

From Imprinting to Adaptation: Building a History of Affective Interaction

Arnaud J. Blanchard Lola Cañamero

Adaptive System Research Group

School of Computer Science, University of Hertfordshire

College Lane, Hatfield, Herts AL10 9AB, UK

{A.J.Blanchard, L.Canamero}@herts.ac.uk

Abstract

We present a Perception-Action architecture and experiments to simulate imprinting—the establishment of strong attachment links with a “caregiver”—in a robot. Following recent theories, we do not consider imprinting as rigidly timed and irreversible, but as a more flexible phenomenon that allows for further adaptation as a result of reward-based learning through experience. Our architecture reconciles these two types of perceptual learning traditionally considered as different and even incompatible. After the initial imprinting, adaptation is achieved in the context of a history of “affective” interactions between the robot and a human, driven by “distress” and “comfort” responses in the robot.

1 Introduction

Imprinting—the phenomenon by which many animals (birds and mammals) form special attachments with objects to which they are exposed very early in life—is a very important learning mechanisms within the developmental process, in particular filial imprinting, in which the imprinting object is treated as a parent, giving rise to affiliative behaviors such as approaching and following. Developing such attachment with a caregiver provides many evolutionary advantages to the newborn in a moment of her/his life in which s/he cannot interact autonomously in the world, providing a basis not only to obtain needed resources and security, but also for social facilitation and learning, and for emotional development.

This phenomenon was for a long time considered to be instantaneous and irreversible, as the term “imprinting” suggests. In the mid 1930’s, ethologist Konrad Lorenz made this phenomenon well-known through his studies of greylag geese. Lorenz raised these animals from hatching, becoming the imprinting—parent-like—object for them. This “unnatural” imprinting to an individual of a very different species initially suggested that the animals had become attached to the first “eye-catching” object they had perceived immediately after hatching. This form

of perceptual learning was also considered to be very different from (and unrelated to) other types of learning arising later in life, such as conditioning or associative learning. Such view has been more recently questioned as over-simplistic. Bateson, for example, postulates a model (see e.g., (Bateson 2000)) in which imprinting is not an instantaneous and irreversible process but a much more flexible and less peculiar phenomenon. The main points of this view can be summarized as:

- Imprinting does not necessarily occur immediately after birth but has a more flexible sensitive period (Bateson & Martin 2000) affected by both experience and species-specific features. This provides some flexibility regarding the exact point in time in which the mother is first “perceived” and imprinted.
- Imprinting is not a monolithic capability but is composed of several linked processes (Bateson 2000): (1) “analysis” or detection of a “relevant” stimulus guided by predispositions of what the animal will find attractive; (2) recognition of what is familiar and what is novel in that stimulus, which involves a comparison between what has already been experienced and the current input; and (3) control of the motor patterns involved in imprinting behavior.
- Although imprinting can be functionally distinguished from learning involving external reward, both types of learning are deeply connected, as suggested by the possibility of transfer of training after imprinting.

In this paper we present a novel Perception-Action architecture and experiments to simulate imprinting in a robot following this latter approach. Starting with a basic architecture that simulates imprinting in the more traditional sense (Section 2), we incrementally modify and extend this architecture to achieve further adaptation, also integrating reward-based learning (Section 3). This adaptation is achieved in the context of a history of “affective” interactions between the robot and a human (Section 4), driven by “distress” and “comfort” responses in the robot.

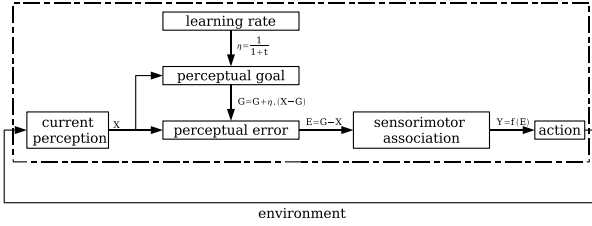


Figure 1: Architecture used to model imprinting.

2 Establishing Attachment Bonds

2.1 Robotic architecture for imprinting

The architecture we have used to implement imprinting follows a “Perception-Action” approach rooted both in psychology (Prinz, 1997) and in robotics (Gaussier *et al.* 1998), and that we have already successfully applied to a movement synchronization task in robots (Blanchard & Cañamero 2005). This approach 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 it is our case here) or learned from experience by the robot (see e.g., (Andry *et al.* 2002) for an example applied to a robot imitation of arm movements).

