Simulating development in a real robot

Gabriel Gómez*, Max Lungarella**, Peter Eggenberger Hotz*, Kojiro Matsushita* and Rolf Pfeifer*
*Artificial Intelligence Laboratory
Department of Information Technology, University of Zurich, Switzerland
Andreasstrasse 15, CH-8050 Zurich, Switzerland
gomez@ifi.unizh.ch
**Neuroscience Research Institute
Tsukuba AIST Central 2, Japan
max.lungarella@aist.go.jp

Abstract

We present a quantitative investigation on the effects of a discrete developmental progression on the acquisition of a foveation behavior by a robotic hand-arm-eyes system. Development is simulated by increasing the resolution of the robot's visual system, by freezing and freeing mechanical degrees of freedom, and by adding neuronal units to its neural control architecture. Our experimental results show that a system starting with a low-resolution sensory system, a low precision motor system, and a low complexity neural structure, learns faster that a system which is more complex at the beginning.

1. Introduction

Development is an incremental process, in the sense that behaviors and skills acquired at a later point in time can be bootstrapped from earlier ones, and it is historical, in the sense that each individual acquires its own personal history [15]. It is well known that newborns and young infants have various morphological (bodily), neural, cognitive, and behavioral limitations, e.g., in neonates color perception and visual acuity are poor (implying a poor tracking behavior) [14]; working memory and attention are initially restricted (giving rise to reduced predictive abilities); motor immaturity is even more obvious, movements have a lack of control and coordination (producing inefficient and jerky movements).

The state of immaturity of sensory, motor, and cognitive systems, a salient characteristic of development, at first sight appears to be an inadequacy. But rather than being a problem, early morphological and cognitive limitations effectively decrease the amount of information that infants have to deal with, and may lead, according to a theoretical position pioneered by [16], to an increase of the adaptivity of the organism. A similar point of view was made with respect to neural information processing by [4]. For instance, it has been suggested that by initially limiting the number of the mechanical degrees of freedom that need to be controlled, the complexity of motor learning is reduced. Indeed, an initial freezing (i.e., not using) of degrees of freedom followed by a subsequent freeing (i.e., release) might be the strategy figured out by Nature to solve the degrees of freedom problem first pointed out by [1], that is, despite the highly complex nature of the human body, well-coordinated and precisely controlled movements emerge over time. In other words, it is possible to conceptualize initial sensory, motor, and cognitive limitations as an adaptive mechanism on its own right, which effectively helps speeding up the learning of tasks, and acquisition of new skills by simplifying the external world of the agent.

The aim of this paper is to provide support for the hypothesis that "starting small" makes and agent more adaptive, and robust against environmental perturbations. Other attempts have shared explicitly or implicitly a similar research hypothesis. [11], for instance, applied a developmentally inspired approach to robotics in the context of joint attention. The authors showed that by having the visual capabilities of a robot mature over time, the robot could learn faster. The effect of phases of freezing and freeing of mechanical degrees of freedom for the acquisition of motor skills was examined by [8] and [2]. For a detailed review of the field of developmental robotics see [9]. Although based on the same research hypothesis, the present study makes at least two novel contributions: (a) it considers the concurrent "developmental changes" in three different systems, i.e., sensory, motor, and neural; and (b) it quantitatively compares a "developing" system to a "nondeveloping" system.

Obviously, an understanding of development cannot be limited to investigate control architectures only, but must include considerations on physical growth, change of shape, and body composition, which are salient characteristics of maturation. Given the current state of technology, however, it is not easy to construct physically growing robots. We propose a method to "simulate" development in an embodied artifact at the levels of sensory, motor, and neural system. We use a high-resolution sensory system and a high-precision motor system with a large number of mechanical degrees of freedom, but we start out by simulating, in software, lower resolution sensors (i.e. by averaging over neighboring pixels in the camera image, or by using only a few pressure sensors) and an increased "controllability" (i.e., by freezing most degrees of freedom). Over time, we gradually increase the resolution of the sensors and the precision of the motors by successively freeing these "degrees of freedom" (i.e. by starting to use the "frozen" joints) and added neuronal units to the neural control architecture.

