Berthouze, L., Prince, C. G., Littman, M., Kozima, H., and Balkenius, C. (2007).
Proceedings of the Seventh International Conference on Epigenetic Robotics: Modeling
Cognitive Development in Robotic Systems. *Lund University Cognitive Studies, 135*.

Developing Sensorimotor Associations Through Attachment Bonds

Antoine Hiolle and Lola Cañamero
Adaptive Systems Research Group
School of Computer Science
University of Hertfordshire
College Lane, Hatfield, Herts AL10 9AB, UK
{A.Hiolle,L.Canamero}@herts.ac.uk

Abstract

Attachment bonds and positive affect help cognitive development and social interactions in infants and animals. In this paper we present a neural architecture to enable a robot to develop an attachment bond with a person or an object, and to discover the correct sensorimotor associations to maintain a desired affective state of well-being using a minimum amount of prior knowledge about the possible interactions with this object. We also discuss how our research on attachment bonds could further developmental robotics in the near future.

1. Introduction

The question of how autonomous robots could be integrated in our everyday life is gaining increasing attention. To that end, robots will have to be able to exhibit adaptive and complex behaviors, and our view is that they should be able to learn without constant instruction from a teacher, and rather develop in interaction with humans and learn from this interaction (Cañamero et al., 2006). Robots will need to constantly learn how to react in different situations and environments with a minimal amount of prior knowledge present in their behavioral systems. A key element towards this goal is the integration of emotional and affective factors in these interactions (Cañamero, 2001, Breazeal, 2003) as a way to guide development and learning. Adding emotional values to different contexts is for example a way to facilitate decision-making (Blanchard and Cañamero, 2005). The formation of attachment bonds is another important aspect, not only to improve human-robot interaction but also as a way to develop further cognitive and emotional capabilities (Nadel and Muir, 2005). According to Bowlby's theory (Bowlby, 1969), a secure attachment bond helps infants during their development. It is known to foster exploratory behaviors, which are essential for the infant to build a coherent and stable internal model of the environment. Furthermore, as stressed in (Kaplan, 2001), it would be possible to build architectures for autonomous robots that could allow us to compare and study the consequences on the development of such attachment bonds using the "Strange Situation Test" (Ainsworth, 1969). From a human-robot interaction point of view, a robot that explores its environment with confidence thanks to its history of affective interactions with humans has the advantage of being self-driven since the robot would have an internal motivation urging it to discover and later understand its environment. A successful robotic implementation of an early model of attachment and its implication in exploratory behavior was presented in (Blanchard and Cañamero, 2005, Blanchard and Cañamero, 2006). This work took inspiration from the imprinting phenomenon first described by Konrad Lorenz in the case of birds (Lorenz, 1935). During the early days of life, an attachment bond develops between young birds and persons or objects to which the animals have been exposed. As a consequence, the birds follow the movements of the imprinted object or person. In this early attachment experience, the imprinted object acts as a sort of security mechanism for birds during exploration; moreover, the simple fact of following the imprinted object helps them discover their environment faster and without any explicit teaching by the imprinted object or person. Modeling this phenomenon with autonomous robots showed that they could benefit from the advantages provided by imprinting to guide their first steps in an unknown environment and as a mechanism to bootstrap affective interactions with humans. However, in our previous model, the robot had already hardcoded or "pre-wired" in its system the know-how to follow the imprinted object. From an epigenetic perspective of development, letting the robot discover and learn by itself how to maintain the imprinted perception—being at the "right" distance from the imprinted object in our case—would be a more plausible approach to model early attachment in humans and other complex mammals, which is closely related to imprinting in birds but slightly different. Indeed, in more complex species in which newborns are less developed when they leave the maternal environment, learning from experience and interactions with the environment plays a crucial role to achieve normal development. In the remainder of this paper, we present such an architecture that allows a robot to imprint a person (or a moving object) present in front of it when it is turned on and then to learn, without any external reinforcement, how to follow the imprinted object. We tested this architecture using two types of robots—an Aibo and a Koala—and here we present and discuss in detail the results obtained in the latter experimental setting.