As depicted in Figure 1, we have used this general approach to model imprinting as an attempt to reduce the difference—i.e., correct the perceptual error, noted E —between the *current perception* (X) and a *goal perception* (G) that, under normal circumstances, would be a beneficial perception related to the caregiver. The speed at which learning takes place depends on time (as reflected by the learning rate η). The choice of actions (Y) to correct perceptual errors is based on sensorimotor associations ($f(E)$) stored in a look-up table.

2.1.1 Learning the goal perception

Intuitively, the most obvious way to implement imprinting in a robot would be to have it learn the first perception that it has when it is switched on (the equivalent of “hatching” in birds) as being his “goal” perception—the perception it will memorize and try to maintain after imprinting. This could be implemented:

$$G(t) = X(t_0) \quad (1)$$

where G is the goal perception (the “goal” values for all sensor readings), t is the time elapsed from “hatching”, X is the current perception (the current values for all

sensors), and t_0 denotes time at “hatching”. This corresponds to the view of imprinting as “stamping” or developing instantaneous and irreversible affiliative bonds with the first “eye-catching” stimulus perceived. This approach presents the advantage that it guarantees that the goal perception of the robot will be reachable, since it corresponds to a perception that has already been reached once. On the contrary, using some sort of “predisposition” to decide which (features of the) stimulus among those perceived at “birth” will become the imprinting object (i.e., the goal stimulus) does not guarantee that a goal stimulus will be found. This could for example happen if no suitable stimulus is present at “hatching” time—e.g., the stimulus is just noise or no stimulus is detected. In this case, using the approach represented by Equation 1, the robot would not be able to start acting, since it would not have acquired a “goal” at birth (it would not be imprinted to anything) and it would not be able to acquire it after the imprinting “time window” had closed.

To solve this problem, instead of memorizing exactly the first perception, the robot could memorize the “average perception” from the beginning of its life, incorporating the history of its interactions with the environment in its perceptual memory. At the beginning, when the robot has few experiences, the average perception will be almost equal to the current perception and this latter will have strong impact on its behavior. However, with time, experiences will accumulate in its memory and the influence of the current perception in guiding its behavior will decrease. Storing all the past perceptions to compute the average perception would be too costly and unrealistic. A more biologically plausible strategy would rather take into account the last “goal perception”, which would incorporate the history of past experiences, and make learning dependent on time by using a decreasing learning rate ($\eta(t) = \frac{1}{1+t}$) to achieve stabilization:

$$G(t) = G(t-1) + \eta(t) \times (X(t) - G(t-1)) \quad (2)$$

Note that this corresponds to the stochastic LMS (Least Means Square) learning rule commonly used in neural networks. This means that using average perception is thus equivalent to learning with a decreasing learning rate, as shown in Figure 2. The learning rate at “hatch-

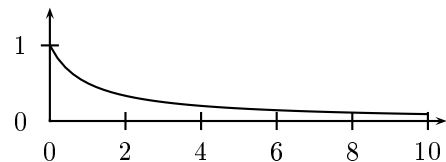


Figure 2: Decrement of the learning rate (y-axis) as a function of time (x-axis). At the beginning, a learning rate of 1 means that the goal perception is equal to the current perception.

ing” or imprinting time is 1; therefore, learning is instantaneous at that moment and the goal perception is equal to the current perception. However, we can easily vary

the size of the “time window” during which imprinting occurs by altering the sharpness of the decreasing rate, since the time unit t is arbitrary.

2.2 Experiments

2.2.1 Apparatus

We have implemented and tested our architecture using a Koala robot (<http://k-team.com/robots/koala>). Only the ring of infrared proximity sensors located around the robot was used to provide perceptual input in these experiments. The average of all the infrared front sensors was used to detect (the proximity of) objects at the front of the robot—we will refer to this averaged reading as “the proximity sensor”. Distance to the perceived stimulus is the only perceptual feature used to form the “goal perception”—i.e., for imprinting. The infrared sensors at the back are used to avoid collision when the robot moves backwards. The only actions of the robot after “hatching” (and therefore imprinting) are forward and backward movements as side-effects of its attempts to achieve the goal perception acquired at imprinting time. As a consequence of this, the robot “approaches”, “follows” or “avoids” (reverses if approached at a distance smaller than the learned distance) the imprinted object as this moves around. We used two types of imprinting stimuli: near objects (high activity of the proximity sensors) and distant objects (lower activity of the proximity sensors). Two types of objects—a human and a cardboard box moved by a human, as shown in Figure 3—were used as near and distant stimuli. Although the experiments worked very satisfactorily with both types of objects, only the results obtained with the cardboard box (10 tests for each condition) were used for analysis purposes due to their higher clarity.



