Conscientious Caretaking for Autonomous Robots An Arousal-Based Model of Exploratory Behavior

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

The question of how autonomous robots could be part of our everyday life is gaining increasing interest. We present here an experiment in which an autonomous robot explores its environment and tries to familiarize itself with its novel elements using a neural-network-based architecture. When confronted with novelty, the lack of stability of its learning structures increases the arousal level of the robot, pushing it to look for comfort from its caretaker in order to reduce this arousal. In this paper, we studied how the behavior of the caretaker—and in particular the amount of comfort it provides to the robot during its exploration of the environment—influences the course of the robot's exploration and learning experience. This work takes inspiration from early mother-infant interactions and the impact that the primary caretaker has on the development of children—at least in mainstream Western culture. The underlying hypothesis is that the behavior of a caregiver, and particularly his/her role in modulating arousal, will influence the development of an autonomous robot, and that arousal regulation will also depend on how accurately the robot signals its internal state and how the caretaker (or human user) responds to these signals.

1. Introduction

The question of how autonomous robots could be part of our everyday life is gaining increasing interest. Our philosophy towards the design of such robot takes an epigenetic approach (Cañamero et al., 2006). Indeed, this approach would help the robot discover and learn affordances in the environment in which it is situated, including the agents it interacts with, as opposed to an approach where the designed architectures would need prior knowledge about the environment and / or extensive explicit instruction by the human user. The main issue that needs to be addressed within this approach is what sort of "bootstrapping" or "built-in" mechanism a robot would need in order to be able to develop its cognitive and social capabilities. To be more precise, what are the "inner drive(s)" and basic principle(s) that will push the robot towards situations in which it will learn what it needs to know in order to be fully operational in the given environment? This problem has many similarities with the development of infants.

The literature in developmental psychology suggests that caretaker-infant attachment bonds are vital to the cognitive and emotional development of infants, see e.g., (Hofer, 2006), especially during the first years of life. Indeed, as John Bowlby (Bowlby, 1969) discovered during his studies of mother-infant interactions, the adult who plays the role of primary caretaker in Western cultures, usually the mother, is utilized by the infant as a secure base in his/her early life, particularly during stressful and/or unusual episodes (Sroufe, 1995). Furthermore, as stressed in (Schore, 2001), if the primary caretaker doesn't act according to the infant's demands in term of interactions, the mental development of the child can be impaired, leading to emotional and cognitive disorders. Therefore, identifying the factors that are particularly relevant during these interactions, as well as their dynamics, is important to understand how the development of a child can lead to many different and uneven outcomes.

In addition to inspiration drawn from developmental psychology, our work is also grounded in previous autonomous robots research, particularly (Cañamero 1997, Avila-Garcia and Cañamero, 2004) regarding affective (hormonal) modulation of behavior selection, and especially (Blanchard and Cañamero, 2006, Cañamero et al., 2006) regarding modeling the caretaker in terms of perceptions that are also used to modulate the robot's affect and thus its behavior. Drawing on these ideas, we have developed a robotic architecture to explore a new environment and learn from it using the robot's caretaker as a "secure base" that provides it with "comfort" to reduce the robot's distress.

The main questions thus addressed in the present study are: (1) How can a caretaker help to shape the development of an autonomous robot in terms of cognitive, emotional and social abilities? And (2) To what extent are psychological theories about mother-infant attachment, and especially the role of the caregiver as a secure base, relevant to the design of an architecture for a developing robot?

One of the key roles that caregivers play in the cognitive-emotional development of infants is to contribute to the regulation of arousal, one of the main aspects of emotional control. Arousal regulation is crucial for the cognitive and emotional development of young infants (Sroufe, 1995, Brazelton and Nugent, 1995). In psychology, and also in this paper, arousal typically denotes a state of heightened physiological activity provoked by either endogenous or exogenous factors such as central nervous system (CNS) fluctuations or external stimulation. This aspect has not yet been properly investigated in developmental robotics; however, we believe that and endowing robots with an internal state akin to arousal that can be affected by external stimulation could contribute to its development by communicating this internal state to its human user or caregiver; this would in turn allow the human caregiver to choose to intervene or not in the robot's own experience.

