The Transfer Problem from Simulation to the Real World in Artificial Evolution

Peter Eggenberger Hotz[1,2] and Gabriel Gómez[1]
[1]AILab, University of Zurich, Switzerland
[2]ATR, Kyoto, Japan
eggen@ifi.unizh.ch

Abstract

In evolutionary robotics artificial neural networks are often used as controllers. As the process to evolve such controllers in the real world is time-consuming, one usually uses simulators to speed up this process. By doing so a new problem arises: The evolved controllers in simulation show often not the same fitness as those in the real world. In order to close this reality gap we propose to evolve networks able to change with experience allowing to adapt to unforseen perturbations. This paper reports on experiments with a robot arm for which a controller was evolved able to adapt to changes of the robot's morphology. The controller was not specified but grown using a developmental approach in which the cells were controlled by artificial genomes.

1 Introduction

In Evolutionary Robotics(ER) artificial neural networks are widely used to construct controllers for autonomous mobile agents [6]. As the evolution in the real world is time-consuming, simulations are used to evolve the controller in a simulated environment and the best individuals are tested in the real world. The flaw of this combined approach is that evolved agents in simulated environments show often a significantly different behavior in the real world due to unforeseen perturbations. A gap between the simulated and real environments exists. Therefore, evolved controllers should adapt not only to specific environments, but should be robust against environmental perturbations.

Other authors tried to overcome the gap between simulations and real-world by different methods [8, 9, 11]. Here, a more biological approach is proposed to evolve large neural networks, which are able to grow and learn; in this method the genetic information no longer directly encodes for specific properties (in this case a neural network), but controls developmental processes generating the functioning structures. This will be explained in the following section. In section 3 the robustness of the learning mechanism is tested, the neural structure and its connections to the robot are described, then the experiments and results are shown in section 4 and we conclude with a discussion in the last section.

2 Evolutionary methods and developmental processes

In order to create the neural controller a developmental evolutionary approach was chosen. The used simulator was developed over the last years starting with simple growth control, neural growth [5], pattern generation and morphogenetic processes [4]. Each cell in the simulator has its own genome, in which each gene is regulated by its own set of regulatory units. A gene will encode either a simple chemical (most important signalling molecules and receptors) or a predefined developmental process such as cell division or axonal outgrowth. Although all the cells have the same genome they become different by exchanging signals. For instance, a cell contains since the beginning a signalling molecule able to activate the cell division gene. If such a signal is distributed asymmetrically during cell division, only one cell will grow, because its cell division gene is active. The other cell will not contain this signal and therefore its cell division gene is silent and no cell division will occur. By selecting for developmental processes able to perform tasks defining the value for the fitness, it is possible to create genomes able to grow meaningful structures such as different shapes (see Figure 1) or neural controllers for robots (see Figure 2).

To be able to explore both, growth and learning, the

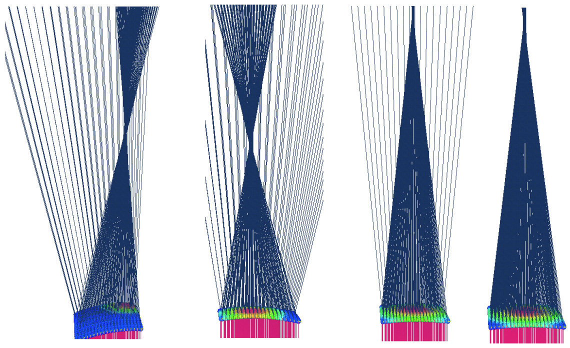

Figure 1: Evolving the shape of an artificial optical lens.

concept of selective recognition between entities, we call it "ligand-receptor interactions" played a crucially important role (Eggenberger[2]). This concept allows (a) a neuron to recognize more or less specifically another partner neuron, (b) an axon to follow cues in the environment and (c) a neuromodulator to diffuse in the environment and influence specific synapses to change their weights. In this way learning rules can be evolved and have not to be pre-specified by the designer of the system.