Figure 3: Experimental setting. In this case, a box located close to the robot is used as imprinting object.

2.2.2 Results and discussion

Using the box, 10 tests were run for each “hatching” condition—near or distant imprinting object. Figure 4

shows one representative example of each condition, with graphs on the left side of the figure (a1 and b1) corresponding to the “near hatching” case, those on the right (a2 and b2) to the “distant” one. Top graphs (a1 and a2) show current (solid line) and goal (dashed line) perceptions, bottom graphs (b1 and b2) show the speed of the robot responding to the random movement of the box. In

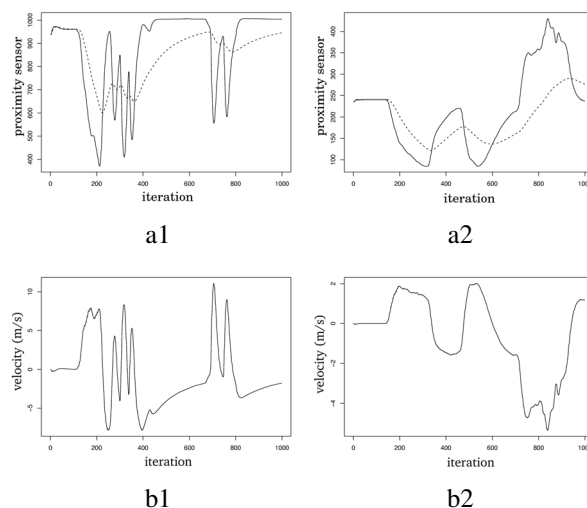


Figure 4: Results of two experiments testing imprinting to near (left graphs) and to distant (right graphs) stimuli. The y-axis shows averaged readings of the proximity sensors on the top graphs and the speed at which the robot moves to correct the perceptual error on the bottom graphs, the x-axis shows time from “hatching” in all graphs.

both conditions, the goal perception (dashed lines in a1 and a2) fluctuates at the beginning, since it is closer to the current perception, but the goal perception becomes more stable with time in both cases, even though the imprinting stimulus moves at different distances at the front of the robot. As a consequence of homeostatic control, the velocity of the robot (graphs b1 and b2) changes in order to decrease the difference between goal and current perceptions. Motor speed is directly proportional to (a fraction of) the magnitude of the perceptual error.

We thus see how the robot learns the imprinting stimulus using a very simple function. Such learning can take place even when the imprinting object (i.e., an object detected at a particular distance within the range of the infrared sensors) is absent at “hatching” time, although learning becomes more slow and difficult with time, corresponding to the limited time window during which the imprinting process is possible in animals. However, this model still implements the simple view of imprinting as “stamping” a permanent and irremovable trace, while, as Bateson points out, “the process is not so rigidly timed and may indeed be undone” (Bateson & Martin 2000). It also disregards the connection between imprinting and other types of (reward-based) learning. For an autonomous robot living in a changing and social environ-

ment, being able to modify or undo what was learned during imprinting is also very important, since (a) it is virtually impossible for the designer to define a priori a time window for the imprinting process that works in all possible environmental conditions, and (b) if the environments (including the social partner) changes, the robot has to adapt to the new features.

3 From Imprinting to Adaptation

In algorithms employed in autonomous robots and neural networks research, it is very common to use a learning rate that decreases with time in order to achieve a good level of stability in memory that consolidates learning. The learning rate must vary with time since, if it were constant, everything that is learned would be replaced by new events, memory contents would change constantly. It is common to use a decreasing learning rate of the type $\eta(t) = e^{-k.t}$, where k is a constant that changes the size of the temporal window. However, learning should change not only as a function of time but also of the *relevance* of the stimulus. The problem now is thus how to make the robot assess what is relevant.