In the remainder of this paper we report on an experiment illustrating how a caretaker can help to modulate the arousal of an infant-like robot by interacting with it and providing it with comfort. The architecture used in this study allows the robot to discover and learn information about its environment, and more specifically to habituate to the presence of certain patterns of stimuli and classify them in a stable manner. During this exploration, the arousal of the robot is stimulated by the novelty and the lack of variability of the patterns it senses. When this arousal level is high, the robot looks for comfort from the caretaker. Arousal thus modulates the behavior of the robot, and the caretaker modulates its arousal.

2. Robotics Model and Experimental Setup

2.1 Experimental Setup

In our experiments we have used an Aibo robot placed on play mat that also contains three cylindrical objects of different colors, as shown in Fig. 1. The robot uses three sensory modalities: color (the main color in the center of its visual field projected into the RGB color space), distance (the distance measurements provided by three distance sensors located in front of the robot), and contact (from one contact sensor on the top of its head and three on its back). Each sensor value (including the 3 RGB components of the color of the centre of its visual field) is discretized and projected into a vector containing ten binary elements. To summarize, the robot has to habituate to a vector aggregating all the elements of the sensory space, i.e., 100 binary elements (30 for the color, 30 for the distance sensors, 30 for the back sensors, and 10 for the head sensor). The caretaker can provide comfort to the robot either by appearing in its visual field and staying in sight, or by touching the sensors on its back. The robot recognizes the caretaker using the color of its clothes (this is hard-coded in this experiment, the caretaker is wearing a black top as it is the only color absent from the experiment room).

Figure 1: Our Experimental Setup

2.2 Robotics Architecture

Our architecture can be described in three main steps. The robot first learns the features encountered in its exploration of the environment, and by habituation and classification. Then the convergence and stability of these structures are evaluated to calculate the arousal level; this arousal level reflects the degree of *Surprise* and *Non-Mastery* of the robot in the current sensorimotor situation. Finally, an appropriate action is selected and executed.

Exploring and Classifying the Environment

To explore and categorize the environment, our architecture uses two different learning systems— . first, a Hopfield-like associative memory neural network to learn the patterns of stimuli encountered during the experiment, then a Kohonen Self-

Organizing map to classify the input pattern vector of sensor values encountered during exploration.

We decided to use a Hopfield-like associative memory system to allow our robot to habituate to the new patterns of stimuli it discovers. Another interesting property of such a system is the ability to recall learned patterns given a non-complete input vector. With this property, a new input pattern that is close to an already learned one, will be learned more quickly by the robot. Since this system learns incrementally and refines its weights to produce an output closer to the input, we can build a real-time measure of convergence and performance of the system. We can therefore measure how the system performs, e.g. the number of time steps the system needs to learn a new pattern, and the Euclidean distance between the output vector of the associative memory and the actual input vector. We will later use this measure to evaluate the arousal of the robot. The associative memory model is based on (Davey and Adams, 2004). The network proposed in this model is a two-dimensional grid of N neurons, with a state or output S_i, locally connected to their four nearest neighbors and randomly connected to four other units of the network with a symmetric connection matrix of weights w_{ij}. The connectivity is a blend of the two configurations represented in Fig. 2. This model is a modification of the standard Hopfield network. The local field h_i of a unit i is given by:

$$h_i = \sum_{i \neq j}^{N} w_{ij} S_j$$

then the next state of the unit i is calculated as:

$$S_i = \begin{cases} 1 & \text{if } h_i > 0 \\ -1 & \text{if } h_i < 0 \\ 0 & \text{if } h_i = 0 \end{cases}$$