In the following, we present quantitative results demonstrating how a concurrent increase of sensory resolution, motor precision and neural capabilities can shape an agent's ability to learn a task in the real world, and speed up the learning process.

In the following section we introduce our experimental setup, we then proceed to specify the robot's task in section 3. The neural network and how it is embedded in the robot are described in section 4. The developmental approach is described in sections 5 and 6. The experiments performed are described in section 7, and the results are discussed in section 8. Finally, we point to some future research prospects in the last section.


(a)   (b)
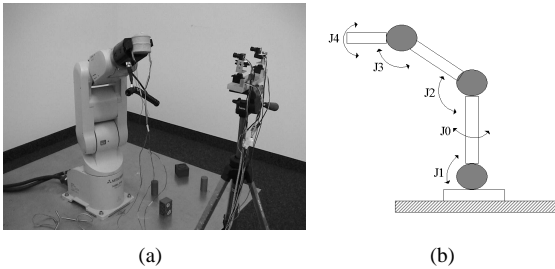
Figure 1. Experimental setup consisting of six degrees of freedom robot arm, four degrees of freedom color stereo active vision system, and a set of tactile sensors placed in the robots gripper.

2 Experimental setup

We performed our experiments by using the experimental setup shown in Figure 1. It consisted of the following components:

- *Robot arm.* An industrial robot manipulator (Mitsubishi MELFA RV-2AJ) with six degrees of freedom (DOF). As can be seen in the Figure 1b, joint $J0$ ("shoulder") was responsible for the rotation around the vertical axis, joint $J2$ ("elbow"), joint $J1$ ("shoulder") and joint $J3$ ("wrist") were responsible for the up and down movements; joint $J4$ ("wrist") rotated the gripper around the horizontal axis. The additional DOF came from the gripping manipulator.

- *Color stereo active vision system.* Two frame grabbers were used to digitalize images with a resolution of 128x128 pixels, down sampled at a rate of 20Hz.

- *Sensory-motor control board.* The communication between the computer and the motor control board that drives the active vision system and gets the tactile information was via a USB controller based on the Hitachi H8 chip.

- *System architecture.* The system architecture was composed of two computers Pentium III/600 MHz and the robot arm controller connected together in a private local area network based on the TCP/IP protocol, one computer controlled the robot arm and the other acquired the tactile input as well as the visual input from the active vision system.
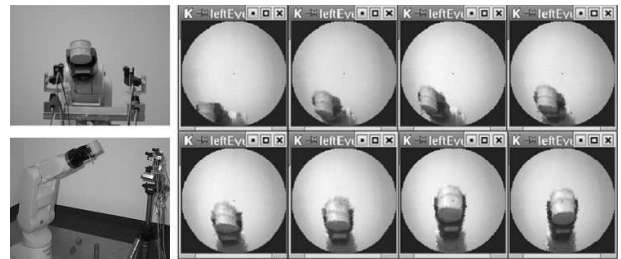

Figure 2. Robotic setup performing an experiment moving an object from the bottom-left corner of its visual field to the center of it. The observer's perspective can be seen on the left side, while the robot's perspective is shown on the right side.

3 Task specification

The task of the robot was to learn how to bring a colored object from the periphery of the visual field to the center of it by means of its robotic arm. It is important to note that although it would have been possible to program the robot directly to perform this task, our aim here is to quantify

the effects of developmental changes on the learning performance. We are not seeking biological plausibility, but biologically inspired mechanisms of adaptive and autonomous behavior.