2. Robot Architecture

Our architecture follows a "Perception-Action" approach (Gaussier and Zrehen, 1995), which postulates that perception and action are tightly coupled and coded at the same level. Action is thus executed as a "side-effect" of wanting to achieve, improve or correct some perception. The perception-action loop can be seen in terms of homeostatic control, according to which behavior is executed to correct perceptual errors. Actions that allow to correct different perceptual errors are selected on the grounds of sensorimotor associations that can be "hardcoded" by the designer (e.g., in a look-up table, as in (Gaussier and Zrehen, 1995, Blanchard and Cañamero, 2005) or learned from experience by the robot, as it is our case here—our robot extracts sensorimotor associations led by its motivation to keep the imprinted object at a constant distance, and using a combination of associative learning and action selection. We have also taken inspiration from (Panksepp, 1998) regarding ideas on the relation between affective states and homeostasis. Figure 1 shows the components of our architecture, implemented using a neural network consisting of neural groups that fulfill different functions, as explained in the remainder of this section.

2.1 Imprinting System

The imprinting system learns the value of the distance sensors. This neural group (Imprinted Perception in Figure 1) contains as many neurons as the number of distance sensors used (in this case two, as shown in Figure 2), which outputs equal to the learned distance value. The imprinting group uses a modified Rescorla-Wagner conditioning rule (Rescorla and Wagner, 1972) with a decreasing

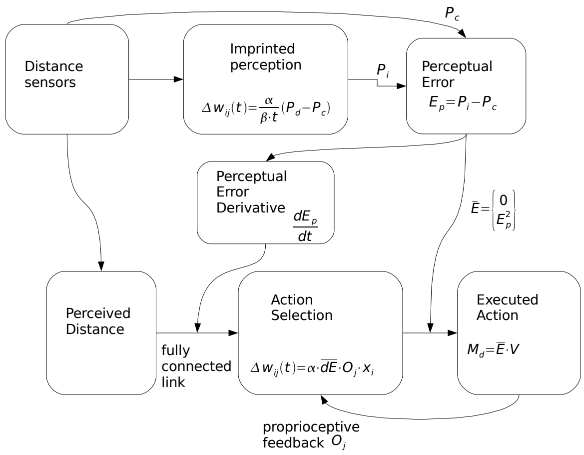

Figure 1: Our architecture for imprinting.

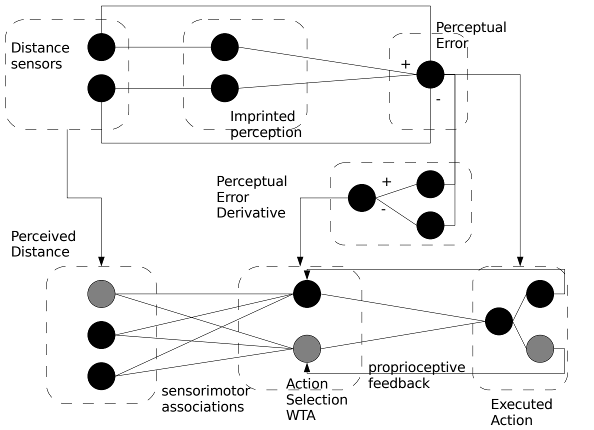

Figure 2: Our detailed architecture with the number of neuronal units in each group.

global learning rate to achieve stabilization:

$$w_{ij}(t) = w_{ij}(t-1) + \frac{\alpha}{(\beta \cdot t)} \cdot (P_d - P_c) \qquad (1)$$

with:

$w_{ij}(t)$ the weight of the link between input neuron i of the Distance Sensors group and neuron j of the Imprinted Perception group.

α the learning rate here equal to 0.2

β the learning rate's decay rate equal to 0.05

P_d the current output of the ith neuron of the Distance Sensors group

P_c the current output of the jth neuron of the Imprinted Perception group

When the global learning rate, $\frac{\alpha}{(\beta \cdot t)}$, reaches a value below 0.001, the output of the imprinting group remains unchanged until the end of the experiment, thus achieving stability in the computation of the perceptual error and its derivative.