Growth mechanisms By varying axonal receptor types, different influences can be easily explored by the evolutionary strategy (ES) and growth patterns resulting in topological maps between interneurons and the sensory input emerged. Figure 2 illustrates the mechanisms by which the gene activities of the cells will determine the types of receptors expressed on their axons and finally determine if a source of signaling molecules has an attracting or repelling influence. Similar ligand-receptor interactions will control which cells will stick together or which will build up synapses or which signalling molecule will change the synaptic weight. By varying the expression of different ligand-receptor interactions the ES can easily explore different connectivity patterns with a small load of genetic information.

To illustrate the proposed evolutionary approach a retinal structure was evolved, the task of the system was to evolve a neural network that could move the eye in such a way that an incoming peripheral sensory stimulus falls in the center of the "eye". The fitness function depended on the number of cells, on the movements, and on the precision of the movements measured by the deviation of the position of a stimulus from the center of the eye after peripheral stimulation. More details about how this controller was evolved and

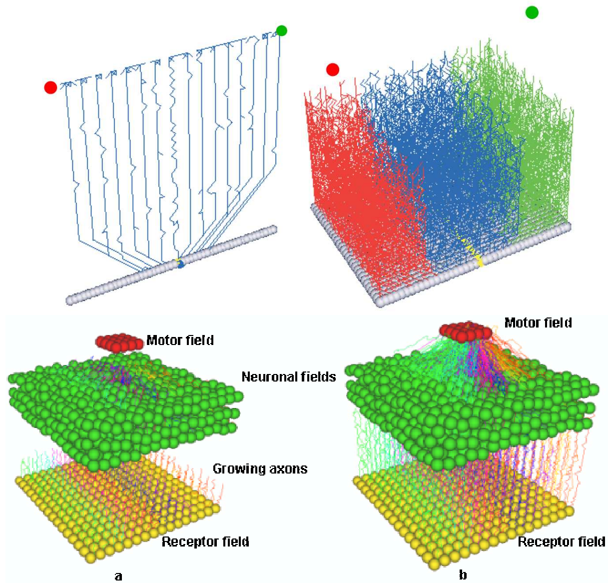

Figure 2: Evolving a neural network for controlling a foveating retina. (a) Growing neural network. (b) Final neural network.

transferred to the robot can be found in the following references: Eggenberger[3] and Eggenberger et al.[5]. Figure 2 gives a schematic overview of the neural network at different stages of evolution. Figure 5a shows the robotic setup and Figure 5c shows the robot tracking a light source.

Learning mechanisms The exploration of learning mechanisms is based on the expression of artificial receptors on the synapses and the interactions with signaling molecules (in Ishiguro et al.[7], this concept was also explored by us to simulate the evolution of a biped robot). In general the change of the weights depends on the activity of different ligand-receptor complexes. By varying the place and time of these interactions, the ES is able to explore learning mechanisms without the need of the designer to declare a set of learning rules (e.g., Hebbian, Anti-Hebbian), because such rules can be evolved. The following example gives a solution of how a Hebbian learning rule could be evolved: Neurons can link the emission of a signaling molecule to their neural activity. Depending on their specific differentiation, the cells produce and emit signaling molecules, the synaptic weight changes if ligands emitted from other cells have a strong enough impact on the synapse. Let us suppose that a neuron emits a substance and the emitted transmitter has an affinity to one of two receptors on a particular synapse and a second neuron emits a substance for a second receptor on that synapse, then the weight will be increased in an AND fashion, a Hebbian learning rule is found. The AES evolved two neuromodulatory systems able to teach the robot about the success of its exploratory movements. The first mechanism localized in the center of the retina cells able to emit neuromodulators when they where activated. This had the effect that the system could learn to move the camera correctly. The second mechanism was more sophisticated and efficient: the neuromodulator was emitted if at any given position the movement of the robot brought the light source closer to the center. This evolved mechanism is quite similar to a value system (Edelman[1] and Sporns et al.[10]). In contrast to their hand-designed methods, where the value system is located in other brain parts, our evolved value system is located directly in the network performing the task.