3.1 Assessing relevance

To assess the relevance of external stimuli, we use the notion of “well-being” or *comfort*: since under normal circumstances, the evolutionary advantage of becoming attached to a caretaker is to foster security, beneficial interactions with the environment, and generally well-being, stimuli that carry some comfort associated with them are thus those stimuli relevant to become attached to. Drawing on Ashby’s view of survival as viability (Ashby 1952) or stability of the internal environment, in our robot comfort is related to the stability of its internal homeostatic variables, following (Cañamero 1997). Closely related architectures have used a similar notion of “comfort” (also termed “well-being” or “satisfaction” in those architectures) and “discomfort” to assess and compare the performance of different behavior selection policies in autonomous robots (Avila-García & Cañamero 2004), and to learn affordances through the interactions of a robot with objects in the environment (Cos-Aguilera *et al.* 2003). There are different ways to calculate comfort when the internal environment consists of several internal homeostatic variables, such as the inverse of the average of the errors (deviations between the actual value and the “ideal value” or “setpoint” of the variable) of all the variables, the variance, etc—see e.g., (Avila-García & Cañamero 2002) for a presentation and discussion of different metrics. A simple way of calculating comfort c at each point in time given n variables, by taking the average of their errors (e_1, \dots, e_n), could be:

$$c = \frac{1}{1 + e_1 + \dots + e_n} \quad (3)$$

In our case, comfort can take values between 1 (maximum comfort) and 0 (minimum comfort). As we will see later (Section 4.4.1), we use tactile contact as a source of comfort. We will thus try to make our robot learn to recognize the stimulus that gives it most comfort. To do that, we modulate the learning rate with the comfort:

$$\eta(t) = c.e^{-k.t} \quad (4)$$

This “perceptual learning” will decrease with time in the same way as learning in the imprinting algorithm, but now it will also depend on the relevance of the stimulus as measured by the comfort it provides to the robot. The more comfort a stimulus provides, the faster the robot will develop an attachment link to it and the stronger this link will be—i.e., the more relevant the stimulus will be for imprinting. However, this modulation of the learning rate does not show the interesting property of instantaneous learning at “hatching” time, even when $e^{-k.t} = 1$. To achieve this, we have to use again the average perception (instead of the current perception), ponderating this time the stimulus X with the level of comfort c . Equation 2 (which calculates the average perception by taking into account the past goal perception) remains the same but the learning rate is now:

$$\eta = \frac{c(t)}{\tilde{c}(t)} \quad (5)$$

where $\tilde{c}(t)$ is the sum of the comfort in all time steps since “hatching”. We can now reproduce the imprinting phenomenon described in Section 2, this time taking into account the relevance that the observed stimulus has for imprinting, since the comfort produced by some stimuli (e.g., a caretaker stroking the robot) amplifies the effects of these relevant stimuli over non-relevant ones (e.g., a static wall). As we will see, in addition to the homogeneity and simplicity of the equations, using this function presents some other advantages for learning.

3.2 Multiple “goal perceptions”

With the function described in Section 3.1, after some time interacting with the environment learning becomes very slow, as \tilde{c} becomes very large. Intuitively, this would correspond to a situation in which the robot has formed an attachment bond with the “caretaker” but cannot learn anything else. However, the robot, like animals, should be able to learn new things while interacting with its environment and which of them are “beneficial” or “noxious” for it while remembering what was learned during imprinting. We are thus facing the problem of how imprinting relates to later forms of learning. A possibility would be to consider further learning as a completely different process that starts once imprinting has finished and for which we could use a learning rate that depends on the comfort (e.g., as in (Cos-Aguilera *et al.* 2003)) but not on the time from “hatching”. However, this would erase useful memories. To make these different types of learning

compatible, we can consider them as related processes to learn what is relevant (beneficial/noxious) for the individual at different time scales. For example, learning to recognize the caretaker serves a goal that is beneficial in the long term, whereas learning about the usefulness of an object to satisfy an urgent need serves an immediate goal. Instead of learning a single “goal perception” that the robot will try to achieve or maintain through its interactions with the environment, it could thus learn different perceptions that will be considered as “goal perception” at different moments depending on the time scale used to remember (seconds, hours, days, etc). We will call these perceptions *desired perceptions*, and we will see how they are selected later on (Section 4.2).

