learnt in simulation can successfully cross the *reality-gap* [6].

2 METHOD

The setup described in [3], [5] was implemented into Roborobo [1]. This included a virtual replica of the environment shown in [4] and a robot sensor configuration that approximates physical layout of the Thymio II robot, i.e 5 sensors are used, 3 in the front and 2 at the back. Unlike the Thymio II, a circular robot was used, which had diameter 55 pixels, in order to scale them to the same relative size as the Thymio's in the physical environment (1cm=5pixels). Each sensor had a cone-shaped layout - where each cone had a range of 77 pixels, with a cone diameter ranging between 8 and 21 pixels.

[1] https://www.cyberbotics.com/overview

GECCO '17 Companion, Berlin, Germany
© 2017 ACM. 978-1-4503-4939-0/17/07…$15.00
DOI: http://dx.doi.org/10.1145/3067695.3082550

The same obstacle avoidance task as described in [3] is used; the same fitness function was adopted as shown in 1.

$$f = \sum_{t=0}^{T} s_{trans} * (1 - s_{rot}) * (1 - v_{sens}) \quad (1)$$

Here the s_{trans} is the combined speed of the two motors. s_{rot} is the rotational speed of the robot normalized between 0 and 1 and v_{sens} is the normalized distance between 0 and 1 of the sensor that is closest to an object.

To simulate the simultaneous nature of multiple robots being evolved in parallel, a multi-threaded evaluator was implemented in which each robot had a dedicated evaluator for determining fitness. The evolutionary, individual learning and social learning algorithms were implemented exactly as described in [4] — the reader is referred to this publication for the details.

3 EXPERIMENTS & RESULTS

To evaluate the contribution of social learning to the quality and speed of learning we perform two experiments, as described by [4]:

- Fixed sensor layout and individual learning of controller
- As above, with added social learning —- broadcasting of memomes by robots.

Experiments were repeated 20 times. Figure 2 shows the mean fitness for each setup per generation, and table 1 gives the mean and standard deviation at the final generation, the distribution of which can be seen in Figure 1. At t-test shows the difference between means is statistically significant (p-value = 1.522e-05) at the 1% level. Figure 2 shows similar trends to that found in using physical robots; the gap between social learning and individual learning only appears enhanced in the simulated system. This is most likely explained by the lack of noise in the system, with perfect reception of broadcast memomes, and a more accurate translation of motor speeds to movement.

4 CONCLUSION

We have described a simulated version of the three-tier framework proposed in [4] which replicates the general trends of the results to a degree that is sufficient to scientifically validate those results. This also enables us to confidently use the simulation to investigate in depth further aspects of socially-related learning. For example, instead of indiscriminate broadcasting of memomes we might share only between related-kin. As with any simulation, the most interesting results will still have to be verfied using real robots and further adaptation made to the simulator as appropriate.

We have already identified that the simulation could be further improved by use of an alternative fitness function: the blind-spots introduced by the sensor models in the simulator enable a high-fitness to be obtained when the motors are running at full-speed (s_{trans} and the robot is in fact colliding with an obstacle, due to the sensor failing to detect an obstacle, i.e. v_{sens} is incorrectly reading 0.

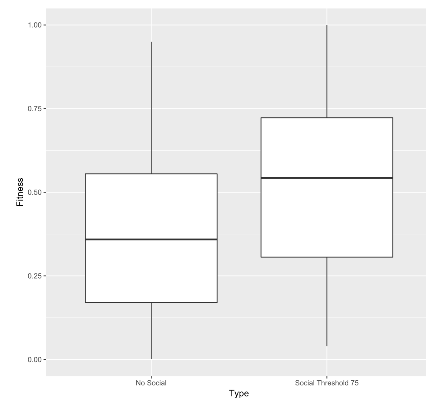

Figure 1: Final Generation Population Fitness

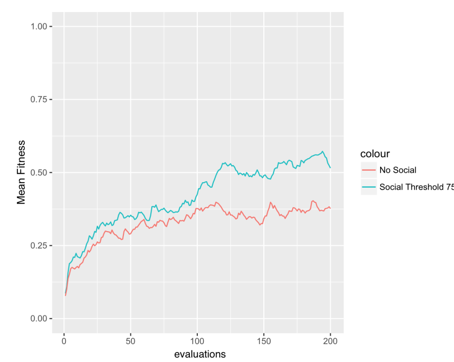

Figure 2: Population Mean Fitness

Table 1: Mean fitness at final generation

Experiment	Mean Fitness	SD
Individual+Social Learning	0.5152	0.2474
Individual Learning	0.3764	0.2388

REFERENCES

[1] N Bredeche, J. M Montainer, B Wheel, and E Haasdijk. 2013. Roborobo! a fast robot simulator for swarm and collective robotics. (2013), CoRR, abs/1304.2888.
[2] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E. (Gusz) Eiben. 2015. Evolutionary Robotics: What, Why, and Where to. *Frontiers in Robotics and AI* 2 (2015), 4. https://doi.org/10.3389/frobt.2015.00004
[3] Jacqueline Heinerman and Massimiliano Rango. 2015. Thymio Swarm Parameters. (2015). https://github.com/ci-group/Thymio
[4] Jacqueline Heinerman, Massimiliano Rango, and A E Eiben. 2015. Evolution , Individual Learning , and Social Learning in a Swarm of Real Robots. (2015). https://doi.org/10.1109/SSCI.2015.152
[5] Jacqueline Heinerman, Massimiliano Rango, and A E Eiben. 2015. Thymio Swarm Github. (2015). https://github.com/ci-group/Thymio
[6] Sylvain Koos, Jean-Baptiste Mouret, and Stéphane Doncieux. 2010. Crossing the reality gap in evolutionary robotics by promoting transferable controllers. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*. ACM, 119–126.