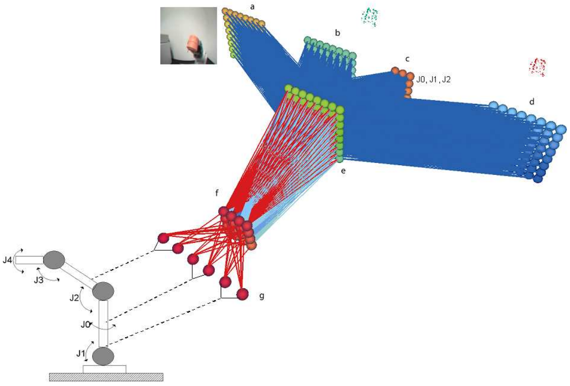

Figure 3: Neural structure and its connections to the robot arm. Neuronal areas: (a) RedColorField. (b) RedMovementToRightField. (c) ProprioceptiveField. (d) RedMovementToLeftField. (e) NeuronalField. (f) MotorField. (g) MotorActivites.

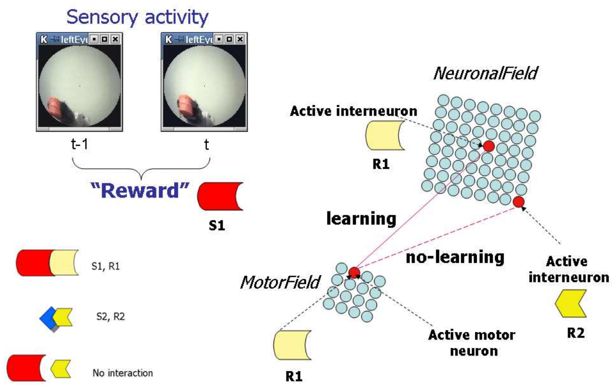

Figure 4: Ligand-receptor concept.

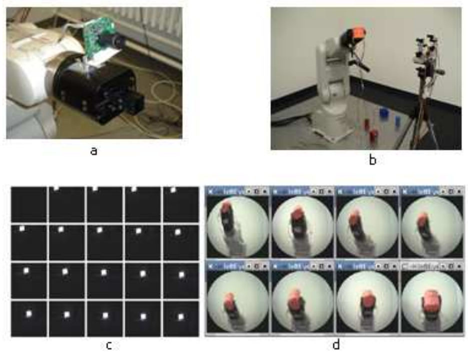

Figure 5: Increasing the sensory-motor capabilities of the robot. (a) First robotic setup. (b) Second robotic setup. (c) First robot task, tracking a light source with the camera. (d) Second robot task, moving a colored object from the periphery to the center of the visual field.

3 Testing the robustness of the learning mechanism

There is an intimate linkage between the neural controller, the sensory-motor system and the task-environment. Therefore to test the robustness of the learning mechanism we increased the sensory-motor capabilities of the robot as well as its task-environment. In our previous work, the task was that a camera which was actuated by two motors should learn to track a light source in a dark environment (see Figures 5a and 5c). We continued testing the evolved controller in the real world increasing the sensory-motor capabilities and the robot's task as follows: (a) The visual system was enhanced to detect color and movement in the environment and a proprioceptive system was added to have feedback of the arm movements. (b) The number of degrees of freedom (DOF) of the robot was increased from two to three. At the beginning of each experiment the active vision system was initialized looking at the center of the visual scene and the position of its motors were held still. The robot arm was placed at a random position at the periphery of the robot's visual field and a colored object was put on its gripper. The robot's task was then to learn how to move the arm in order to bring the object from the periphery of the visual field to the center of it (see Figures 5b and 5d). Figure 4 shows "where" and "when" the synaptic connections between the neuronal areas *NeuronalField* (see Figure 3e) and *MotorField* (see Figure 3f) were changed. The active neurons controlling the robot arm were "rewarded" if the movement of the arm brought the colored object closer to the center of the visual field and "punished" otherwise.