In our network we use asynchronous random-order updates. To learn the presented input pattern vector, we use a modified version of the following procedure from (Davey and Adams, 2004):

Begin with a zero-weight matrix

Repeat either until all local fields are correct or for M time steps

Set the state of the network to one of the input patterns ξ

For each unit i in turn

Calculate $h_i \xi_i$

If this is less than a threshold T, then change the weights between unit i and all

other connected units j, according to:

$$\forall j \neq i \quad w'_{ij} = w_{ij} + \frac{\xi_i \xi_j}{N}$$

Here, ξ is the binary input pattern of the discretized value of the sensors. It is the input pattern vector to be learned by the robot. The point in which our algorithm differs from the original (Davey and Adams, 2004) is the repetition until all local fields are correct. In our experiment, the number of steps used to learn the current pattern is fixed (10 steps in the current settings). Therefore, the pattern is learned correctly and completely if the robot stays in its current position, in front of the exact sensory input pattern; if all the local fields are correct before ten time steps, then learning stops as described above. We use the Euclidean distance between ξ and the output of the system to compute the measure named *Surprise*, which is used to evaluate the arousal level. Therefore a small of distance means a

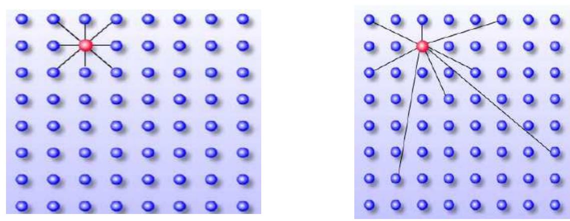

Figure 2: Associative memory network connectivity (locally connected on the left and randomly connected on the right, from (Calcraft et al., 2007))

The second learning algorithm we use is a classical Kohonen Self-Organizing map (Kohonen, 1997). The goal of this module is to classify the input pattern vector of sensors values encountered during exploration. Although we used the classical algorithm, we don't have a decreasing learning rate or neighborhood size over time; therefore, the map is constantly learning but has nevertheless a satisfying stability for already encountered patterns, and it also keeps its plasticity. We have chosen this particular learning system also because of its property of incremental convergence and because classification is an important cognitive ability that our robot could use, especially to discover and associate contingencies related to a certain sensorimotor context. We compute the measure of the stability of the map by summing the variations of every weights $w_i j$ of the map. This measure is used later to compute the arousal level of the robot.

Arousal Model

To compute the arousal of the robot we use two different contributions. First, we evaluate the discrepancy between the current pattern of stimuli and the

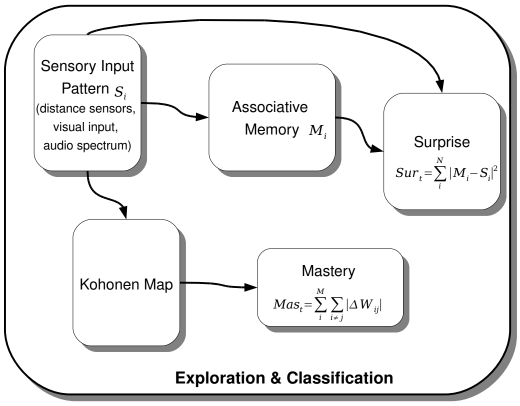

Figure 3: The robot explores and classifies the environment using a Hopfield-like associative memory and a Kohonen Map.

output of the associative memory, a value we call *Surprise* Sur_t, since it decreases as a function of the familiarity of the current pattern of stimuli. Indeed, since the associative memory has a fixed number of time steps to learn the pattern, more than one presentation is needed. When a pattern is familiar enough, the network converges fast and the *Surprise* value is close to zero. We also use Mas_t, a value we call *Non-Mastery*, which is the sum of the variations of the weights of the Kohonen map. This value shows the ability of the robot to classify the current pattern and how these classes evolve. If the values of the weights are stable, the current input pattern has already been correctly classified and the *Non-Mastery* is low.

