Available online at www.sciencedirect.com
SCIENCE DIRECT®
Robotics and Autonomous Systems xxx (2004) xxx–xxx
Robotics and Autonomous Systems
www.elsevier.com/locate/robot

From behaviour-based robots to motivation-based robots

Riccardo Manzotti*, Vincenzo Tagliasco

Laboratory for Integrated Advanced Robotics (LIRA-Lab), Department of Communication, Computer and System Sciences, University of Genoa, Viale Causa 13, Genoa I-16145, Italy

Received 28 December 2002; received in revised form 12 October 2004; accepted 20 October 2004

Abstract

The appearance on the market of entertainment robots for children and families has ipso facto created the new category of motivation-based robots. A taxonomy of the architectures of different robot categories is proposed. The architecture of motivation-based robots is phylogenetic and ontogenetic. A tentative architecture for a specific experimental setup is described. The results of the experiment show that a new motivation arises from the interaction between the robot and the environment. Motivation-based robots equipped with ontogenetic architecture might provide the foundation for a new generation of robots capable of ontogenetic development.

Keywords: Associative learning; Development; Entertainment robots; Behaviour-based robots; Motivation-based robots

1. Introduction

Manufacturers design entertainment robots capable of interacting with humans. This interaction occurs at several levels: from the selection of a set of in-built behaviours to the capability of being independent and acting on its own. Entertainment robots learn actions through touch sensors, switches and voice recognition modules. The most sophisticated robots are said to develop unique personalities through the interaction with a specific environment. Robots that go through a series of development phases (real or simulated from toddler, to child, to adult) appeal to consumers. Moreover, robots must show emotions like happiness, sadness, anger and surprise, in different degrees. Entertainment robots must be curious and must be able to explore their surroundings on their own: these robots develop in relation to their personal history. We define this class of robots as motivation-based robots because they aim at re-creating the motivational structure of biological beings. The time has now come to move from behaviour-based robots [1] to motivation-based robots [2–4].

Recently, in neuroscience and robotics, the problem of what motivation is has been investigated in the more general framework of what a subject is [5,6]. We

* Corresponding author. Tel.: +39 010 353 2817/3478257675; fax: +39 010 353 2948.
E-mail addresses: manzotti@dist.unige.it (R. Manzotti), vincenzo@dist.unige.it (V. Tagliasco).

0921-8890/$ – see front matter © 2004 Published by Elsevier B.V.
doi:10.1016/j.robot.2004.10.004

propose an engineering approach to create motivations in robots that does not require such a broad framework. Behaviour-based robots make use of fixed "motivations" hardwired in their structure at design time [7–9]. Even systems that are capable of learning new behaviours must pursue a target of some kind programmed at design time; for instance, if a robot has to learn to reach a given target with its arm, it will learn to move according to a predefined "motivation". On the contrary, the motivation-based robots must be able to perform actions driven by motivations which they did not possess at design time, but which they have developed by interacting with the environment.

For instance, an "intelligent" electronic device like a last-generation digital photo camera performs a long list of "intelligent" tasks: it selects the best program depending on the light, it applies a complex procedure for each program in order to select the right exposure, the right focus and a long list of related parameters. It modifies its behaviour on the basis of the environmental conditions in order to optimise the end result. Yet, notwithstanding what has been stored in its internal memory during a long journey, its behaviour does not change. No external event can modify its internal procedures as they were originally designed. Alternatively, if a 3-year-old child came on the same long journey, s/he would change. The events that happened to/around him/her would change not just his/her memory but also his/her future development, his/her internal criteria and the way in which future events will modify him/her. On an intermediate level between the 3-year-old child and the camera, there are classic artificial neural network implementations, such as speech recogniser programs. They are clever devices; they recognize normal speech pronounced by an average male or female voice. They store individuals' voices and modify their internal parameters in order to learn how to improve their performance. In this respect they are better than the camera: what happens to them modifies their behaviour. If we take two different instances of speech recogniser programs used by two different individuals, they are different: each is specialized on its owner's voice. On the other hand, if we take two cameras used by two different photographers, they are exactly the same. Even classic artificial neural networks are lacking something: their goals remain the same. Independently of their experiences they do not change their goals. On the other hand, the 3-year-old child develops new goals at any time. Experiences modify both his/her behaviour(s) and his/her criteria.

From the previous example, it is clear that there is a difference between motivation-based beings and behaviour-based beings. In the next paragraph this difference will be described in detail and a candidate architecture for artificial beings will be proposed.

A mallard duckling before its imprinting process has no idea of the visual appearance of its mother; however, since the bird sees its mother under favourable conditions, it develops a strong motivation to see the mother duck again. Before the imprinting there was no interest whatsoever for that kind of visual object, but immediately afterwards, the mallard duckling tries to keep the image of its mother inside its visual field. The motivation is 'to have the mother's image inside the visual field'. All its following actions are performed in order to make this event occur as frequently as possible. If that particular mother-bird had not shown itself to the mallard duckling, the newborn bird would not have developed any interest in it. If a different image had been shown instead of the real mother, let us say the face of Konrad Lorenz, the newborn bird would have tried to maximize the event 'to have the face of Konrad Lorenz inside the visual field'. More complex behavioural patterns are based on the same concept of repetition of an occurred event (motivation). Peter had a nice evening with Susan so he invites her again in order to repeat the pleasant experience. Mary had a pleasant time in Venice and so she plans a new holiday there.

