In Berthouze, L., Kozima, H., Prince, C. G., Sandini, G., Stojanov, G., Metta, G., and Balkenius, C. (Eds.)
Proceedings of the Fourth International Workshop on Epigenetic Robotics
Lund University Cognitive Studies, 117, ISBN 91-974741-3-4

Artificial Immune Networks for Robot Control

Juergen Rattenberger, Patrick M. Poelz, Erich Prem
Austrian Research Institute for Artificial Intelligence
Freing 6/6, 1010 Vienna, Austria
{juergen, patrick, erich}@oefai.at

Emma Hart, Andrew Webb, Peter Ross
Napier University, School of Computing
Edinburgh EH10 5DT, UK
{e.hart, a.webb, p.ross}@napier.ac.uk

Abstract

We investigate how a robot can be provided with an architecture that would enable it to developmentally 'grow-up' and accomplish complex tasks by building on basic built-in capabilities. The paper introduces into the basic principles of AIS and presents experimental results from a real robot. To our knowledge, this is the first implementation of an AIS architecture for controlling a real mobile robot.

1. Introduction

Much of current research in learning robots is devoted to highly specialized tasks and equally specialized solutions for learning the desired tasks. In an epigenetic approach, however, the focus should be on much more versatile robots that allow the incremental construction of robot capabilities with increasing complexity. The work presented in this paper aims at developing a process through which complex perceptual structures and control components emerge as a result of the robot's interaction with its environment, see (Ross et al., 2003). The technology chosen here is inspired by biological immune systems.

2. The AIS approach

The artificial immune system approach used here was particularly inspired by Jerne's immune network theory (Jerne, 1973).

In the beginning the robot is driven by basic instincts such as a 'desire to avoid collisions' and a 'desire to seek novelty'. The robot then should learn through experience, and the learned behaviours should gradually take over control from the instinct-driven initial system. The robot therefore needs to capture some minimal details of its experiences. This is realised by so called rule like associations (RLAs). A detailed specification of RLAs can be found in (Hart et al., 2003). RLAs consist of three parts: condition, action and expectation.

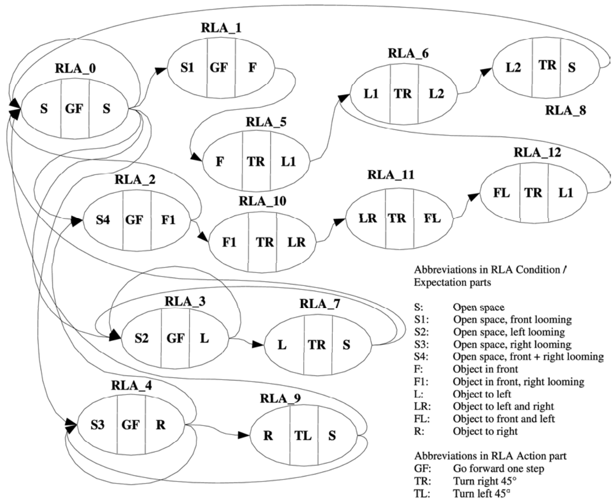

Figure 1: Flow diagram showing sequences of chosen RLAs. Taken from (Hart et al., 2003).

In analogy to the immune system metaphor the RLA corresponds to the antibody and the sensory data corresponds to the antigen. The condition and action parts of a specific RLA can be regarded as paratope and epitope. According to Jerne's immune network hypothesis the antibodies recognise each other by the paratope of one antibody and the epitope of another and in this way stimulate or suppress each other.

Figure 1 depicts a typical sequence of RLAs to be generated in a developmental process of adaptation. Paths through the network can be interpreted as an episodic long-term memory. This memory maintains a record of relationship between sensory conditions, actions taken, and the consequences of those actions.

The basic learning algorithm works as follows: the system is presented with environmental features. The algorithm then selects an RLA whose condition part is closest to the environmental input situation (antigen). The system then takes appropriate action based on the RLA (antibody). The algorithm evaluates whether the RLA correctly predicted the expected outcome from taking this action. If the system behavior is in line with the desired outcome, the RLAs which produced this system behavior receive positive reinforcement. This increases an RLA's weight corresponding to concentration in the immune system analogy. In the case of negative reinforcement, the RLA chosen is cloned and mutated and the concentration is decreased. Also, RLAs that are connected to the RLA which generated negative reinforcement receive negative feedback.