The formulas of how these values are calculated are shown in Fig. 3. At each time step, the arousal of the robot is computed as:

$$A_t = \begin{cases} \frac{Sur_t + Mas_t}{2} \text{ if } T_{Care} = 0 \\ A(t-1) - \alpha \cdot T_{Care} \text{ otherwise} \end{cases}$$

where T_{Care} takes the value 0.5 when the caretaker is in sight, 0.8 when he/she touches the back sensors, 1 when both conditions are met, and 0 otherwise. Here α is the decay rate of the instantaneous arousal when the caretaker is interacting (set to 0.2). $A(t)$ is then used to evaluate a smoothed value of the arousal that we call *instantaneous arousal*, as follows:

$$A_{inst}(t) = \frac{\tau_a \cdot A_{inst}(t-1) + A(t)}{\tau_a + 1}$$

This value allows us to calculate an average of this arousal, called *sustained arousal*,

$$A_{sus}(t) = \begin{cases} \frac{\tau_{sus} \cdot A_{sus}(t-1) + A_{inst}(t)}{\tau_{sus}+1} \text{ if } T_{Care} = 0 \\ \text{and } \quad A_{inst}(t) > 0.4 \\ 0 \text{ otherwise} \end{cases}$$

where $\tau_a = 30$ is the time window on which the instantaneous arousal is calculated, as an average of $A_{inst}(t)$, and $\tau_{sus} = 10$, the time window on which the sustained arousal is calculated, as an average of the instantaneous arousal.

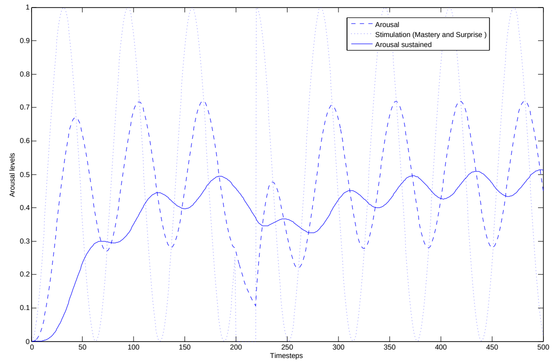

Figure 4: Dynamics of the Arousal System.

We show in Fig. 4 how the arousal level changes according to the degree of stimulation (*Surprise* and *Non-Mastery* values). This stimulation is here a sinusoidal function oscillating between 0 and 1. On time step 200, this stimulation is set down to 0 for 100 time steps. We can see how the arousal levels vary. Especially, the sustained arousal, which is used to demonstrate a high level of stimulation, decreases slowly and takes into account the overall state of the robot's arousal level. In a situation where the robot meets a high level of stimulation (new unlearned features), a decrease of the stimulation level has to be long enough for the robot to return to baseline, and another small stimulation will trigger the call for the caretaker's intervention.

Choice of Actions

The actions that the robot takes are based on the levels of both, instantaneous and sustained arousal.

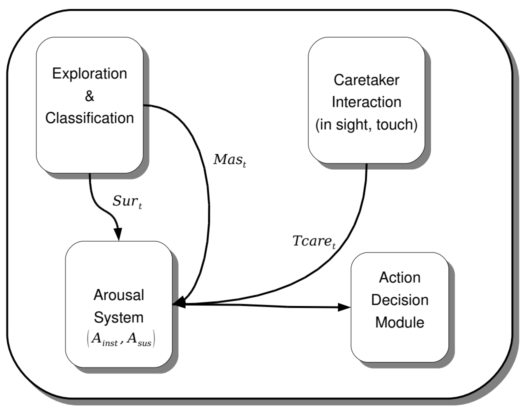

Figure 5: Entire Architecture

A_{inst}	A_{sus}	Actions
< 0.25	–	turn to explore
≥ 0.25 and ≤ 0.7	–	stay still and learn
> 0.7	–	bark to get attention
> 0.7	> 0.6	search for the caretaker