2.2 Perceptual Categorization System

Our robot must learn to associate its relative position with respect to the imprinted object to the action to be taken to correct its perceptual error. For this, it must first calculate its perceptual error, then evaluate in which category of perception it is (Distance perceived), to be able to choose the right corrective action.

To modulate the response of the system according to the discrepancy between the current perception and the imprinted one, we compute the current perceptual error (E_p) between the imprinted perception (a distance) and the current one:

$$E_p = \sum P_i - P_c \tag{2}$$

with:

P_c the current perception value (the current value of the distance sensors in the case of the Koala setup)

P_i the imprinted perception value (the value of the distance sensors during the imprinting phase)

The neuronal group computing these values has 1 output neuron for the perceptual error, E_p and is linked to the perceptual error derivative group which contains three units, two that are computing an average of the perceptual error on two different time windows and the discrepancy between these two units is used to evaluate the derivative of the perceptual error. First the perceptual error is thresholded like :

$$\bar{E_p} = \begin{cases} 0 & \text{if } |E_p| < \theta_1 \\ E_p & \text{otherwise} \end{cases}$$

where θ_1 is chosen to provide an interval where the system considers its perception to be the correct one, i.e. the imprinted perception.

We now use this value to evaluate the two average values of the error on two time windows :
$e(\tau) = \frac{e(\tau-1)\cdot\tau+\bar{E_p}}{\tau+1}$ the average value of $\bar{E_p}$ over τ time steps. Then we calculate the derivative of the error:

$$\frac{dE_p}{dt} = e(\tau_1) - e(\tau_2) \tag{3}$$

with $\tau_1 = 2$, and $\tau_2 = 4$, the two time windows.

Since we want our robot to be able to associate its relative position with respect to the imprinted object to the action to be taken in order to correct its perceptual error, we project the actual value of the distance sensors into three categories: too far from the object, too close to it, and correct distance (the distance for which $\bar{E_p} = 0$). Therefore this neural group (Distance Perception in Figure 1) contains three neurons, one for each category, and only one neuron is activated at each timestep. Although this categorization could have been achieved on line by the system itself, we decided to use a fixed one in this case in order to focus on the problem that is our object of study here—the perception-action pairing. The output of this neural group is used as input for the action selection one.

2.3 Action Selection and Learning

The task of the action selection module is to learn how to maintain the desired perception learned by the imprinting module. Therefore, it needs to select the correct action according to the actual perceived distance category. To this end, the latter is fully connected to a Winner-Take-All (WTA) group of neurons (Action Selection Group in Figure 1). This group receives also a modulatory input from the perceptual error group, dE_p, and proprioceptive feedback from a motor output group which displays the real action that has been executed. This signal acts as the teaching signal for the learning module. The input dE_p is used as a kind of reinforcer helping the system to learn associations between the active perceptual category and the action that has been produced. The association between a perceptual category (Perceived Distance in Figures 1 and 2) and an action that makes the perceptive error decrease ($dE_p < 0$) will be strengthened, whereas the association between a perception and an action that makes the perceptive error increase ($dE_p > 0$) will be weakened. The initial weights between the perceptual categories and the WTA are initialized to small random values. The WTA group contains two output neurons, one for the action of going forward and one for going backward. Hence the WTA group learns using a modified Hebbian rule and produces outputs as follows:

$$w_{ij}(t) = w_{ij}(t-1) - \alpha \cdot dE_p \cdot O_j \cdot x_i \tag{4}$$

with:

$w_{ij}(t)$ the weight of the link between input neuron i of the distance system group and neuron j of the WTA action selection group initialized with random positive values between 0 and 1.

