Collective Sharing of Knowledge in a DREAM

Jacqueline Heinerman
VU University Amsterdam
Amsterdam, The Netherlands
j.v.heinerman@vu.nl

ABSTRACT

Generalising on-line learned knowledge in evolutionary robotics results in robots that can accomplish tasks in varying circumstances. This is the goal of the DREAM project. Even faster accomplishment of tasks and understanding of the environment can be realised when there is a collective of robots that share information–*social learning*. In this paper, we propose research questions and scenarios for investigating social learning within the DREAM project.

1. INTRODUCTION

The DREAM project aims to incorporate sleep and dream-like processes within a cognitive architecture to achieve the capability to generalise knowledge gained in on-line learning, in particular in learning through evolutionary methods. Through generalisation of knowledge, robots are able to identify chains of behaviours that solve (sub-)tasks in varying circumstances. This will allow the robots to go beyond the rote learning of appropriate behaviour for very particular tasks and environments.

Within a collective of robots, sharing knowledge between robots can lead to faster learning and the identification of more efficient behaviours through a consolidation of the knowledge and experiences acquired by the different individuals. The topic of this paper is this learning by sharing knowledge–*social learning*, particularly our ideas for researching the impact and possibilities of social learning in the context of DREAM.

Implementing social learning implies a number of design choices that define with whom, when and what knowledge is shared. DREAM will investigate the impact of these aspects (*who*, *when* and *what*) on the speed of learning and the efficiency of the behaviours that arise when implementing a social learning algorithm in a cognitive architecture that allows for generalised behavioural patterns at a more abstract level than direct sensory-motor loops.

It has been noted that social learning implements an evolutionary system [1, 2]. In fact, many of the considerations in this paper apply as much to on-line evolutionary systems such as Embodied Evolution [4]. We use the term social learning rather than embodied evolution to differentiate from evolutionary processes at other levels that are also part of the cognitive architecture in DREAM (which we cannot describe here for lack of space).

Many of the choices we consider are similar to choices in evolutionary algorithms, particularly regarding various selection schemes. Therefore, some options for our design choices are inspired by research into evolutionary robotics, in particular on-line evolutionary systems such as EE.

This paper outlines the research questions and experimental scenarios we will conduct to investigate social learning in the initial phases of the DREAM project.

2. SOCIAL LEARNING ALGORITHMS

We consider three important components when investigating a social learning algorithm: *who*, *when* and *what*. There are more components describing a social learning algorithm but these are considered out of scope for the DREAM project. In this section we explain possible choices of each component and describe the research questions that we will investigate.

Who

To decide who can send and receive information we consider two possibilities. Either everyone has the possibility to send and receive information (egalitarian) or robots can have a predefined role to be a teacher or student (teacher-student).

Using a teacher-student mechanism has been studied in [1]. In this paper social learning proves to spread knowledge pieces over the population, acting as an accelerator for individual learning and as a knowledge repository of individually discovered knowledge that otherwise would be lost after an agent's death.

In [3] the researchers investigate an egalitarian model of social learning (ESL) in which agents are not labelled as teachers or students, instead allowing any individual receiving a sufficiently high reward to teach other agents. ESL promotes diverse behaviour in the overall population and prevents premature convergence.

For our research, we investigate both the egalitarian and the teacher-student model for different purposes. The teacher-student model helps to first investigate how good knowledge can spread through the group. Then we use the egalitarian model to see how the group can build the required knowledge together, making the teachers redundant and therefore the learning open-ended.

GECCO '15, July 11 - 15, 2015, Madrid, Spain

DOI: http://dx.doi.org/10.1145/2739482.2768491

When

The *when* means to decide on the moment when communication takes place and the selection mechanism. There are many options for choosing a communication moment. Communication can for example arise when two robots are within communication range or at the end of each generation. For the rules that tell whether communication actually takes place, a selection mechanism has to be chosen where the quality of the robot controller influences the participation in social learning of the robot.