Table 1: Actions taken based on arousal levels

The robot can turn in only one direction, to discover a new pattern of stimuli when the arousal is low and the robot is in a "bored state". If the arousal is neither low nor high the robot remains still and tries to learn the current pattern of stimuli. If the arousal level is high, the robot barks to attract the caretaker's attention, and if the arousal is high and sustained, the robot looks for the caretaker by moving is head from top to bottom and left to right, trying to attract the caretaker in sight. Numerically speaking, the actions described above are taken when the conditions below are met:

3. Results

At every time step, we recorded the values described in the model section, namely *instantaneous arousal*, *sustained arousal*, caretaker interventions, associative memory error, and variations of the Kohonen map's weights.

We have represented the results of two typical experiments in Fig. 6 with two different caretaking styles: an active caretaker, responding almost constantly to the robot's demands (results on the right-hand side of the figure) and always staying on the right of the robot to appear in sight every time the robot is looking for him/her, and a caretaker who only interacts at the beginning and then leaves the robot on its own intervening only a few times (once every two minutes). Both experiments start in the same way: when the robot is put on the play mat, it is almost instantly asking for the caretaker, since all the features are new and highly stimulate its arousal (the sustained arousal level is high and oscillates above the 0.6 threshold). Then the caretaker appears in sight and touches its back sensors to calm it down. We can observe on the graphs that for both caretaking styles, the *Non-Mastery* and *Surprise* values are high and sustained in the case of the non-caring caretaker, since the "non-caring" caretaker then backs away immediately after putting the robot down. On the contrary, for the other type of caretaking, the experimenter stays close during the whole experiment. In the case of the non-intervening caretaker, the robot is surprised and quickly stimulated by the new environment, and the levels of arousal (sustained and instantaneous) urge it to look for the caretaker quickly. By doing this, the robot actually sees the colors of the upper environment, which are novel stimuli, and tries to learn them, and this results in an even higher increase of its arousal levels. As for the experiment with an active caretaker, since he interacts and provides comfort, the arousal levels are lower and the robot can explore without.

Style	$\bar{Mas}$	$\sigma(Mas)$	$\bar{Sur}$	$\sigma(Sur)$
Caring	0.5987	0.0355	0.3456	0.0565
Not Caring	0.6427	0.0407	0.6455	0.0324

Table 2: Results for 10 runs for each caretaking style.($p < 5\%$)

To find out how the two caretaking styles differ in terms of stability and performance of the exploration and classification system, we ran our experiment 10 times for each of the scenarios. The results for the average values and standard deviations for *Non-Mastery*, *Surprise* and *Sustained arousal* for the entire experiment are presented in Table 2. These values are used as a measurement of the quality of the learning process, to evaluate how each caretaking style affect the learning experience of the robot. Each run lasted 50,000 time steps and started from the exact same position. We can see that in terms of the Kohonen Map stability (the *Non-Mastery* value), the caring caretaker behavior does not outperform the non-caring one by a large difference. However, there is a large difference in terms of *Surprise* (the associative memory's performance) between the different caretaking styles. This means that by intervening during surprising episodes, the human agent has managed to decrease the time the robot spent in front of truly new situations, allowing the robot to skip for the time being. If these patterns were too far from the already learned ones, it is clear that the robot would need more time to learn them. Instead of having the robot stuck in front of them for a long time, human intervention allows the robot to try some other cases first, and return later to the difficult pattern to learn, offering a break in stimulation that might be too sustained, and therefore stressful. The sustained arousal gives coherent results since the robot without the caretaker has to deal on its own to reduce its arousal by mastering the situation and becoming habituated to the patterns. We can only conclude with this sample that the behavior of either caretaker is not necessarily the "optimal" one, and that finding the correct trade-off between staying close and not caring requires more investigation. As an end result, in all our runs the robot had learned and classified all the patterns encountered; therefore, its arousal always remained below the lower threshold and caused the robot to keep turning in the arena in a "bored state", looking for new features to learn.