2. Architectures for building robots: a taxonomy

Not all the motivations of biological systems are fixed at birth: they only possess a very limited, survival driven, built-in set of motivations. As they grow and develop, biological systems continuously generate new motivations on the basis of two separate factors: their genetic background and their past experience. Both are necessary in order to select a particular motivation. A mallard duckling does not have the motivation to follow its genetic mother. Yet, via its genetic background, the bird possesses the capability of choosing a bird and selecting it as a motivation. That particular bird (hopefully its mother) will become the motivation that will control the learning of the bird.

The behaviour of behaviour-based artificial structures depends on experience and motivations (goals) defined elsewhere at design time [1,2]. In complex biological systems, behaviour still depends on experience and motivations; yet, motivations are not fixed. Motivations are the result of the interaction between experience and a limited number of hardwired instincts (the ones provided by genes). In many complex biological systems, it is possible to distinguish between phylogenetic aspects and ontogenetic ones, nature versus nurture [10–12]. In general, phylogeny refers to those processes that produce new structures (genes, bodily features, behaviours, instincts) in a time scale larger than that of single individuals. On the contrary, ontogeny is limited to the life span of single individuals [10]. Furthermore, ontogeny can be driven by the phylogenetic repository (genes or instincts) or by the unpredictable contingencies of the environment. Here we endorse the view that is necessary to distinguish between goals which are determined before the actual development of an agent or subject, and those goals which are specified after the birth of the agent. We will call the former instincts and the latter motivations. The objective of this paper is to illustrate a simple set of procedures which produce motivations during development, as in the case of the imprinting procedure of birds.

Is it possible to implement instincts and motivations in an artificial system? We propose a taxonomy of architectures: a fixed control architecture, a learning architecture and an ontogenetic architecture (Fig. 1). In the first case (Fig. 1a), the system has no capability of modifying how it does what it does. There is a simple Decision Maker module, which take the input signal and produces the output on the basis of some a priori hard-wired module. Examples of this structure are simple control devices or machine automata. In the second case (Fig. 1b), the system is capable of modifying its behaviour to fulfil some a priori target. The system is capable of modifying *how* it behaves. The Decision Maker module is flanked by a Rule Maker module. The Rule Maker module can modify the a priori rules contained in the Decision Maker module on the basis of a priori hard-wired criteria. Examples of this structure are reinforcement learning or supervised learning artificial neural networks. In the third case (Fig. 1c), the system is capable of modifying not only *how* it does what it does, but also to define *what* it does. The Motivation Maker module sets the goals that have to be pursued by the Rule Maker module.

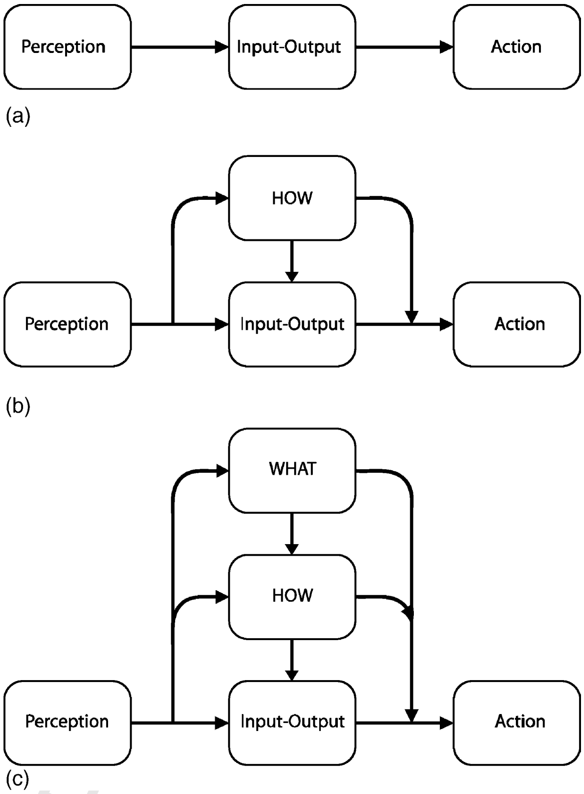

Fig. 1. Three possible architectures. In the first case (top) both *what* and *how* the system does is defined a priori; in the second case (middle) the system modifies *how* it behaves but not *what* it is doing; in the last case (bottom) the system modifies both *what* and *how* it does.

2.1. Fixed control architecture

In this case, the causal structure of the system is fixed (see Fig. 1a). There is no ontogenesis whatsoever. Notwithstanding the behavioural complexity of the system, everything happens because it has been previously coded within the system structure. A mechanical device and a complex software agent are not different in this respect: both are pre-programmed in what they must achieve and how they must achieve it. Nothing in their structure is caused by their experiences. Suitable examples of this category are Tolam's artificial sow bug [13], Braitenberg's thinking vehicles [14], Brooks' artificial insects [15,16] and recent entertainment robots like Sony's AIBO and Honda's humanoid ASIMO.