α the learning rate, here equal to 0.2

dE_p the derivative of the perceptual error

O_j the proprioceptive feedback from the motor output group

x_i the output value of the ith neuron of the Perceived Distance group

The group (Executed Action) uses the output of the WTA to compute the speed of the robot, in this case the direction of the movement. However, if the perceptual error is null, we want the system to remain static, as in (Blanchard and Cañamero, 2006). For this, we use the value of $\bar{E}$ to modulate the value of the motor output. The motor output value is a real number, and will have the effect of going forward when positive, backward when negative. In order to avoid abrupt changes in the speed of the robot, we need to produce a smooth motor output; the value of the motor output is filtered as follows:

$$M(t) = M(t-1) + \alpha(M_d - M(t-1)) \tag{5}$$

with the selected motor output M_d computed as:

$$M_d = \bar{E} \cdot V \tag{6}$$

where V, the current direction of the robot, equals -1 when going backwards, 1 when going forward. This value is directly computed using the outputs of the WTA group. The two output neurons of this group are then updated to provide feedback to the Action Selection group, so if the current motor output M_d is negative, one neuron is going to be activated and if it is positive, the other one will be activated; both neurons will be inhibited if the perceptual error is null.

At the beginning of the "life" of the robot (for a short period after it is turned on) no action is taken in order to allow imprinting to take place. Then the system works in two phases. During the first (action selection) phase, all groups have their outputs updated and then an action is executed. During the second (learning) phase, the perceptual error and its derivative are updated, and the WTA Action Selection group learns the consequences of its last action—the weights from the Perceived Distance group to the Action Selection are updated.

3. First Experiments and Results

3.1 Experimental setup

To test our system, we used two different types of robots and settings: an Aibo, Figure 3, and a Koala Figure 4. In both cases, we used a one-dimensional task in which the robot became imprinted to an experimenter playing the role of a caretaker placed in front of it. In the case of the Aibo, using the camera, the robot became imprinted to a ball held and moved by the experimenter and it learned to follow the ball with movements of its head, while attempting to correct the perceptual error—the difference between the actual position of the ball in its visual field, and the position it had when it was imprinted. Using the Koala, the experimenter was standing and moved forward and backwards in front of the robot; the Koala used its infrared sensors to detect the experimenter and had to learn to move with (follow or back up from) the experimenter while trying to maintain the distance at which it had been imprinted. In this paper we report our experiments and results using the Koala scenario.

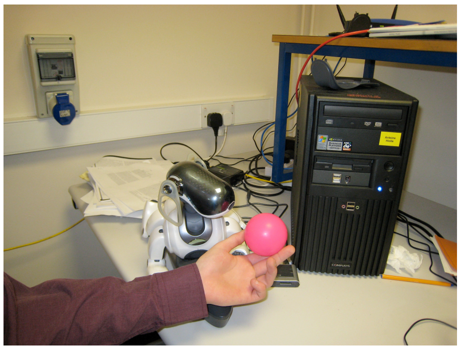

Figure 3: Setup using Aibo learning to follow laterally the ball.

Figure 4: Setup using a Koala robot learning to follow the experimenter.

The experiment starts by turning on the robot in front of the caretaker, and none of them moves[1] for a small period of time during which the initial imprinting takes place. After this phase, if the caretaker doesn't move, the perceptual error of the robot remains equal to 0 and therefore no movement is produced. Then, after a few seconds, the caretaker

[1] We haven't discussed here the case in which the experimenter moves during the imprinting period, the obvious result being that the imprinted perception stabilizes at an average value of the distances experienced during this period.

moves away from the robot. The robot will then execute the action selected as winner output by the Action Selection group. If the action executed makes the perceptual error decrease, then the robot will learn that this action is the correct one to execute in that situation—in this particular example approach the caretaker, resulting in a following behavior. If the action executed is not the correct one, after few timesteps the robot will choose to execute another action and, if it corrects the perceptual error, it will learn that it is the correct action to execute in that situation.