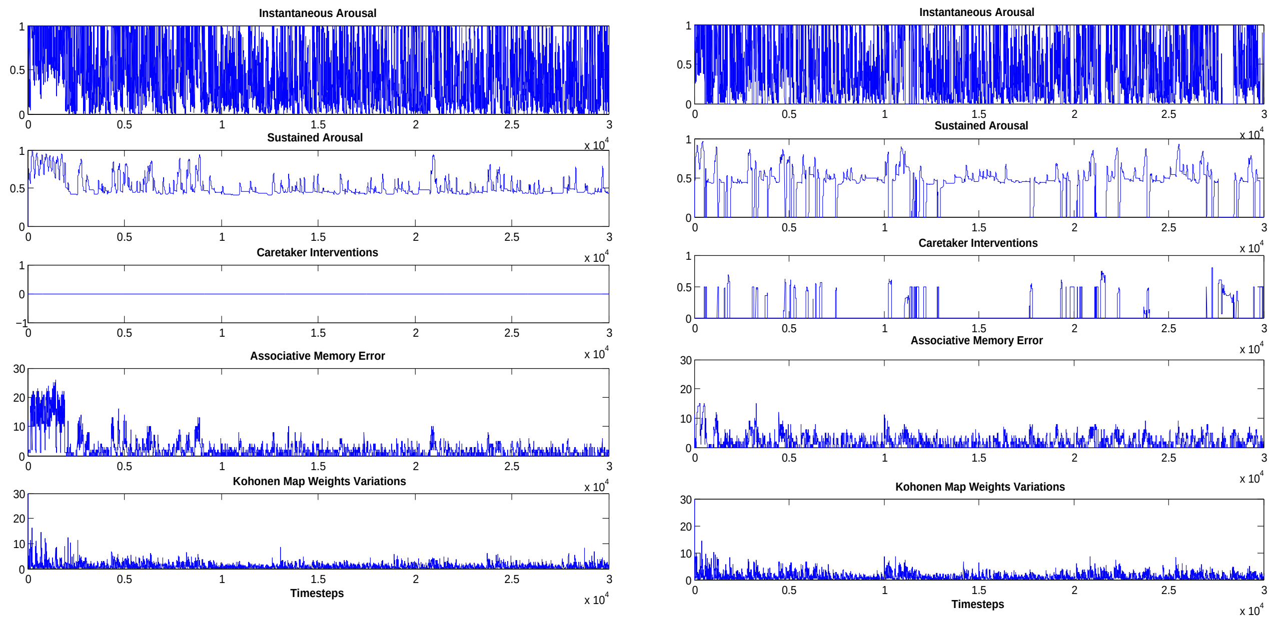

Figure 6: Evolution of (from top to bottom): instantaneous arousal, sustained arousal, caretaker interventions, associative memory error, and variations of the Kohonen map's weights. The graphs on the left-hand side correspond to an experiment with a caretaker only available at the beginning of the experiment whereas the ones on the right-hand side correspond to an active caretaker often providing comfort to the robot.

4. Discussion and Related Work

The architecture used in our experiment allows a robot to explore an unknown environment as a function of the dynamics of its interactions with the caretaker and the behavior of this latter. We have seen that even using such a simple architecture, the outcomes of every experiment are different depending on the type of interactions. The developmental approach we have followed reproduces mother-infant interactions. Our results show how using the caretaker as arousal—and indirectly as behavior—modulator is actually possible without having a complex architecture. Furthermore, apart from these two opposite caretaking styles, our architecture allows to actively choose whether a situation—a pattern of stimuli—has to be learned or avoided. Indeed, if the caretaker wants the robot to really learn the pattern, he/she can provide a small amount of comfort for the robot to have its instantaneous arousal in the middle level, between the two thresholds. This way the robot remains in its current position, without looking for the caretaker or moving away. In the opposite case, the caretaker can provide comfort to the robot so that it continues to look for another situation, keeping the instantaneous arousal below the lower threshold, and therefore preventing the robot from learning a situation that the caretaker considers irrelevant.