2.2. Learning architecture

A different level of structural dependency with the environment is provided by the architectures that can learn *how* to perform a task (see Fig. 1b). Behaviour-based robots can be classified in this category. Systems based on artificial neural networks are well-known examples of this kind of architecture. These systems determine how to get a given result once they have been provided with a specific motivation. The motivation can be given either as a series of examples of correct behaviour (supervised learning) or as a simple evaluation of the global performance of the system (reinforcement learning) [17,18]. In both cases some kind of learning is applied. These systems lack the capability of creating new motivations. By controlling its motors a behaviour-based robot can learn how to navigate avoiding static and dynamic obstacles. However the motivation behind this task is defined by the *a priori* design of the system. There are several examples of this kind of learning agent: Babybot at LIRA-Lab [19,20], Cog at MIT [7,21].

2.3. Ontogenetic architecture

A system that learns both *how* to perform a given task and *what* task must be performed, corresponds to an ontogenetic architecture (see Fig. 1c). This is the case for most, if not all mammals; it is true for primates and for human beings. They are systems capable of developing new motivations that do not belong to their genetic background. In the field of artificial systems there has been a series of attempts to address this problem [22–25] as well as attempts to locate similar structures in the cortical architecture of humans [26]. For their development, these systems depend more on the environment than the previous two categories. A system belonging to the first category does not depend on the environment for what it does or for how it does what it does. A system belonging to the second category does depend on the environment for how it does what it does, but not for what it does, which is phylogenetically determined. A system belonging to the third and last category depends on the environment both for what and for how it does what it does.

3. A motivation-based architecture

The proposed architecture is ontogenetic according to the previously defined taxonomy. The underlying idea is to have a physical structure (that implements the proposed architecture), which is activated by incoming events and develops motivations on this basis. The proposed architecture makes use of elementary associative processes, simple Hebbian learning and case-based reasoning.

The architecture receives an incoming stimulus and produces a signal (Relevant Signal) which depends on the value the system gives to the incoming stimulus. For instance, if the incoming stimulus corresponds to the mother's face, the system will produce a strong Relevant Signal. If the incoming stimulus corresponds to a dull grey object, the Relevant Signal will be weaker.

The architecture is made of three main modules: the Category Module that is basically a pattern classifier; the Phylogenetic Module that contains the *a priori* criteria; the Ontogenetic Module that applies Hebbian learning and develops new criteria by using the patterns stored in the Category Module. The incoming stimuli are stored in the Category Module on the basis of the Relevant Signal coming from the Phylogenetic Module and the Ontogenetic Module. At the beginning, the Relevant Signal depends on those properties of the incoming signals that are selected by the Phylogenetic Module. Subsequently, the Relevant Signal is flanked by the new signals coming from the Ontogenetic Module.

The architecture is aimed at mimicking the development of motivations in human beings. For instance, a human develops an interest for cars even if nothing in his/her phylogenetic code is explicitly directed towards cars. On the contrary, an insect cannot develop new motivations but must follow its genetic blueprint: it has no ontogenetic development. One of the issues of this architecture is to explicitly divide the ontogenetic part from the phylogenetic part.

3.1. Category Module

The category module has the role of grouping in clusters, classes and categories of stimuli coming from the external events. A discrete flow of incoming signals is the input of the category module. No hypothesis is required for their timing; no hypothesis is required

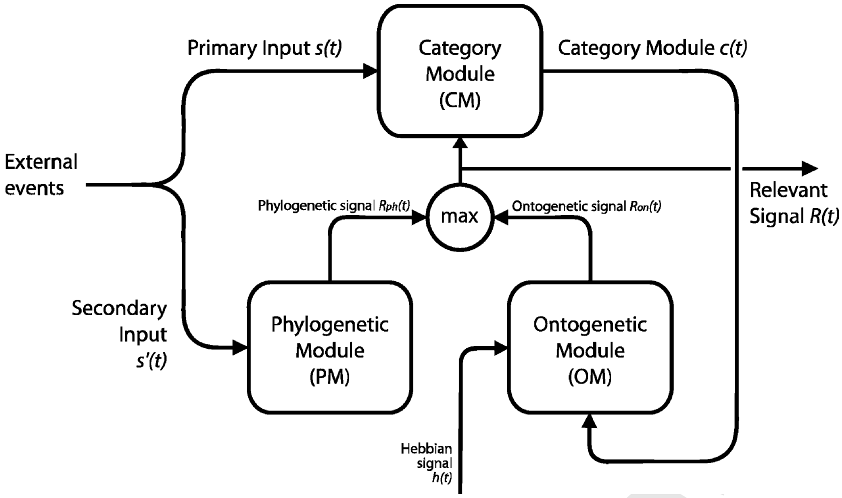

Fig. 2. A scheme for a motivation-based architecture.

for their nature. These signals could be of any kind (chunks of auditory signals, visual images, filtered visual images). Each signal is represented by a vector $\vec{s}$ of real numbers ($\vec{s} \in \mathfrak{R}^n$). CM creates a series of clusters C_i grouping classes of stimuli where each cluster C_i is a set of stored stimuli.