During this experience, we recorded the values of the distance sensors, the perceptual error between the desired perception and the current one, the derivative of the latter and the values of the weights between the categorized perception (Distance Perceived) and the action to do. Figure 5 shows an example in which the caretaker approached the robot, getting closer than the distance the robot was imprinted to. As we can see, the weight value, associated here with the action of backing up from the caretaker, increased correctly during the experiment. More specifically, if we look closer in the rectangular boxes labeled 1 in the figure, we can observe that the weight value increases when the derivative of the perceptual error is negative, which happens when the square perceptual error decreases—in this case, when the caretaker slowly approached the robot. The caretaker then stopped moving, the robot went slowly backwards and, since the derivative of the error was negative, the association between this situation and the action of going backwards was strengthened, and the robot reached again the desired perception and stopped moving. When the caretaker tried again to move closer to the robot—the moment inside the rectangular box labeled $1'$—the robot quickly reached the desired perception again, showing us that the correct association had been learned. The caretaker then moved closer to the robot again, but this time very quickly. We can see in the box labeled 2 that inducing this quick perturbation provoked a decrease in the association learned, the weight value decreases first. But since no other action had been associated as the correct action in that situation, the robot moved backward again, and the association is again reinforced with an even higher value than before. The same effects were observed with the opposite perturbation—the caretaker quickly moving away from the robot. It is interesting to note that our learning system is influenced by the intensity of the perturbation and its length. If the experimenter were to move further and further away from the robot, this system would not be able to learn how to follow her/him. But in a step-by-step manner, it learns with increasing accuracy the correlations between the perception it wants to reach and the correct action to do.

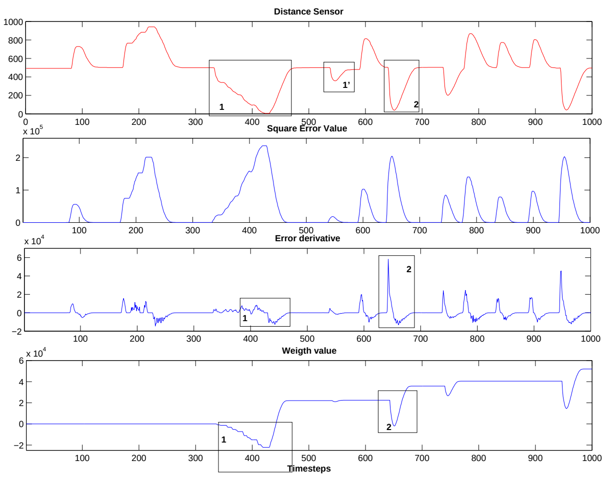

Figure 5: Evolution of (from top to bottom): perception of the distance, square and derivative of the perceptual error, and the association weight between perception and action, producing in this case the behavior of backing up from the imprinted object as it gets too close.

4. Increasing Complexity: Experiments and Results

After we successfully tested our architecture on a one-dimensional problem, we modified the experiment to test how the capabilities of such a system could be extended. We added a second dimension to the experiment by allowing the robot to turn. Since the purpose of these experiments was to provide the smallest amount of prior knowledge, duplicating the whole system on the other dimension would not be an adequate challenge for the system. Therefore we modified the system components as follows:

- The sensory input is a vector including four components—the values of the two frontal sensors and of the two side sensors; thus the Sensory Group and the Imprinting Group have each four components now.
- The perceptual categorization (previously Perceived Distance) group has 5 categories; the two new ones are activated by the sign of the difference between the two sides sensors.
- There are four possible actions in the Action Selection group—going backward, forward, turning left and turning right.
- The Executed Action is still chosen according to the Action Selection group winner, but the inhibiting system is also driven by the perceptual error derivative. When the squared perceptual error ($\bar{E}$) is null, all motor actions are inhibited