At the outset of each experiment the active vision system was initialized looking at the center of the visual scene (x_c, y_c) and the position of its motors were kept steady throughout the operation. 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. Once the object was detected by the pressure sensors the robot started to learn how to move the arm in order to bring the object from the periphery of the visual field (x_0, y_0) to the center of it (x_c, y_c). In other words, the eyes should teach the robot arm to solve the task, the object was the visual stimulus and the way to solve the task was the movement of the robot arm. A typical experiment is shown in Figure 2. For more details see [5, 6].
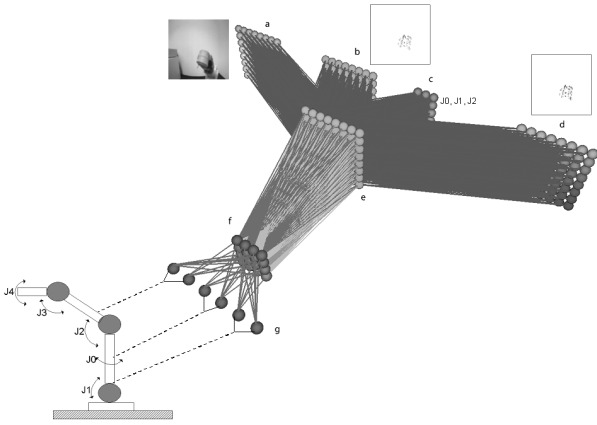
Figure 3. Neural structure and its connections to the robot's sensors and motors. Neuronal areas: (a) RedColorField. (b) RedMovementToRightField. (c) ProprioceptiveField. (d) RedMovementToLeftField. (e) NeuronalField. (f) MotorField. (g) MotorActivites.

4 Neural control architecture

The components of the neural structure and its connections to the robot arm are depicted in Figure 3.

4.1 Sensory field

- *Color information*. Three receptor types are considered: red (r), green (g), and blue (b). A "broadly" color-tuned channel was created for red:

$$R = r - (g + b)/2 \qquad (1)$$

This channel yields maximum response for the fully saturated red color, and zero response for black and white inputs. The negative values were set to zero. Each pixel was then mapped directly onto the 8x8 neuronal units of area *RedColorField* (see Figure 3a). The activity S_i of the i-th neuron of this area was calculated as:

$$S_i = \begin{cases} 1.0 & : \quad if \ R_i > \theta_1 \\ 0.0 & : \quad otherwise \end{cases} \qquad (2)$$

Where R_i is the value of the red color-tuned channel for the i-th pixel; and θ_1 is a threshold value.
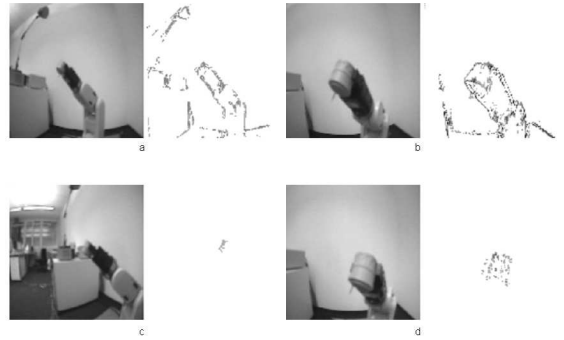
Figure 4. Motion detection. (a) Movement was detected from right to left. (b) Movement was detected from left to right. (c) and (d) Motion detectors reacting only to red objects moving in the environment.

- *Motion detection*. Motion detectors were created to detect movements of red objects in the environment. These motion detectors are based on the well-known elementary motion detector (EMD) of the spatio-temporal correlation type [10], a description of the model implemented, can be found in [7]. Motion detectors reacting to red objects moving to the right side of the image were mapped directly to neuronal units of the area *RedMovementToRightField* (see Figure 3b) and the motion detectors reacting to red objects moving to the left side of the image were mapped directly to neuronal units of the area *RedMovementToLeftField* (see Figure 3d). Both neuronal areas have a size of 8x8. The activities of the neurons in these areas were calculated as:

$$S_i = \begin{cases} 1.0 & : \quad if \ |EMDOutput_i| > \theta_2 \\ 0.0 & : \quad otherwise \end{cases} \qquad (3)$$

Where S_i is the activity of the i-th neuron; $EMDOutput_i$ is the output of the motion detector at position i-th; and θ_2 is a threshold value.