The process of cluster definition is based on an internally built-in criteria for clustering and on the presence of a Relevant Signal (see Fig. 2).

Whenever an incoming signal is received, a Categories Vector $\vec{c}$, which is the output of the CM, is computed. The Categories Vector contains as many elements as the clusters inside the CM at the time in which the incoming signal is analysed; the elements of $\vec{c}$ provide an indication of which cluster best represents the current stimulus. The ith element c_i is equal to the normalized difference between the maximum possible distance, usually 1, and the actual distance d_C (which will be explained below in this paragraph) between the incoming signal $\vec{s}$ and the cluster C_i. In this way, the element c_i with the greatest value corresponds to the cluster C_i that best matches the incoming signal:

$$\vec{c} = \begin{pmatrix} 1 - d_C(\vec{s}, C_1) \\ 1 - d_C(\vec{s}, C_2) \\ \vdots \\ 1 - d_C(\vec{s}, C_n) \end{pmatrix}.$$

There is no unique way to determine the distance functions d_C ($d_C : (\mathfrak{R}^n \times D) \mapsto \mathfrak{R},\ D$ cluster domain) between a vector and a cluster. The process of c_i updating requires the definition of two thresholds: one to define the minimum distance from cluster (*mcd*) and another to define the maximum distance from a cluster (*Mcd*).

The CM tunes its activity on the basis of the Relevant Signal. As shown in Fig. 3, the Relevant Signal ($R(t)$) is the sum of two different signals: the Relevant Ontogenetic Signal ($R_{\mathrm{on}}(t)$) and the Relevant Phylogenetic Signal ($R_{\mathrm{ph}}(t)$), according to

$$R(t) = \max(R_{\mathrm{on}}(t), R_{\mathrm{ph}}(t)).$$

If and only if the Relevant Signal is active, every time a signal is received, the CM performs the following actions:

(i) If the stimulus is too similar to the already stored stimuli, do nothing ($d_C(\vec{s}, C_i) < mcd$).
(ii) If the stimulus is sufficiently similar to one of the previously created clusters ($mcd \le d_C(\vec{s}, C_i) \le Mcd$), the stimulus is added to that cluster.
(iii) If the stimulus is not sufficiently similar to any of the stimuli already stored, a new cluster is created ($d_C(\vec{s}, C_i) > Mcd$).

By storing a stimulus only if the Relevant Signal is active, the system does not assign new resources for every incoming signal (the first rule is useful to avoid to store equivalent stimuli).

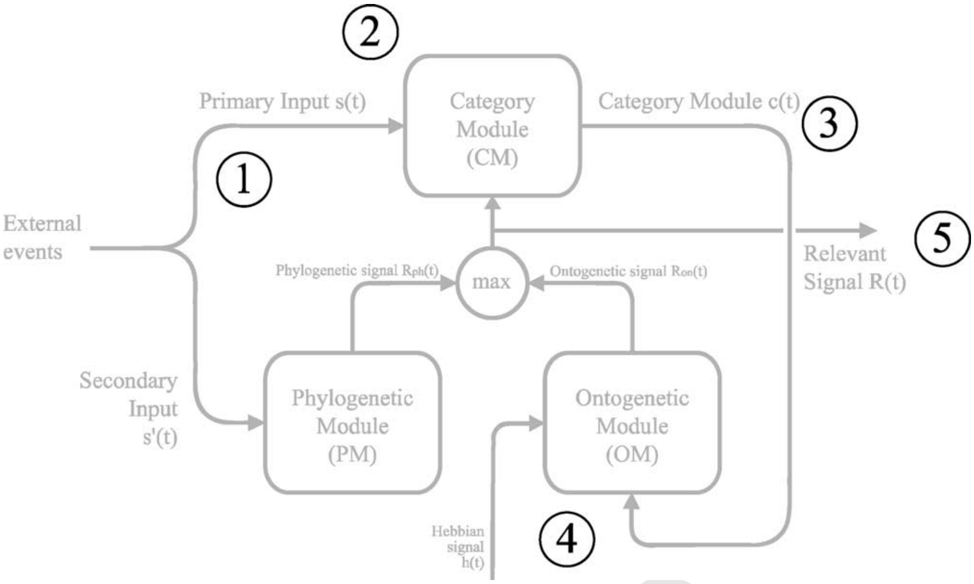

Fig. 3. Timing of operations.

3.2. Phylogenetic Module

The Relevant Phylogenetic Signal, $R_{\mathrm{ph}}(t)$ is produced by the Phylogenetic Module (PM, Fig. 2). This module is the only one that has some built-in criteria concerning the relevant properties of the incoming signal (for instance, the structure of the Category Module does not present any similar feature). Functionally, it has the same role as the genetic instincts in biological systems. It is similar to saliency systems or attention mechanisms [27]: it selects which stimuli are worth the attention of the system. A Phylogenetic Module works in two different ways: (i) it autonomously produces a signal on the basis of some internal criteria; (ii) it produces a signal on the basis of some external events. In the second case the PM needs some kind of elementary capability in order to recognize particular occurrences of events in the external environment (the presence of the mother, the presence of soft or brightly coloured objects).