4 Adaptation via Affective Interaction

As Bateson points out (Bateson 2000), imprinting should not be regarded as an irreversible process that was completed once and for all when the appropriate “time window” closes to the world. Even if learning about the features of the imprinting object becomes more difficult after the “sensitive period” (Bateson & Martin 2000), the effects of imprinting are not irreversible. Increased learning difficulty presents a mechanisms to protect the learned object “representation” from change after imprinting. However, leaving the possibility of further learning open also presents evolutionary advantages. Think of an animal or a robot initially imprinted to a very devoted and “close” caretaker; if the caretaker is replaced by a “colder” and more “distant” one with very different interaction patterns, our “infant” will be much better off if it is able to adapt to this new circumstances by learning from its experience, since otherwise it would keep “making mistakes” in its interactions with the new caregiver and would feel permanently miserable. From the perspective of learning, this implies trying to reconcile imprinting and reward-based learning, and this presents problems such as conflicting requirements regarding the learning rates needed for each process. Our approach thus differs from reinforcement learning algorithms such as Q-learning and TD-learning since it deals with several learning rates and makes a selective use of memory—only the “best” perception related to each time scale is kept. To provide a common framework for imprinting and reward-based (in our case comfort-based) learning, we have to reconcile the following ideas:

- At the beginning (i.e., during the imprinting process) we want the learning rate to decrease with time to consolidate memory and “protect” what was learned about the caregiver.
- It is useful to continue learning new things. Since we don’t know in advance which is the best learning rate for each particular case, it might be useful to remember “goal perceptions” at different time scales. However, this process cannot work at the beginning (during the imprinting process) since the robot has

not accumulated enough experiences. Also the learning rate needed (closer to a constant rate) seems in conflict with the decreasing learning rate above.

To take advantage of the benefits of both cases, we can modulate the learning rate by rising it to different powers depending on the time scale of the modulation. The time scale k is defined by γ_k , which can take values between 0 and ∞ :

$$G(t) = G(t - 1) + \left(\frac{c(t)}{\bar{c}(t)}\right)^{\gamma_k} \cdot (X(t) - G(t - 1)) \quad (6)$$

If γ_k tends to ∞ , the learning rate tends to 0—after “hatching”, there is no further adaptation of what has been learned about the imprinting object. If γ_k tends to 0, the learning rate tends to 1—there is no stability and the desired perception tends to correspond to the current perception. Between these two extremes, we have different intermediate learning modes available. Examples are provided in Figure 5, which shows the evolution of the learning rate (modulated by the comfort, which is kept constant) under three different time scales.

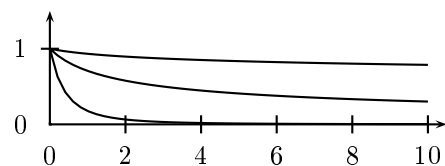


Figure 5: Evolution of the learning rate on three different time scales. The y-axis shows learning rate values, the x-axis time from “hatching”. Parameter values defining the time scale of the learning rates are $\gamma_0 = 0.1$ for the top curve, $\gamma_1 = 0.5$ for the middle curve, and $\gamma_2 = 2.5$ for the bottom curve.

4.1 The effects of comfort

Making the learning rate depend on the comfort can present disadvantages depending on how the comfort modulates this rate. It is very difficult to know in advance what the average comfort of the robot will be. If we use different fixed learning rates modulated by the level of comfort, adaptation could be very low if the environment is “difficult” or “hostile” (producing very low levels of comfort), but learning could also be unstable if the environment is highly “positive” (i.e., providing very high levels of comfort). The strong influence of the comfort level can thus be problematic because this level would have to be chosen depending on the hostility of the environment, and neither the robot nor the designer have this information in advance.

The method that we propose and use here does not present this problem, since the learning rate is not modulated by the level (absolute value) of comfort but by its *variation*. Comfort is therefore not regarded as having an “ideal value” that the robot should try to maintain or achieve, but as a relative notion that changes under different circumstances. The “background goal” of the robot

will still be to maintain an acceptable level of comfort, but what “acceptable” means can change, i.e., the “set-point” or the “threshold” setting that goal is variable. This allows the robot to learn to adapt its perceptual goal (or to learn different perceptual goals) depending on what is considered as “acceptable” comfort at that moment. This also means that, as a result of this learning, the robot will adapt its interactions to the interaction styles of different caretakers. This adaptation is not something that only takes place “then and there”, but it also depends on the history of the interaction. The use of the sum (“past history”) of the comfort \tilde{c} in the denominator of the learning rate allows to modulate the effect that the current comfort has on it as a function of past experiences.

Our robot is now able to memorize different “desired perceptions” related to different time scales. Let us see how to select among them the “goal perception” that it will actually try to reach.