As for the related work, a comparable model of arousal modulation and mother-infant interaction, although, not applied to robotics, can be found in (Smith and Stevens, 1996, Smith and Stevens, 2002). In these contributions, the authors used a similar approach to modulate arousal based on neurophysiological data (Hofer and Sullivan, 2001) regarding how endogenous opioids modulate arousal in infants. However, their architecture did not have any cognitive system related to the interactions and their qualities, but was focused on the dynamics of the dyadic interaction. Also related to our study is (Likhachev and Arkin, 2000), in which the notion of comfort and object of attachment is used by a robot to remember its "comfort zones". One of the important aspects that differs in our work is that we use a person instead of an object, and also that the comfort of our robot is not a function of the distance between the robot and the object of attachment.

Finally, in (Thomaz and Breazeal, 2007), an interesting experiment is described showing how a human can help a robot learn a certain task. In this contribution, a robot can explore and learn on its own, but has also the opportunity to use human guidance to adapt to new tasks, changes in the environment, and to generalize one task to similar ones. The robot communicates its internal state with basic facial expressions and gestures. This "Socially Guided Exploration" presents some similarities with the work presented here; in both experiments the interactions with a human are used to enhance the learning process, and also in both cases the human teacher/caretaker has to pay attention to the feedback from the robot in order to intervene to help and guide the robot. However, what differs between the two experiments is the modalities the human uses to interact with the robot. In the experiment presented in this paper, the human caretaker orients the robot's behavior by touching its back sensor to reduce its arousal level in order for the robot to move to another sensorimotor context, or appear in sight, whereas in the contribution discussed here, the human teacher can either point with his/her finger to a certain region of the environment or even give verbal instructions to the robot. We argue that the simple non-verbal way of interacting we used in our experiment is sufficient to bias the behavior and improve the learning process of an autonomous robot.

5. Conclusion and Future Work

In the experiments described above, we have shown how it is possible to modulate the exploratory behavior of an autonomous robot using notions like *Surprise* and *Non-Mastery* to take into account its cognitive development, and especially using a caretaker as a secure base to provide comfort and reduce its arousal. We have seen that this architecture allows a robot to base its behavior on an internal state related to its own learning experience. Then the human agent—the robot's caretaker or user—can shape the development of the robot by either letting it cope with the current situation or, mediated by physical or visual contact, signal to the robot that it can discard the current situation for the time being. This non-verbal way of biasing the experience and exploration of the robot presents the advantages of needing less complex sensor processing compared to other interaction modes (e.g., speech processing and recognition), and looking and feeling more natural, especially if the robot is to be considered as an infant robot.

To provide a more autonomous and adaptive solution, we could integrate also material from previous work modeling the imprinting phenomenon, using a perception or a compound of them as "desired perceptions" (Blanchard and Cañamero, 2005, Hiolle et al., 2007). These perceptions could be the voice of the caretaker and his/her face. We could then add to our architecture the possibility for the robot to learn how to attract the attention of the caretaker and keep him/her close enough, as has been done in (Hiolle and Cañamero, 2007).

We would also like to investigate different behavioral profiles oscillating between exploring, learning, and demanding the caretaker's presence, and for this we need to explore several configurations of the parameters we have used such as the decay rates of arousal levels. We think that the use of even earlier experiences of the robot could help evaluate these parameters. Using this as grounding for an early shaping of the personality of the robot would help us build a more complex robot, and assess its attachment style using for example an Ainsworth-like Strange Situation Test (Ainsworth, 1969, Kaplan, 2001). To improve the autonomy of our robot's development, adding a curiosity drive (Oudeyer et al., 2007) would guide the robot's exploration towards more interesting situations, acting in order to increase its "learning progress". Another interesting possibility would be to modify our architecture using arousal to directly modulate the cognitive abilities of the robot, particularly learning.