For instance, a baby looks with more curiosity at brightly coloured objects than at dull colourless objects, independently of any past experience. This behaviour requires the existence of a hardwired function looking for a relevant property of images (saturated colours). This module provides criteria that can be used to select correct actions (for instance those actions that maximize the presence of the interesting stimuli).

The performance of the Phylogenetic Module is implemented by the function $f_{\mathrm{phylogenetic}} : \Re^n \mapsto [0, 1]$ applied to the input $\vec{s}'(t)$ that is a signal from which it is possible to know if something relevant is happening. The signal $\vec{s}'(t)$ comes from the external environment. For instance it could be a verbal approval for a specific event; or it could be a reward/punishment following a behaviour. The resulting output is:

$$R_{\mathrm{ph}}(t) = f_{\mathrm{phylogenetic}}(\vec{s}'(t)).$$

A system could contain one phylogenetic function for each kind of event the designers want the system to react to. For instance, there could be a function to detect the presence of round-shaped objects (a prototype for faces), a function to detect the presence of objects with highly saturated colours and a function to detect the presence of moving objects. At every instant there could more than one function to signal that something interesting is going on: more than one $f_{\mathrm{phylogenetic}}$ function can be evaluated. The output of the Phylogenetic Module is the maximum among the outputs of the different $f_{\mathrm{phylogenetic}}$ functions whose input is always

$\vec{s}'(t)$:

$$R_{\text{ph}}(t) = f_{\text{phylogenetic}}(\vec{s}'(t)) = \max_{i=1,\ldots,m}(f^{i}_{\text{phylogenetic}}(\vec{s}'(t))),$$

where m is the number of kinds of events which the system is capable of reacting to from the beginning. So m is the number of elementary instincts (each corresponding to a separate phylogenetic function) that the system possesses. It is important to outline that (i) the Phylogenetic Module is incapable of adaptability and that (ii) the Phylogenetic Functions might be very simple because their role is to orient the attention of the CM towards certain classes of objects, albeit making mistakes.

In a multisensory system, each sensory modality can be used as an alternative source of information for another sensory modality. In real biological systems, there are plenty of sources of information (like pain, skin receptors, tactile information) that can be the input $\vec{s}'(t)$ of the PM. The same sensory modality can be the input both for the PM and for the CM. If this happens, it is possible to assume that

$$\vec{s}'(t) = \vec{s}(t).$$

If the system were composed of just the PM and the CM, the system would be a reinforcement learning system.

3.3. Ontogenetic Module

Whereas the Phylogenetic Module has built-in criteria about the nature and the relevant properties of the incoming signal, the Ontogenetic Module selects new criteria on the basis of experience. Functionally it has the same role as the acquired ontogenetic criteria in biological systems.

The Ontogenetic Module acts as a gate for the incoming output of the CM $\vec{c}(t)$. The gating procedure is implemented by means of an internal vector $\vec{g} = (g_1, \ldots, g_n)^t$ which has the same number of elements as the clusters in CM. $\vec{g}$ is contained inside the Ontogenetic Module. The output of the OM is computed as the maximum among the elements g_i times the elements c_i of the CM:

$$R_{\text{on}}(t) = \max_{i=1,\ldots,n}(g_i \cdot c_i). \tag{1}$$

The g_i have the role of gates (hence the use of the letter g) in order to let or to prevent the effect of the output of the CM to propagate further. If the g_i are positive, the corresponding c_i contribute to the Relevant Signal. Since the c_i represent the stored categories acquired during the experiences of the system, the R_{on} is the result of the ontogenetic development.

The result of the architecture is to produce a new reinforcement signal $R_{\text{on}}(t)$, which depends only on the actual experiences of the system (i.e. on the received input signals). Here $R_{\text{on}}(t)$ is called the Relevant Ontogenetic Signal because it derives from the actual experiences of the system. It is the result of the development of an individual system and its history; hence it pertains to its ontogeny.

The vector $\vec{g}$ is the result of a Hebbian learning implementation with respect to the simultaneous occurrence of signals $(h(t), \vec{c}(t))$; learning happens when $h(t)$ and $c_i(t)$ fire simultaneously. The value of g_{i} approaches the value 1 if the signal $h(t)$ and the component $c_i(t)$ are correlated in time. A possible function is the following:

$$g_i = \frac{2}{\pi}\arctan\left(\int_{t_0}^{t}(h(\tau)\cdot c_i(\tau))^q \,\mathrm{d}\tau\right),$$

where $q \in [0, 1]$ can be used to tune the speed of learning. The element $c_i(t)$ corresponds to the ith elements of the output of the CM, and $h(t)$ is the signal that controls the performance of the Ontogenetic Module.

Four different choices are possible for $h(t)$: (i) $h(t)$ is set to equal a positive constant; (ii) $h(t)$ is an a priori time variant function; (iii) $h(t)$ is set to equal the output of the PM ($h(t) = R_{\text{ph}}(t)$); (iv) $h(t)$ is connected to some independent sources of signals that are linked to the environment.