4.2 Selecting the time scale

The choice of the time scale (and therefore of the desired perception that will be sought as “goal perception”) can be directly driven by the comfort. The goal perception will be mainly associated with a desired perception in a short time scale when the comfort is high, and it will be associated with a desired perception in a long time scale—of which the imprinting perception is an example—when the comfort is very low. The use of a short time scale allows the robot to be very reactive to external changes, which in principle is advantageous for its survival, but on the other hand the lack of experience puts it in an “insecure” position that should be avoided when the comfort is already low (Avila-García & Cañamero 2004). Intuitively, when the robot feels comfortable, it will have a more open stance towards the external world and will tend to “live in the present”. On the contrary, in a situation of discomfort it will be more closed to the world and the present situation, to look back for past memories.

The final goal perception will be a combination of the different desired perceptions weighted by a “filter”—see Figure 6. The position of the maximum value in that filter depends on the comfort.

4.3 Explore or exploit?

Our architecture now allows the robot to continue learning after the initial imprinting, and it can have at its disposal a rich repertoire of past “desired perceptions” that can be used to search for new “goal perceptions”. However, if the robot is continually trying to achieve its “best” perception looking into its multiple time-scales memory, it will avoid any new perception and therefore will not be able to learn from new experiences. This can be seen as an instance of the well known “exploitation/exploration” dilemma (Wilson 1996) in autonomous learning, i.e., how to decide between using the knowledge

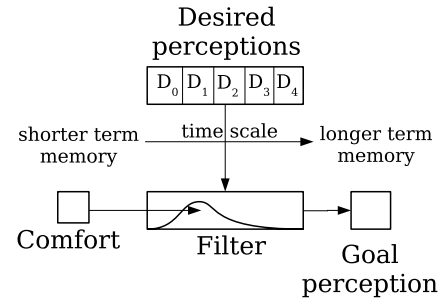


Figure 6: The goal perception is formed from desired perceptions at different time scales by means of a filter that weights the contributions of these desired perceptions. In this filter, the maximum value is defined by the value of the comfort.

already acquired in order to solve a problem, or continuing exploring to acquire new knowledge. We thus need a mechanism to solve this problem.

Comfort can also be used to provide such mechanism, since there is evidence that a good level of comfort (e.g. postural comfort (Kugiumutzakis *et al.* 2005)) facilitates learning in infants. This also makes sense in our architecture. When the robot has a low level of comfort, it will look for a “better” perception. If it is unable to reach it, its comfort will continue to decrease and it will keep hopelessly trying indefinitely to reach it. With time, the situation will become so bad that the caretaker will not even be able to approach the robot to provide it comfort, since the “best” perception that the robot has at the moment is one of a distant imprinting object and it will reverse when the caretaker approaches it too much, to try to keep that “best” perception. Conversely, in a situation of high comfort the robot will have no reason to change its current perception. A good strategy seems thus to let the current perception change (i.e., to “pay attention” to new perceptions) when the robot has a good level of comfort; this is achieved by inhibiting (modulation by “activity” in Figure 7) its attempts to attain its desired perception. On the contrary, when the comfort is low, the robot will try to actively reach memorized good perceptions. Openness to the world, activity, and learning would thus have an inverted-U shape as a function of comfort.

Figure 7 summarizes our global Perception-Action architecture described in this section, combining imprinting and reward-based (comfort-based) adaptation.

4.4 Experiments

4.4.1 Apparatus

The setting of these experiments is very similar to the one presented in Section 2.2, but this time we need to add some new features to manage comfort. The robot receives comfort as a result of tactile contact on its leftmost infrared proximity sensor. We also added to the architecture an internal homeostatic variable, “tactile contact” that the robot must keep close to an “ideal value” and that

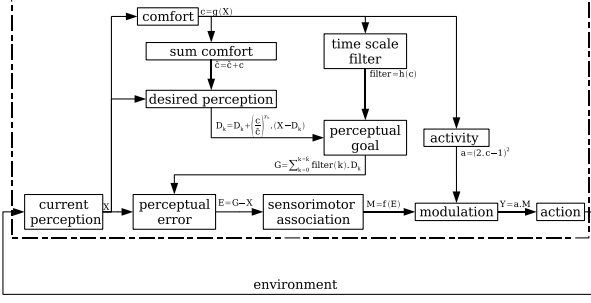


Figure 7: Global Perception-Action architecture for imprinting and adaptation.

decays with time in the absence of contact on the infrared sensor mentioned above. We adapt Equation 3 to calculate the robot’s comfort using this variable:

$$c = \frac{1}{1 + e^{-contact}} \quad (7)$$

The robot will try to keep c as high as possible given its present circumstances and the history of its interactions. To facilitate interaction with humans, the robot emits beeps with a frequency that depends on the level of “distress”, i.e., the frequency of the beeps increases as the comfort decreases. This is akin to a “separation distress” response in animals, and is intended to “flag” the need for action on the part of the human—tactile contact that will produce a “comfort response” in the robot (see e.g., (Panksepp 1998) for a discussion of the separation distress and comfort responses in animals).