Finally, on another level, accurate and consistent metrics to qualify and even quantify the behavior of the caretaker need to be studied to systematically measure how a caretaker interacts and assess the effects of the different caretaking styles. This would also open up a door to the investigation of how a robot could develop bonds with several caretakers and exhibit preferences for a given caretaker as a function of the given context or situation.

Acknowledgements

We are grateful to Neil Davey and Rod Adams for discussions about their neural networks model of associative memory. 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.

Avila-Garcia, O. and Cañamero, L. (2004). Using hormonal feedback to modulate action selection in a competitive scenario. In Schaal, S., Ijspeert, J., Billard, A., Vijayakumar, S., Hallam, J., and Meyer, J. A., (Eds.), *From Animals to Animats 8: Proceedings of the 8th Intl. Conf. on Simulation of Adaptive Behavior*, pages 243–252. Bradford Book.

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.

Brazelton, T. B. and Nugent, J. (1995). *Neonatal Behavioral Assessment Scale*. Cambridge University Press.

Cañamero, L.D. (1997). Modeling Motivations and Emotions as a Basis for Intelligent Behavior. In W.L. Johnson, ed., *Proc. First Intl. Conf. Autonomous Agents*, 148–155. ACM Press.

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

Calcraft, L., Adams, R., and Davey, N. (2007). Optimal connection strategies in one- and two-dimensional associative memory models. In *Proceedings of the 7th Intl. Wksp. on Information Processing in Cells and Tissues, Oxford*.

Davey, N. and Adams, R. (2004). High capacity associative memories and connection constraints. *Connection Science*, 16(1):47–65.

Hiolle, A. and Cañamero, L. (2007). Developing sensorimotor associations through attachment bonds. In Prince, C., Balkenius, C., Berthouze, L., Kozima, H., and Littman, M., (Eds.), *Proc. 7th Intl. Wksp. on Epigenetic Robotics*, pages 45–52. Lund University Cognitive Studies.

Hiolle, A., Cañamero, L., and Blanchard, A. (2007). Learning to interact with the caretaker: A developmental approach. In Paiva, A., Prada, R., and Picard, R., (Eds.), *Proc. of the 2nd Intl. Conf. on Affective Computing and Intelligent Interactions*, pages 422–433.

Hofer, M. (2006). Psychobiological Roots of Early Attachment. *Current Directions in Psychological Science*, 15(2):84–88.

Hofer, M. and Sullivan, R. (2001). Toward a neurobiology of attachment. *Handbook of developmental cognitive neuroscience*, pages 599–616.

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

Kohonen, T. (1997). *Self-Organizating Maps*. Springer-Verlag.

Likhachev, M. and Arkin, R. (2000). Robotic comfort zones. In *SPIE Sensor Fusion and Decentralized Control in Robotic Systems III*, pages 27–41.

Oudeyer, P.-Y., Kaplan, F., and Hafner, V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 2(11).

Schore, A. (2001). Effects of a secure attachment relationship on right brain development, affect regulation, and infant mental health. *Infant Mental Health Journal*, 22(1-2):7–66.

Smith, T. and Stevens, G. (1996). Emergence, Self-Organization, and Social Interaction: Arousal-Dependent Structure in Social Systems. *Sociological Theory*, 14(2):131–153.

Smith, T. and Stevens, G. (2002). Hyperstructures and the Biology of Interpersonal Dependence: Rethinking Reciprocity and Altruism. *Sociological Theory*, 20(1):106–130.

Sroufe, L. A. (1995). *Emotional Development: The Organization of Emotional Life in the Early Years*. Cambridge University Press.

Thomaz, A. and Breazeal, C. (2007). Robot learning via socially guided exploration. In *Proceedings of the 6th Intl. Conf. on Development and Learning, London*.