- *Proprioceptive information.* The movements of each joint of the robot arm were encoded using eight neuronal units. During the experiments the size of the neural area *ProprioceptiveField* (see Figure 3c) was increased. The minimum size was 8x1 when it encoded the joint: $J0$, it had a medium size of 8x2 when it encoded the joints: $J0$ and $J2$, and it had a maximum size of 8x3 for encoding the joints: $J0$, $J1$, and $J2$. Joint $J0$ had a range of movements from -60 to 60 degrees, joint $J1$ moved in a range from -25 to 25 degrees, and joint $J2$ moved in a range from 0 to 100 degrees.

4.2 Neuronal field and motor field

The size of the neuronal area *NeuronalField* (see Figure 3e) was 8x8 and its neuronal units had a sigmoid activation function.
During the experiments the size of the neuronal area *MotorField* (see Figure 3f) was increased. The minimum size was 4x4 and the maximum was 16x16 and its neuronal units had a sigmoid activation function whose outputs were passed directly to the *MotorActivities* (see Figure 3g) for controlling the joints of the arm: $J0$, $J1$ and $J2$. The size of the neuronal area *MotorActivities* was 6x1.

4.3 Synaptic connections

Neuronal units in the areas *RedColorField*, *RedMovementToLeftField*, and *RedMovementToRightField* were connected retinotopically to the neuronal units in area *NeuronalField*. The neuronal units in the area *ProprioceptiveField* were fully connected to the neuronal units in area *NeuronalField*. The neuronal units in area *NeuronalField* were fully connected to the neuronal units in area *MotorField*, which in turn were fully connected to the *MotorActivities*.

4.4 Learning Mechanism

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. In this way the synaptic connections between the neuronal areas *NeuronalField* (see Figure 3e) and *MotorField* (see Figure 3f) were changed. A learning cycle (i.e., the period during which the current sensory input is processed, the activities of all neuronal units are computed, the connection strength of all synaptic connections are computed, and the motor outputs are generated) had a duration of approximately 0.35 seconds. For more details see [3] and [5, 6].

5 Simulating development in a real robot

Because we are dealing with embodied systems, there are two dynamics, the physical one or body dynamics and the control one or neural dynamics. There is the deep and important question of how the two can be coupled in optimal ways. It has been hypothesized that given a particular task environment, a crucial feature of adaptive behavior is a balance between the complexity of an organism's sensor, motor, and control system (this is also referred to as principle of ecological balance) [13] and [12]. Here, we extended this principle to developmental time, and attempted to comply to it by simultaneously increasing the sensor resolution, the precision of the motors, as well as the size of the neural structure. Such concurrent changes are thought to simplify learning processes providing the basis for maintaining an adequate balance between the complexity of the three sub-systems, which reflects the development of biological systems.

## 5.1 Increasing the motor capabilities of the robot

The development of the robot's controllability was achieved by an initial freezing of mechanical degrees of freedom and gradual releasing of them. At the beginning only joint $J0$ was used, during the second developmental stage two joints were used (i.e., $J0$ and $J2$) and during the third developmental stage three joints were used (i.e., $J0$, $J1$, and $J2$).

a. Left camera

b. Right camera
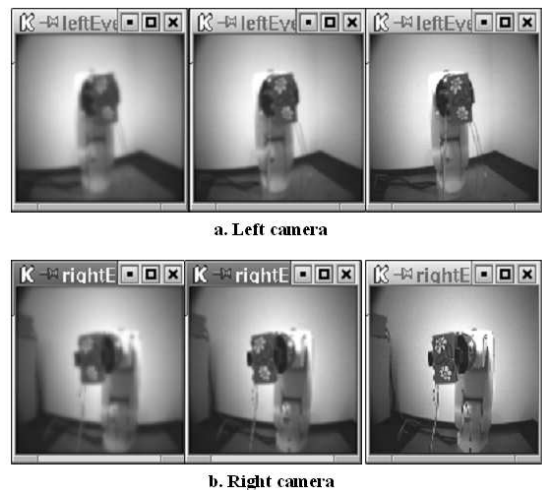
Figure 5. Gradual Increase of the sensory resolution. From left to right the image develops from blurred to high resolution.