In the first case, since $h(t)$ is a constant, each g_i is proportional to how much the corresponding category has been represented in the input stimulus $\vec{s}$. The more frequent and the more intensely a category matches the input, the greater its effect on the Relevant Ontogenetic Signal will be.

In the second case, $h(t)$ varies in time according to an a priori time variant function. Each g_i will correspond to those categories that are representative of the input during those periods in which $h(t)$ is larger. For instance, $h(t)$ might be high in an initial period and then it might vanish: the Ontogenetic Module will accept only those categories that are representative of the input during the initial period.

In the third case ($h(t) = R_{\text{ph}}(t)$), in an early stage, each g_i will be representative of those categories that occur at the same time as the activations of the Phylogenetic Functions. Eventually, there can be a drift from the categories selected by Phylogenetic Functions to the new categories selected by the Ontogenetic Functions.

In the fourth case, $h(t)$ is assigned to a separate source of signals; different sensor modalities can be associated. For instance, the incoming signal $\vec{s}$ might be visual, while $h(t)$ might be the result of the tactile sensory modality. As a result, the g_i would be higher when the two different sensory modalities are simultaneously present.

The OM produces a new reinforcement signals that are indirectly related to the phylogenetic structure of the system. The interaction between the OM and the CM generates a new set of functions, which are the ontogenetic equivalent of the phylogenetic functions:

$$f^i_{\text{ontogenetic}}(\vec{s}) = g_i(t)(1 - d_C(\vec{s}, C_i)).$$

At each instant, the ontogenetic functions $f^i_{\text{ontogenetic}}(\vec{s})$ compute the relevant ontogenetic signal. Their form depends on the information stored in the g_i and in the C_i, which is the result of the past history of the system. We can rewrite Eq. (1) as follows:

$$\begin{aligned} R_{\text{on}}(t) &= \max_{i=1,\ldots,n} (g_i \cdot c_i) \\ &= \max_{i=1,\ldots,n} (f^i(\vec{s}(t))) = f_{\text{ontogenetic}}(\vec{s}(t)). \end{aligned}$$

3.4. How the architecture works

The main goal of the architecture is to create a structure that can be changed completely by its own experiences. In the architecture there is a clear-cut division between the phylogenetic part (the a priori section) and the ontogenetic part produced by the interaction with the environment.

As it is possible to see in Fig. 3, the timing of operations is the following. First the incoming stimulus (1) is compared to each cluster of stored vectors (2) and, as a result, the output vector is computed on the basis of the current structure of the network (3). Then the Ontogenetic Signal is computed by the Ontogenetic Module (4). Finally, the Ontogenetic Signal is combined with the Phylogenetic Signal to produce the Relevant Signal that is sent to the Category Module and to the output (5). Only at this stage the Category Module modifies its clusters on the basis of both the incoming stimuli and the Relevant Signal. If the Ontogenetic Module were not active, the architecture would stop its development and become a pure feed forward network.

4. Experimental results: the emergence of motivations

To test the architecture, an experiment was carried on in which a robot embodying the proposed motivation-based architecture develops a new motivation on the basis of its own experiences. In the experiment, an incoming class of visual stimuli (not coded inside the architecture) produces a modification in the system's behaviour differently from what happens in behaviour-based robots. In behaviour-based robots the transition between different behaviours elicited by a motivation is defined by the designer and does not depend on a newly produced self-motivation. By interacting with the environment, the system adds a new motivation that changes not only *how* (behaviour) but also *what* (motivation at the basis of behaviour) the system is doing. The system has, in this preliminary experiment, a single behaviour: directing or not its gaze towards objects. This behaviour is not what is learned by the architecture; it is used by the architecture to show the effects of its new motivation.

A series of different shapes associated with colours were presented to the robot. The system is equipped with a phylogenetic motivation that is aimed at very coloured objects; a colourless stimulus, independently of the shape, does not elicit any response. Since the system has an ontogenetic module it develops further motivations directed towards classes of stimuli different from those relevant for its phylogenetic module. After a period of interaction with the visual environment (constituted by a series of elementary coloured shapes), the robot is motivated by colourless shapes also. The system shows the capability to develop a motivation (by directing its gaze towards the stimulus) that was not envisaged at design time and that is the result of the ontogenetic development.