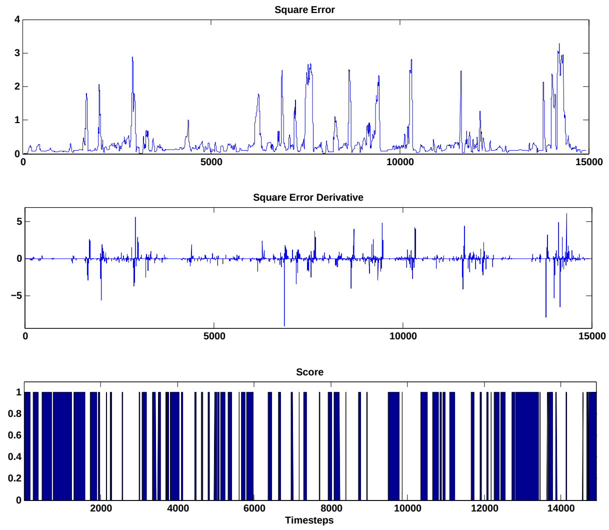

Figure 6: Evolution of the perceptual square error, its derivative, and score of the robot in reaching the correct perception(from top to bottom).

> but when the absolute value of the derivative of the perceptual error is below a low threshold, the other degree of freedom of the action that has not yet been tried—frontal or lateral—is the only one that can be activated.

We introduced this modification in the inhibition of the actions in order to let the robot switch to other actions when no progress in reaching the homeostasis is experienced during a fixed amount of timesteps. This helps the robot try another action when the perceptual error is high and stable, e.g., when the caretaker/experimenter is out of the range of its sensors.

The experimental setup to test this was the same as in the one dimension experiment. The weights are initialized to random values within a close interval (between 0.2 and 0.6). We recorded the same data as in the first experiment. In Figure 6, we show the evolution of the squared perceptual error, the derivative of it, and what we call the score of the robot, which is 1 when the correct perception is reached, 0 otherwise. We can observe that the robot is still able to reach the correct perception rapidly. However, it now has several ways to achieve it and some configurations need the experimenter to move toward an easier position for the robot to solve the problem. For instance, when the experimenter backs away from the robot, depending on the initial weights, this sometimes causes the robot to alternate between turning left and right, which results in reaching the correct perception in a non-optimal manner. We observed the same effects when the experimenter moved to the sides: the non-optimal solution that appeared here was to go backwards for a few timesteps, then forward during the same amount of time, and the robot got stuck into a local minima where the error cannot be reduced to its ideal value below the fixed threshold, and since the derivative is still fluctuating the system does not switch the action to another degree of freedom. This also affects the stability of the association weights: like in the first experiment, when the experimenter moves, the perturbation rapidly induces a decrease of the weights and there is no guarantee that the robot will converge to the correct perception.

5. Related Work

The architecture presented develops simple sensorimotor associations using a modified Hebbian learning rule modulated by the derivative of the perceptual error. In previous experiments by Andry and colleagues (Andry et al., 2001), a similar architecture was used to learn new sensorimotor associations without any explicit reinforcement in the case of a teacher-student interaction. The system had to react to inputs produced by the teacher, with the sole guidance of the quality of its prediction of the rhythm of the teacher's input. The prediction error was then used to compute a reinforcer in order to signal the system whether the current sensorimotor associations were likely to be more or less accurate. More precisely, the temporal derivative of this reward was directly modulating the probabilities, named credibility, for each weight to change using the Probabilistic Conditioning Rule (Gaussier et al., 1997). In their contribution, Andry and colleagues note problem that using a simple correlation rule (e.g. a Hebbian rule) would require the teacher to wait until the student gives a correct output according to the input. They point that, since that was not the aim of their experiment, using the prediction of the rhythm to tune the sensorimotor associations would ease and speed up the interactions. However, in our experiment this type of learning rule is particularly suitable since we want to build the sensorimotor association according to the behavior of the caretaker. If the caretaker does not wait for the robot to acquire the correct associations, then this caretaker would not be able to stand as a good reference for the robot to explore and learn from its environment correctly; hence, the attachment bond with this particular caretaker would not be classified as a secure one according to Bowlby's categorization. Meanwhile it is interesting to see that using a reinforcer computed with the derivative of the perceptual error—in our case the perceived distance, in theirs the rhythm of the interaction—helps to solve sensorimotor associations learning problem during human robot interactions.