5.2 Increasing the sensory capabilities of the robot

Increasing the resolution of the cameras was achieved by means of a gradual increase of the sharpness of a Gaussian blur lowpass filter applied to the original image captured by the cameras (see Figure 5(right)). Figures 5(left) and 5(center) show the result of applying a 5x5 and a 3x3 Gaussian kernel to the original image respectively.

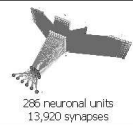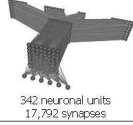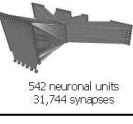The number of pressure sensors mounted on the gripper of the robot was also increased over time.
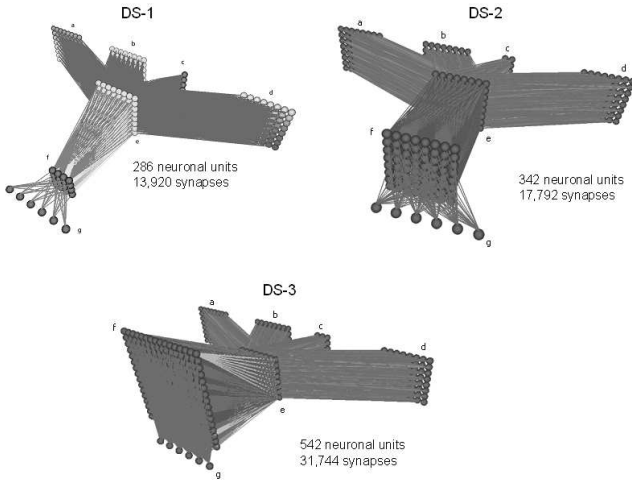
Figure 6. Gradual increase of the neural structure to cope with more sensory input and with more degrees of freedom of the motor system.

5.3 Increasing the complexity of the neural structure

In Figure 3 an overview of the neural network and its connections to the sensory-motor system is given. The neural network was gradually enhanced to cope with more sensory input and with more degrees of freedom of the motor system by (a) adding eight neuronal units to the area *ProprioceptiveField* (see Figure 3c) in order to encode another DOF and (b) making the size of the neuronal area *MotorField* (see Figure 3f) four times larger. The new weights were initialized randomly and the old weights were kept at their current values in order to preserve the previous knowledge acquired by the robot. The process is shown in Figure 6 and summarized in Table 1.

6 Developmental schedule

Development, in contrast to mere learning, implies on the one hand changes in the entire organism (not only

Figure 7. Configuration of the sensory, motor and neural components of the robot through the developmental approach. From top to bottom: DS-1 (immature state), DS-2 (intermediate state) and DS-3 (mature state).

the neural system) over time, and on the other hand a long-term perspective. The robot's movements were continuously shaped by the aforementioned learning mechanism, and "developmental" changes were triggered by the robot's internal performance evaluator (see definition of index "P" for the robot's task performance in Section 7). Such changes consisted in advancing the present developmental stage (DS-$_i$) to the next one. We defined a set of three different developmental stages (DS) in which the robot "grew up" as follows:

6.1 Developmental stage number 1 (DS-1)

At this stage, the sensory input to the robotic agent's neural structure consisted of a blurred, low resolution image (a 5x5 Gaussian kernel was applied to the original image captured by the cameras, see Figure 5(left)), and the activity of one pressure sensor. The neural network had 286 neuronal units and 13,920 synaptic connections, and controlled one single degree of freedom (i.e., joint $J0$). This developmental stage corresponds to the "*immature*" state of the robot. See Figure 7(DS-1).

6.2 Developmental stage number 2 (DS-2)

At this stage the robotic agent consisted of a medium level blurred image (a 3x3 Gaussian kernel was applied to the original image captured by the cameras, see Figure 5(center)), two pressure sensors, two DOF (i.e., joint $J0$ and $J2$), and the neural network had 342 neuronal

units and 17,792 synaptic connections. This corresponds to the "*intermediate*" state of the robot. See Figure 7(DS-2).