4.4.2 Results

Figure 8 shows the results of one example of interaction with the robot. We began the interaction (the moment of “hatching”) without any object at the front of the robot. The robot therefore starts with a “noisy” imprinting situation in which there is no imprinting object—point ‘a’ at the top of Figure 8. Therefore, when we try to approach it (point ‘b0’) it moves backwards (marked as ‘b1’ in the bottom graph of the figure). We then give it some comfort (point ‘c1’ in the middle graph) by touching its side sensor and we observe (also in the middle graph) that the activity level or “arousal” decreases. When we approach the robot again (point ‘c0’) it does not reverse (“avoid us”) anymore, as we can observe in the “plateau” in the lower graph. We then remain close to the robot for some time, touching its sensor simultaneously in order to make it learn that in fact, and contrary to its initial experience, it is beneficial to have a “stimulus” in front of it. When this “stimulus” disappears, we also stop touching its side sensor; the comfort then starts to decrease while the activity level or ‘arousal’ increases (d1), and the robot will give a high weight to a long time scale (d0) and therefore it will try to reach a long-term desired perception (e0): it will move forwards (d2) to try find something at

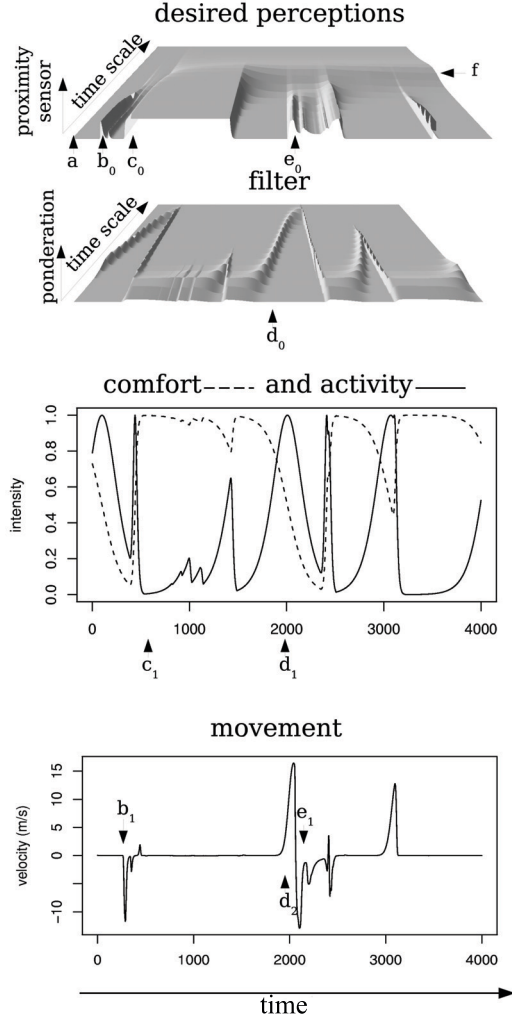


Figure 8: Evolution of the different internal states and movements of the robot during an interaction of about 2 minutes. See text for explanation.

its front. When it finds it, it stops (e1). It is interesting to note a very stable shape (denoted by ‘f’) on a rather long time scale of the desired perceptions graph. This means that, globally, the presence of something at the front of the robot is positive even if locally (on a short time scale) it is not always the case. In fact, continuing this experiment (approaching an “object” to the robot and giving it comfort) for a longer period, we would assist to a slow propagation of that stable shape (f) to the very long-term scales, eventually modifying the memory of the imprinting stimulus.

5 Conclusion and Perspectives

We have presented a Perception-Action architecture and experiments to simulate imprinting in a robot. Following recent theories about imprinting in animals, we do not consider imprinting as rigidly timed and irreversible but as a more flexible phenomenon that allows for further adaptation as a result of experience. Our architec-

ture reconciles two types of perceptual learning traditionally considered as different, and even incompatible, due to apparently conflicting features and functions: the establishment of an initial attachment to a “caregiver” (an imprinting object) and reward-based learning as a result of experience, that we have grounded in the notion of internal comfort. Adaptation is achieved in the context of a history of “affective” interactions between the robot and a human, driven by “distress” and “comfort” responses in the robot.