Furthermore, the use of the derivative of a prediction error has been described by Oudeyer and Kaplan as a way to help a robot exhibit a curious behavior in an open-ended environment. In (Oudeyer and Kaplan, 2004), these authors presented an architecture permitting the robot to choose to move towards a situation were actual learning could be achieved. To this end, the system learned from each situation (categorized in similar sensorimotor states) the value of its learning progress (e.g. the opposite of the derivative of its prediction error) and then decided which next states would assure the higher decrease of its prediction error. It was the value of this prediction error that acted as an internal reward. The architecture we present here uses a similar approach to correct the sensorimotor associations using the opposite of the derivative of the perceptual error, therefore, we can note that this kind of reinforcer can be applied to address several problems in autonomous robotics. Moreover, it might be possible to unify these results to come up with a more complete architecture allowing a robot to exhibit a curious behavior emerging from the experience of its interactions with the caretaker. Since in our experiment the robot is learning to correct its sensorimotor associations in order to maintain its proximity to the caretaker, in the meantime it is actually reducing its prediction error, even though the notion of prediction is not explicitly used in our system. Indeed, after having learned the correct associations and decreased its perceptual error, the robot has managed to reach a desirable position in terms of affective state. A simple associative architecture would thus allow it to learn that correcting its prediction error has lead to a desirable state, and that, continuing to decrease this predictive error would have a positive emotional value. Moreover, letting a robot develop an attachment bond with a caretaker as a desirable perception to reach to feel secure, could also help to deal with the problem of choosing between several situations that predict the same learning progress during the exploration phase. Indeed, if a conflict during exploration arises, instead of picking one action randomly, the robot could choose the action that would lead it to the perception that is the closer to the desirable one.

6. Conclusion and Future Work

We have shown that the system presented in this paper is able to learn the consequences of its actions led only by the tendency to maintain a perception associated with the presence of a caretaker. Our robot can learn to maintain a desired perception, and therefore to follow its caretaker around, without any prior knowledge of how to do so, rather than having this knowledge "pre-wired" like in previous and related work. We have seen that the robot learns fast relatively to the complexity of the task to learn. However, each time a perturbation in the homeostasis of the system is induced by the experimenter, there is no way for the robot to distinguish whether this perturbation is a consequence of its own action or not. That is why the association weights decrease during this perturbation. When the task to learn is more sensitive to these perturbations (i.e. when the weights of the associations have values which are very close among them), our system has to learn again the associations after each external perturbation. Adding a way for our robot to discriminate between perturbations due to external causes (e.g. the actions of the caretaker) or internal causes (typically the actions of the robot), although far from being a trivial problem, would be a natural future extension to our system from a developmental perspective.

This problem of external perturbations is also related to how caretakers respond to infants' demands. It seems natural that the experimenter acting as a caretaker would have to adapt his/her behavior to that of the robot. For example, if the caretaker were not to wait for the robot to learn how to follow him/her, we could say that the caretaker would not be responding correctly to the needs of the robot in terms of interactions. The appropriate behavior for the caretaker would be to wait for the robot to reach the desired perception and to have the time to learn how to reach that perception (i.e. which action to execute to do so in that situation), and to follow the caretaker at a constant distance, by trying different actions in a sort of "motor babbling", so that the robot can be in a good emotional state, without the distress of the absence of the caretaker. The interactions involved in these simple learning tasks are comparable to mother/infant interactions during the first year, and are particularly relevant to investigate Bowlby's notion of secure-insecure attachment (Bowlby, 1969) and its influence in the development of emotional and cognitive capabilities, such as openness towards the world and curiosity. The use of such an architecture to build stable and relevant low sensorimotor associations does not appear to be biologically plausible as it is. We studied this system in a real robot/human interaction in order to evaluate its dynamics. Using a real continuous value—the derivative of the perceptual error — to reinforce behavioral responses could be compared to the effects of specific chemicals released in the brain during infant/mother interactions. It is known that comfort and proximity of the mother induces significant release of opioids in the infants' brains (Panksepp, 1998), which inhibits the effects of cortisol, the hormone released during stressful episodes. These mother/infants interactions are believed to shape the brain areas responsible for coping with emotional situations (Schore, 2001). The behavioral responses the infant will exhibit are a consequence of this emotional learning process and the interactions between hormones and neural substrates. In future work, our architecture could prove itself useful in biasing more complex behavior, gazing at someone, exploring the environment or looking for the attachment figure.