6.3 Developmental stage number 3 (DS-3)

At this stage the robotic agent consisted of the full high resolution image from the cameras (see Figure 5(right)), four pressure sensors, three DOF (i.e., $J0$, $J1$ and $J2$), and the neural network had 542 neuronal units and 31,744 synaptic connections. This corresponds to the "*mature*" state of the robot. See Figure 7(DS-3).

6.4 Control setup

The control setup had the same configuration of the fully matured robotic agent at stage number 3.

The schedule on how the robot was changed over time was determined by the learning mechanism, every time that the robot was considered to have learned to solve the task its configuration was changed moving from one developmental stage to the next one. This was achieved as follows:

- The resolution of the camera image was increased.

- one or two pressure sensors were added.

- another degree of freedom came into operation and the size of the neuronal area: "*ProprioceptiveField*" (see Figure 3c) was increased in 8 neuronal units.

- the size of the neuronal area: "*MotorField*" (see Figure 3f) was increased by a factor of four, the new weights were initialized randomly and the old weights were kept at their current values in order to preserve the previous knowledge acquired by the robot.

Figure 7 presents a summary of the configuration of the robot at each developmental stage. The number of neuronal units in each neuronal area at each developmental stage can be found in Table 1.

Through this simulated development (from DS-1 to DS-3) the initial setup with reduced visual capabilities, noisy motor commands, low number of degrees of freedom, a few pressure sensors and a neural control architecture with a reduced number of neuronal units, was converted into an experimental setup with good vision, larger number of degrees of freedom, larger number of pressure sensors and a neural control architecture with a sufficient number of neuronal units.

At developmental stage number 3, the robotic agent reaches the same sensory, motor and neural configuration than the control setup. At this point, their performances could be

Table 1. Neural structure at each developmental stage

Neuronal Area	stage 1	stage 2	stage 3
RedColorField	64	64	64
RedMovementToRightField	64	64	64
ProprioceptiveField	8	16	24
RedMovementToLeftField	64	64	64
NeuronalField	64	64	64
MotorField	16	64	256
MotorActivites	6	6	6
Total neuronal units	286	342	542

compared to see whether the learning was affected or not by the developmental approach described above.

7 Experiments and results

Figure 8 shows a typical experiment where the robot learned to move the object from the periphery of its visual field to the center of it by means of its robotic arm. To evaluate the change of the robot's task performance over time, at each time step i, we computed the cumulated distance covered by the center of the object projected onto one of the robot's cameras (x_i, y_i):

$$\hat{S} = \sum_{i=0}^{N-1} \sqrt{((x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2)} \qquad (4)$$

Thus, (x_0, y_0) is the initial position of the object as perceived by the robot, and (x_N, y_N) = (x_c, y_c) is the center of the robot's visual field (assuming that the robot learns to perform the task).

The shortest possible path between (x_0, y_0) and (x_c, y_c) is defined as:

$$S = \sqrt{((x_0 - x_c)^2 + (y_0 - y_c)^2)} \qquad (5)$$

By using S and $\hat{S}$, we defined an index for the robot's task performance:

$$P = \frac{S}{\hat{S}} \qquad (6)$$

The closer P is to 1, the more straight the trajectory, and therefore the better the robot's behavioral performance.