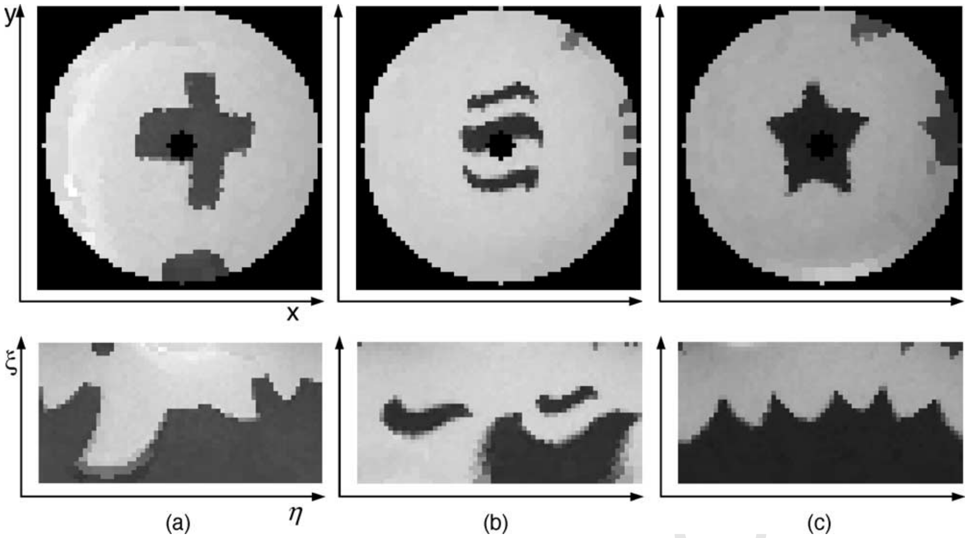

Fig. 4. The Cartesian (upper row) and log-polar (lower row) images for a cross (a), a wave (b), and a star (c).

4.1. *Robotic setup*

A robotic head with four degrees of freedom has been adopted as robotic setup. We used the EuroHead developed for navigation (Pan: range = 45°, velocity = 73°/s, acceleration = 1600°/s, resolution = 0.007°; Tilt: range = 60°, velocity = 73°/s, acceleration = 2100°/s, resolution = 0.007°) [28]. However, we only used two degrees of freedom of the head since, for the purpose, of this experiment only a pointing device was needed. Robots characterized by more sophisticated morphologies could have been used to perform more complex tasks. However, in this preliminary stage of research, an exceedingly complex combination of morphological, behavioural and computational factors would have been extremely difficult to be interpreted.

4.1.1. *Sensory Module*

The robotic head was equipped with a videocamera capable of acquiring log polar images [29,30]. Log polar images (Fig. 4) are defined by

$$x = \rho\cos(\theta), \qquad y = \rho\sin(\theta),$$

$$\theta = k \cdot \eta, \qquad \rho = r_0 \cdot a^{\xi},$$

together with

$$\rho = \sqrt{x^2 + y^2}, \qquad \theta = \arctan\left(\frac{y}{x}\right),$$

$$\eta = \frac{\theta}{k}, \qquad \xi = \ln_a\left(\frac{\rho}{r_0}\right).$$

These images offer two main advantages among the others: (i) invariance with respect to rotation and scaling; (ii) reduced number of pixels with wide field of view. Furthermore, in this case the use of log-polar images allows an implicit selection of a target (due to the space-variant distribution of receptors). In foveated visual apparatus, the central part of the image corresponds to the majority of pixels and thus when an object is fixed, its image is much more important than the background. As a result, there is no need to perform explicit selection of a target; the direction of the gaze implicitly selects its own target.

The robotic head has two degrees of freedom: the camera is capable of a tilt and pan independent motion (Fig. 5). Since the head was able to move only in a limited span with the pan and the tilt (40° each) it was possible to determine which point on the board was looked at. By measuring the angle position of each saccade is possible to measure which region of the visual stimulus is more frequently observed by the head.

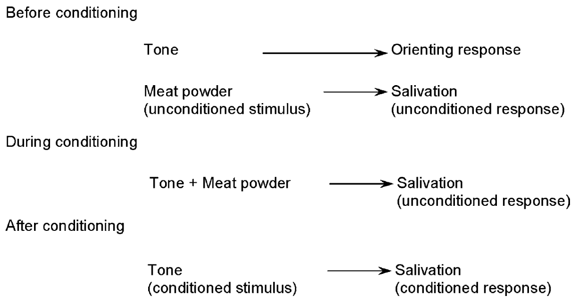

Fig. 5. Sensory and motor setup.

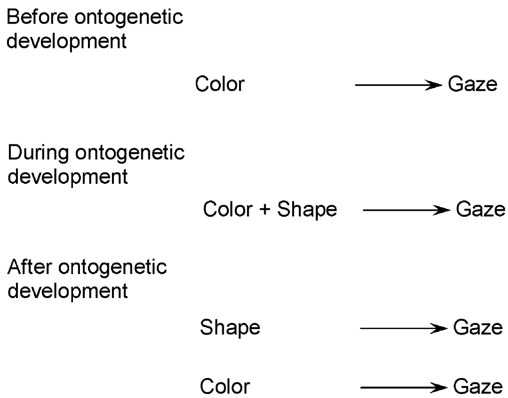

Fig. 6. The probability density function on the basis of the control parameter λ.

4.1.2. Motor Module

The robotic head is programmed to make random saccades; a Motor Module generates saccades on the basis of an input signal λ that controls the probability density of the amplitude r. The motor input λ is the only signal needed by the Motor Module in order to control its actions. The probability function of the angle has a uniform distribution from 0 to 2π. The probability function of the amplitude is equal to

$$p(r, \lambda) = \frac{1}{\int_{-r_{\max}}^{+r_{\max}} \mathrm{e}^{-\lambda \cdot \rho^2}\, \mathrm{d}\rho} \mathrm{e}^{-\lambda \cdot r^2},$$

where r is the random variable for the amplitude (Fig. 6). If λ is low (near to 0), the probability density is almost constant, therefore there is an equal probability for each amplitude. If λ is higher, a small amplitude is more probable.