Our implementation made some simplifications that we would like to improve in the future to achieve richer human-robot interactions. First, we only used one feature (distance to the perceived stimulus) to learn about the “caretaker”. Proper treatment of learning about the imprinting object would require considering multiple features that the robot would have to analyze in order to recognize the “caretaker” from different perspectives and in different situations. Second, at present the robot only stores a desired perception per time scale in its memory. However, taking into account other contextual factors would necessitate learning and handling different desired perceptions within each time scale. Third, including more potential sources of comfort (e.g., adding more internal needs such as “feeding”, “keeping warm”, etc.) would create richer social interactions between the robot and the “caretaker”. Finally, desired perceptions provide the robot with a mechanism to “decide” what it should reach, but further development would also require a mechanism to “decide” what it should avoid (“avoided perceptions”), something like the basis of “fear” system.

Acknowledgments

Arnaud Blanchard is funded by a research scholarship of the University of Hertfordshire. This research is partly supported by the EU Network of Excellence HUMAINE (FP6-IST-2002-507422).

References

- Andry, P., Gaussier, P., and Nadel, J. (2002). From sensorimotor development to low-level imitation. In C.G. Prince, Y. Demiris, Y. Marom, H. Kozima, and C. Balkenius (Eds.), *Proceedings 2nd Intl. Wksp. on Epigenetic Robotics. Lund University Cognitive Studies*, 94. Lund: LUCS.
- Ashby, W.R. (1952). *Design for a Brain: The Origin of Adaptive Behaviour*. London: Chapman & Hall.
- Avila-García, O. and Cañamero, L. (2002). A Comparison of Behavior Selection Architectures Using Viability Indicators. In R. Damber (Ed.), *Proc. of the EPSRC/BBSRC International Workshop Biologically-Inspired Robotics: The Legacy of W. Grey Walter*, 86–93. August 14–16, 2002, HP Labs Bristol, UK.
- Avila-García, O. and Cañamero, L. (2004). Using Hormonal Feedback to Modulate Action Selection in a Competitive Scenario. In S. Schaal, A.J. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam and J.-A. Meyer (Eds.), *From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior (SAB'04)*, 243–252. Cambridge, MA: The MIT Press.
- Bateson, P. (2000). What must be known in order to understand imprinting? In C. Heyes and L. Huber (Eds.), *The Evolution of Cognition*, 85–102. Cambridge, MA: The MIT Press.
- Bateson, P. and Martin, P. (2000). Sensitive Periods. In P. Bateson (Ed.), *Design for a Life : How Behavior and Personality Develop*, NY: Simon & Schuster.
- Blanchard, A. and Cañamero, L. (2005). Using visual velocity detection to achieve synchronization in imitation. In Y. Demiris (Ed.), *Proc. of the AISB'05 Third International Symposium on Imitation in Animals and Artifacts*, University of Hertfordshire, UK, April 12–15, 2005. SSAISB 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. New York: ACM Press.
- Cos-Aguilera, I., Cañamero, L., and Hayes, G. (2003). Learning Object Functionalities in the Context of Action Selection. In U. Nehmzow and C. Melhuish (Eds.), *Towards Intelligent Mobile Robots, TIMR'03: 4th British Conference on Mobile Robotics*. University of the West of England, Bristol, UK, August 28–29, 2003.
- Gaussier, P., Moga, S., Banquet, J., and Quoy, M. (1998). From perception-action loops to imitation processes. *Applied Artificial Intelligence*, 1(7).
- Kugiumutzakis, G., Kokkinaki, T., Makrodimitraki, M., and Vitalaki, M. (2005). Emotions in Early Mimesis. In J. Nadel and D. Muir (Eds.), *Emotional Development*. New York, NY: Oxford University Press.
- Panksepp, J. (1998). *Affective Neuroscience: The Foundations of Human and Animal Emotions*. New York, NY: Oxford University Press.
- Prinz, W. (1997). Perception and action planning. *European journal of cognitive psychology*, 9(2).
- Wilson, S.W. (1996). Explore/exploit strategies in autonomy. In P. Maes, M. Mataric, J. Pollack, J.-A. Meyer, and S. Wilson (Eds.) *From Animals to Animats: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, 325–332. Cambridge, MA: The MIT Press.