Figure 9 shows how the robot's behavior improved over time for the last part of the experiment number 1 (see Figure 8 interval d.) and gives the performance measure over time.
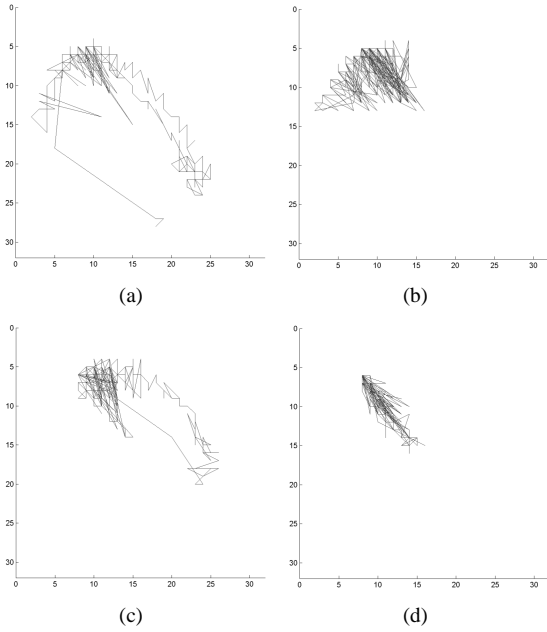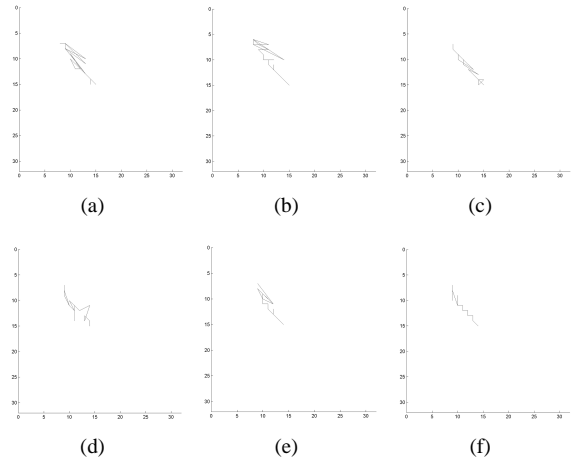
Figure 8. Experiment number 1. Learning to move a colored object from the upper left corner of the visual field to the center of it. Position of the center of the object in the visual field during the learning cycles in the interval (a) [1, 400]. (b) [401, 800]. (c) [801, 1200]. (d) [1201, 1602].

Figure 9. Robot's internal performance evaluator *"P"* during the learning cycles in the interval (a) [1232, 1266], *P*=0.2898; (b) [1313, 1340], *P*=0.3574; (c) [1370, 1393], *P*=0.5114; (d) [1438, 1455], *P*=0.5402; (e) [1502, 1519], *P*=0.6569;(f) [1565, 1582], *P*=0.9176. (see Figure 8d).

A total of 15 experiments were performed with two types of robotic agents: one subjected to developmental changes (i.e., DS-1, then DS-2 and finally DS-3), and one fully developed since the onset (control setup). The results clearly show that the robotic agents that followed a developmental path took considerably less time to learn to perform the task. These robotic agents started with the configuration of the developmental stage number 1 and learned to solve the task during the learning cycle 483 ± 70 (where $\pm$ indicates the standard deviation), then they were converted to robotic agents with a configuration as described by the developmental stage number 2 which subsequently learned to solve the task around the learning cycle 1671 ± 102 and finally they become to be in the developmental stage number 3 (with the same configuration than the control setup) and solve the task around the learning cycle 4150 ± 149 (this is a cumulative value).

The control setup agents with full resolution camera images, four pressure sensor, three DOF (i.e., $J0$, $J1$ and $J2$), and a neural network with 542 neuronal units (randomly initialized synaptic connections) learned to solve the task around the learning cycle 7480 ± 105.

In other words, a reduction of about 44.5 percent in the number of learning cycles needed to solve the task can be observed in the case of robotic agents that followed a de-velopmental approach when compared to the control setup agents.

8 Discussion and conclusions

We set out to investigate if the immaturity of sensory, motor, and neural system, which at first sight appears to be an inadequacy, might speed learning and task acquisition. In other words, we hypothesize that rather than being a problem, immaturity might effectively decrease or even eliminate excessive information and its potentially detrimental effects on learning performance.

This might be indeed the case as shown by the results presented in this paper. A system starting with low resolution sensors and low precision motor systems, whose resolution and precision are then gradually increased during development, learns faster than a system starting out with the full high resolution high precision system from scratch. For this particular case, by employing a developmental approach the learning was speeded up by 44.5 percent. To our knowledge this is the first time that this point is actually shown in a quantitative way.

There is a trade-off between finding a solution following a developmental approach and the potentially better solution, when starting out from the full high resolution high precision system from scratch.