3. Experiments

In the experiments reported here, the robot is driven by the desire to 'avoid obstacles' or better 'try to avoid obstacles' and the 'desire to seek novelty'. By reinforcement the AIS then learns to perform wall following. In order to gain more knowledge about the AIS approach, we take a closer look on the network at different points in time. Thus, we examine the RLA network after 4,000 iterations, then after 6,000 and finally after 12,000 and judge how good the robot learned wall following.

The test bed used for the experiments with the real robot consists of a circular barrier made from wooden boards aligned in a way that the ongoing process of learning wall following can be easily monitored. The robot used in our experiments is a 6-wheeled KURT2 robot originally designed for sewage pipe inspection. In our experiments we use its 12 IR and 2 ultrasonic sensors only.

We performed several tests with KURT2. The aim of the first series of tests was to investigate the structure of RLA networks after several iterations and how long it would take for KURT2 to successfully perform wall following. To find out, what the required population size for a stable network would be, several experiments with different maxima of the population size have been performed and analyzed. A typical result is depicted in Figure 2.

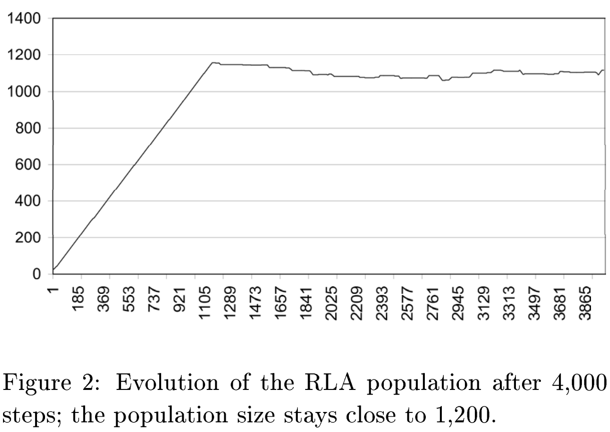

Figure 2: Evolution of the RLA population after 4,000 steps; the population size stays close to 1,200.

In order to find out a good maximum size for the RLA population or the smallest size for a functioning RLA network, we performed experiments with different maxima, ranging from 20, 200, 450 to more than 1,000. The first experiment performed with 20 RLAs and 6,000 iterations took only about 7 minutes. The resulting behaviour was completely instable and the robot once even turned in circles for a while. It also bumped against the wooden barrier several times, even after this 6,000 iterations. By increasing the number of iterations to 12,000, the robot performed wall following on the left side from time to time but, again, no stable behaviour could be produced. Stability was only achievable for populations of 1,000 RLAs and more. Although, that does not mean, that all of the 1,000 RLAs are being used for simple wall following, only less than 20 have been used on a regular basis, the others were only needed for the learning-process.

These first results are encouraging in that this is the first robot implementation of AIS for behaviour generation which we are aware of. Future work will focus on more complex behaviours generated using this AIS principles. Following (Piaget, 1952) it is also planned to introduce additional levels of abstraction so as to equip the robot with more stringent conceptual knowledge.

Acknowledgements

This research is supported by the EC as IST project SIGNAL. Partners are University of Bonn, Napier University, National Research Council Genova, and the Austrian Research Institute for Artificial Intelligence - also supported by the Austrian Federal Ministry for Education, Science, and Culture and the Austrian Ministry for Transport, Innovation, and Technology.

References

Hart, E., Ross, P., Webb, A., and Lawson, A. (2003). A role for immunology in next generation robot controllers. In *Proc. of ICARIS 2003, Edinburgh, UK*, volume 2787, pages 46–56. Springer-Verlag Heidelberg.

Jerne, N. (1973). The immune system. *Scientific American*, 229(1):52:60.

Piaget, J. (1952). *The Origins of Intelligence in Children.* Norton, N.Y.

Ross, P., Webb, A., Hart, E., Prem, E., Poelz, P. M., and Morgavi, G. (2003). Requirements for getting a robot to grow-up. In *Lecture Notes in Computer Science*, volume 2801/2003, pages 847–856. Springer-Verlag Heidelberg.