Acknowledgements

We are grateful to Arnaud Blanchard for his feedback and support during the first stages of this work. This research is supported by the European Commission as part of the FEELIX GROWING project (http://www.feelix-growing.org) under contract FP6 IST-045169. The views expressed in this paper are those of the authors, and not necessarily those of the consortium.

References

Ainsworth, M. D. S. (1969). Object relations, dependency, and attachment: A theoretical review of the infant-mother relationship. *Child Development*, 40(4):969–1025.

Andry, P., Gaussier, P., Moga, S., Banquet, J.-P., and Nadel, J. (2001). Learning and communication in imitation: An autonomous robot perspective. *IEEE Transactions on Man, Systems and Cybernetics, Part A: Systems and humans*, 31(5):431–442.

Blanchard, A. and Cañamero, L. (2006). Modulation of exploratory behavior for adaptation to the context. In Kovacs, T. and J., M., (Eds.), *Biologically Inspired Robotics (Biro-net) in AISB'06: Adaptation in Artificial and Biological Systems.*, volume II, pages 131–139.

Blanchard, A. and Cañamero, L. (2005). From imprinting to adaptation: Building a history of affective interaction. *Proc. of the 5th Intl. Wksp. on Epigenetic Robotics*, pages 23–30.

Bowlby, J. (1969). *Attachment and loss*, volume 1:Attachment. New York : Basics Books.

Breazeal, C. (2003). Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies*, 59:119–155.

Cañamero, L. (2001). Building emotional artifacts in social worlds: Challenges and perspectives. *Emotional and Intelligent II: The Tangled Knot of Social Cognition.*

Cañamero, L., Blanchard, A., and Nadel, J. (2006). Attachment bonds for human-like robots. *International Journal of Humanoïd Robotics.*

Gaussier, P., Revel, A., Joulain, C., and Zrehen, S. (1997). Living in a partially structured environment: how to bypass the limitations of classical reinforcement techniques. *Robotics and Autonomous Systems*, 20:225–250.

Gaussier, P. and Zrehen, S. (1995). Perac: A neural architecture to control artificial animals. *Robotics and Autonomous Systems*, 16:291–320.

Kaplan, F. (2001). Artificial attachment : Will a robot ever pass ainsworth's strange situation test. In *Proceedings of Humanoids 2001 : IEEE-RAS International Conference on Humanoid Robots*, pages 125–132.

Lorenz, K. (1935). Companions as factors in the bird's environment. In *Studies in Animal and Human Behavior*, volume 1, pages 101–258. London: Methuen & Co., and Cambridge, Mass.: Harvard University Press.

Nadel, J. and Muir, D. (2005). *Emotional Development.* Oxford University Press.

Oudeyer, P.-Y. and Kaplan, F. (2004). Intelligent adaptive curiosity: a source of self-development. In Berthouze, L., Kozima, H., Prince, C. G., Sandini, G., Stojanov, G., Metta, G., and Balkenius, C., (Eds.), *Proc. of the 4th Intl. Wks. on Epigenetic Robotics*, volume 117, pages 127–130. Lund University Cognitive Studies.

Panksepp, J. (1998). *Affective Neuroscience.* Oxford University Press.

Rescorla, R. and Wagner, A. (1972). A theory of pavlovian conditioning: Variations in effectiveness of reinforcement and nonreinforcement. In Black, A. and Prokasy, W., (Eds.), *Classical Conditioning II*, pages 64–99. New York: Appleton-Century-Crofts.

Schore, A. (2001). Effects of a secure attachment relationship on right brain development, affect, regulation, and infant mental health. *Infant mental health Journal*, 22:7–66.