4 Experiments and results

Our experiments were performed using the experimental setup showed in Figure 5b. In order to compare the effects of the learning mechanism different sets of experiments were performed. The result of a non-sensory motor coordinated experiment can be seen in Figure 6a, in this experiment the learning mechanism was not activated and therefore the robot arm made only random movements without being able to solve the task. In contrast when the learning mechanism was activated, the robot was engaged in a sensory motor coordinated loop, as can be seen in the Figure 6b. In this experiment the robot was able to solve the task bringing the object to the target position several times in a stabilized trajectory.

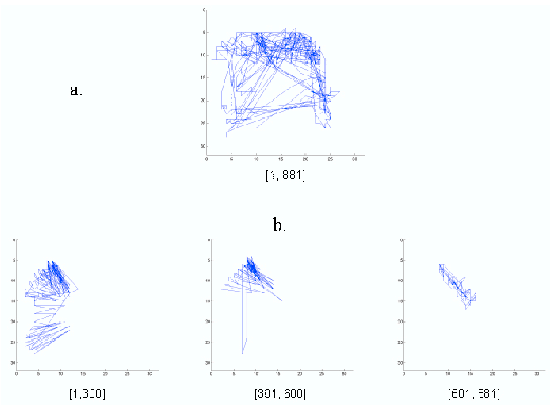

Figure 6: Position of the center of the object in the visual field during the experiments. (a) Random exploration. (b) Sensory-motor coordinated experiment.

5 Discussion

With the above results we could show that the evolved learning mechanism can be transferred to a robot arm, that the task could be learned by the robot and that the eyes can teach the robot arm to bring an object to the center of the visual field not only for two DOF, but also for three DOF. These results point to a new direction of how to overcome the reality gap: the goal should be to evolve learning mechanisms (not fixed neural networks) that can control the performance of a robot interacting with the real world.

Acknowledgments

Peter Eggenberger was sponsored by the EU-Project HYDRA (IST-2001-33060). Gabriel Gómez was supported by EU-Project ADAPT (IST-2001-37173).

References

[1] Gerald M. Edelman. *Neural Darwinism.* New York:Basic Books, 1987.

[2] Peter Eggenberger. Evolving neural network structures using axonal growth mechanisms. In *Int. Joined Conference on Neural Networks*, 2000.

[3] Peter Eggenberger. On evolving a simulated foveating retina. *3D Forum: The Journal of Three Dimensional Pictures*, 15(4):164–169, 12 2001.

[4] Peter Eggenberger. Genome-physics interaction as a new concept to reduce the number of genetic parameters in artificial evolution. In *Proc. of CEC2003)*, pp. 191–198, 2003.

[5] Peter Eggenberger, Gabriel Gomez, and Rolf Pfeifer. Evolving the morphology of a neural network for controlling a foveating retina - and its test on a real robot. In *Proc. of ALIFE VIII*, pp. 243–251, 2002.

[6] Dario Floreano and J. Urzelai. Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13:431–443, 2000.

[7] Akio Ishiguro, Akinobu Fujii, and Peter Eggenberger Hotz. Neuromodulated control of bipedal locomotion using a polymorphic cpg circuit. *Adaptive Behavior*, 11(1):7–17, 2003.

[8] N. Jacobi, Phil Husbands, and Inman Harvey. Noise and the reality gap: The use of simulation in evolutionary robotics. *Proc. of ECAL95*, pp. 704–720, 1995.

[9] H. Lund, J. Hallam, and W. Lee. Evolving robot morphology. In IEEE Neural Networks Council, editor, *Proceedings of IEEE Fourth Int. Conf. on Evolutionary Computation*, pp. 197–202. Piscataway, NJ:IEEE Press, 1997.

[10] Olaf Sporns and William H. Alexander. Neuromodulation and plasticity in an autonomous robot. *Neural Networks*, 15:761–774, 2002.

[11] T. Kondo, A. Ishiguro, Y. Uchikawa, and P. Eggenberger. Autonomous robot control by a neural network with dynamically-rearranging function. In *Proc. of AROB99*, pp. 809–813, 1999.