An example where the range and the quality of the controller influence the participation in social learning is [4]. A robot broadcasts genes to other robots within a certain range at a rate proportional to its energy level. Conversely, a received gene is accepted at an inversely proportional rate to the agent's energy level. They conclude that the evolved behaviors are better than the best hand-coded behaviors.

We will investigate a similar setup where the frequency of sending information is dependent on the performance of the robot. We will also investigate a mechanism where it is possible for the robot to ask for specific knowledge from others.

What

The *what* decides the explicit information that is transferred between two robots. Therefore, transferring knowledge by imitation is out of scope for our research.

In on-line evolutionary robotics, a commonly used method for implementing the *what* component is by sending (part of) the robot's controller.

As a result of the knowledge generalisation in the DREAM project, we will have different levels of knowledge to share among the robots. We will investigate the influence of the type of shared information on the behavior of the robots. We also investigate whether we can find knowledge pieces that are general enough to send between physically different robots.

3. SCENARIO

The environment we use to investigate social learning in the initial phases of the DREAM project is *foraging*: searching for objects to bring to a target.

The robot we will use is a mobile Thymio II robot extended with a Raspberry Pi with a camera and WiFi. In total, we will use 4 environments with increasing complexity. The environments and robot are shown in Figure 1.

Environment 1 has a switch that has to be pushed to open the door and environment 2 has one puck to bring to the target. Environment 3 is a combination of the previous two where there is a switch and a puck to collect. Environment 4 is a variation of environment 3 where the button, puck and target have a different position.

The reason for using multiple environments is to show that the on-line learned knowledge can be generalised over different environments which makes social learning beneficial even when robots are in different situations.

There are different objects the robot needs to recognise (button, puck and target) and act accordingly. The robot will be able to recognise these object with the camera. Recognising an object will result in choosing a certain behavior. Behaviors are implemented as neural networks.

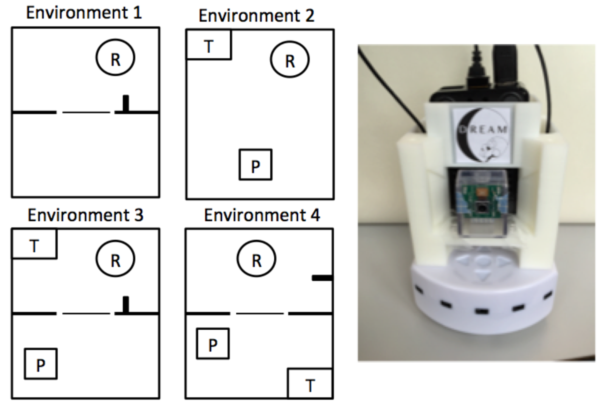

Figure 1: Left: the four environments where R=robot, P=puck and T=target. Right: the Thymio II robot extended with Raspberry Pi and camera module (3D support designed by ISIR/UPMC)

The weights of these neural networks and the corresponding behaviors need to be learned and are (partially) shared by the robots.

4. ACKNOWLEDGMENTS

This work was made possible by the European Union FET Proactive Initiative: Knowing, Doing, Being: Cognition Beyond Problem Solving funding the Deferred Restructuring of Experience in Autonomous Machines (DREAM) project under grant agreement 640891.

The author would like to thank Nicolas Bredèche, Evert Haasdijk, Gusz Eiben and other partners in the DREAM project for many inspirational discussions on the topics presented here.

5. REFERENCES

[1] E. Haasdijk, P. Vogt, and A. E. Eiben. Social learning in population-based adaptive systems. In Z. Michalewicz and B. Reynolds, editors, *Proceedings of the 2008 IEEE Congress on Evolutionary Computation*, Hong Kong, China, 2008. IEEE Computational Intelligence Society, IEEE Press.

[2] C. Marriott and J. Chebib. The effect of social learning on individual learning and evolution. *arXiv preprint arXiv:1406.2720*, 2014.

[3] W. Tansey, E. Feasley, and R. Miikkulainen. Accelerating evolution via egalitarian social learning. In *Proceedings of the 14th annual conference on Genetic and evolutionary computation*, pages 919–926. ACM, 2012.

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