Important is to keep in mind that the motor abilities should be increased gradually with the sensor abilities, since this significantly reduces the learning problem.

9 Future research

We will add proprioceptive information about the position of each motor of the active vision system and one possible task for the robot would be to not only bring the object to the center of the visual field, but also to normalize the size of the object in the camera image (i.e., a big object would be presented by the arm to the cameras further away than a smaller one) providing the robot with an "Embodied concept of size". In a future set of experiments we will put the developmental schedule under the control of an Artificial Evolutionary System.

Acknowledgments

References

[1] N. Bernstein. *The coordination and regulation of movements.* Pergamon, Oxford, England., 1967.

[2] L. Berthouze and M. Lungarella. Motor skill acquisition under environmental perturbations: on the necessity of alternate freezing and freeing of degrees of freedom. *Adaptive Behavior (To appear)*, 2004.

[3] P. Eggenberger Hotz, G. Gómez, and R. Pfeifer. Evolving the morphology of a neural network for controlling a foveating retina - and its test on a real robot. In *Standish, R. K., Bedau, M. A., and Abbass, H. A., editors, Artificial Life VIII: Proceedings of the 8th International Conference on the Simulation and Synthesis of Living Systems, Sydney, Australia*, pages 243–251, 2002.

[4] J.L. Elman. Learning and development in neural networks: The importance of starting small. *Cognition*, 48:71–99, 1993.

[5] G. Gómez and P. Eggenberger Hotz. An evolved learning mechanism for teaching a robot to foveate. In *Sugisaka Masanori and Tanaka Hiroshi, editors (AROB 9): Proceedings of the 9th Int. Symp. on Artificial Life and Robotics, Beppu, Oita, Japan*, pages 655–658, 2004.

[6] G. Gómez and P. Eggenberger Hotz. Investigations on the robustness of an evolved learning mechanism for a robot arm. In *In Groen, F., Amato, N., Bonarini, A., Yoshida, E., and Krose, B., editors (IAS 8): Proceedings of the 8th International Conference on Intelligent Autonomous Systems, Amsterdam, The Netherlands*, pages 818–827, 2004.

[7] F. Iida. Biologically inspired visual odometer for navigation of a flying robot. *Robotics and Autonomous Systems*, 44:201–208, 2003.

[8] M. Lungarella and L. Berthouze. On the interplay between morphological, neural and environmental dynamics: a robotic case-study. *Adaptive Behavior*, 10:223–241, 2002.

[9] M. Lungarella, G. Metta, R. Pfeifer, and G. Sandini. Developmental robotics: a survey. *Connection Science 15*, 4:151–190, 2003.

[10] David Marr. *Vision. A Computational Investigation into the Human Representation and Processing of Visual Information.* W.H. Freeman and Company, 1982. MAR d 82:1 1.Ex.

[11] Y. Nagai, K. Hosoda, S. Morita, and M. Asada. A constructive model for the development of joint attention. *Connection Science 15*, 4:211–229, 2003.

[12] R. Pfeifer, F. Iida, and J. Bongard. New robotics: design principles for intelligent systems. *Artificial Life Journal (To appear)*, 2004.

[13] Rolf Pfeifer and Christian Scheier. *Understanding Intelligence*. MIT Press, 1999.

[14] A. Slater and S. Johnson. Visual sensory and perceptual abilities of the newborn: beyond the bloomig, buzzing confusion. In *Simion, F. and Butterworth, G., editors, The Development of Sensory, Motor and Cognitive Capabilities in Early Infancy: From Sensation to Cognition. Hove, Psychology Press*, pages 121–141, 1997.

[15] E. Thelen. Dynamics mechanism of change in early perceptuo-motor development. In *In McClelland, J. and Siegler, S., editors, Mechanims of cognitive development: Behavioral and neural perspectives. Proceedings of the 29th Carnegie Symposium on Cognition.*, 1999.

[16] G. Turkewitz and P.A. Kenny. Limitation on input as a basis for neural organization and perceptual development: A preliminary theoretical statement. *Developmental Psychobiology*, 15:357–368, 1982.