The rationale of this probability schema resides on the fact that the motor unit should mimic an exploratory strategy. When a visual system explores a field of view, it makes large random saccades. When it fixates an interesting object, it makes small random saccades.

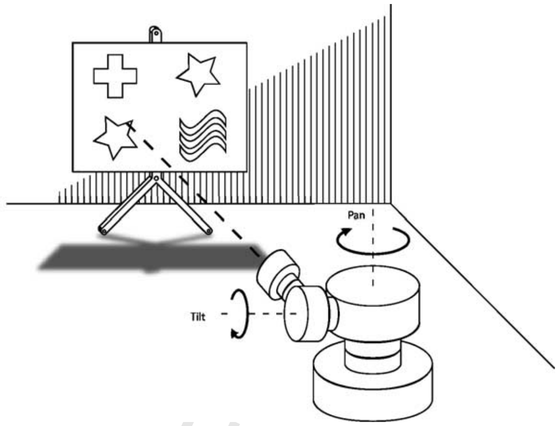

Fig. 7. The proposed architecture (named 'Artificial Motivations') is independent of the sensor and motor parts. It contains some basic information about the relevant signal to bootstrap the system.

4.2. Architecture implementation

It is important to note that no modification has been made to the architecture on the basis of the particular properties of the robotic setup. The architecture could be used in a completely different robotic setup, with completely different input and output signals without having to change (Fig. 7).

4.2.1. Category Module

The Category Module creates clusters of incoming stimuli on the basis of the Relevant Signal. Each of these clusters corresponds to a category. Further than the Relevant Signal, the CM uses an internal criteria to control the cluster creation: the distance function $d_C(\vec{v}, C)$ between a vector and a cluster. This distance is derived from a distance function between vectors $d(\vec{v}, \vec{w})$: $d : (\mathfrak{R}^n \times \mathfrak{R}^n) \mapsto \mathfrak{R}$, d continuous. must be a distance between vectors. Suitable candidates for this function are the Minkowski function or the Tanimoto distance or the correlation function [31]. In the experiment the function is implemented as such

$$d(\vec{v}, \vec{w}) = C(\vec{v}, \vec{w})$$

$$= \frac{1}{2}\left(1 - \frac{\sum(v_i - \mu_v)(w_i - \mu_w)}{\sqrt{\sum(v_i - \mu_v)^2 \cdot \sum(w_i - \mu_w)^2}}\right).$$

The advantages of this function are that it is more robust to change in average value, more resistant to noise.

On the basis of $d(\vec{v}, \vec{w})$ it is possible to define the distance function between a vector and a set of vectors.

Two solutions are easily implemented. First, the distance between a vector and a cluster is computed as the minimum distance between a given vector $\vec{v}$ and all the vectors belonging to a given set C:

$$d_C(\vec{v}, C) = \min_{\vec{w} \in C}(d(\vec{v}, \vec{w})).$$

Yet the above approach is computationally expensive since it entails that, for a given set, all vectors must be stored somewhere. A different approach is based on the assumption that it is possible to compute the average distance, which is equal to the distance with the centre of gravity. If M is the number of elements of set C, and $\vec{c}$ is its mean vector:

$$d_C(\vec{v}, C) = \frac{\sum_{\vec{w} \in C} d(\vec{v}, \vec{w})}{M} = d(\vec{v}, \vec{c}).$$

This approach has the advantage that is sufficient to keep in memory only the mean vector of each set. This means that each set can be stored as a vector. The results are based on this solution. It is important to note that no specific information about the nature of the vectors is part of the Category Module.

4.2.2. *Phylogenetic and Ontogenetic Modules*

The Phylogenetic Module contains the built-in criteria to bootstrap the system. In this case the built-in criterion consists in selecting brightly coloured objects. This module implements the following phylogenetic function:

$$R_{\mathrm{ph}} = \frac{\sum \mathrm{Saturation}(\eta, \xi)}{N},$$

where R_{ph} is the Relevant Signal, Saturation(η, ξ) the colour saturation at the pixel (η, ξ) in log polar coordinates, and N the total number of pixels in the image. Therefore R_{ph} is proportional to the average level of colour saturation. This phylogenetic function represents the only built-in part of the architecture. It corresponds to the phylogenetic contribution to the development of the system. The Relevant Signal R_{ph} is used to control the motor behaviour: even if the architecture were composed only by the phylogenetic module, it would drive the system towards highly colour saturated targets. In the neighbourhood of a coloured object oscillations of this function are possible, however there will always be a maximum in correspondence of an image centred on the coloured target. When the target is in the fovea of the log polar image, it corresponds to the maximum number of pixels.

The Ontogenetic Module corresponds to the definition we gave in Section 3.3; no modifications were needed.

4.3. *A comparison with Pavlov's classic conditioning*

As a final argument, we would draw a comparison with Pavlov's classic experiment of conditioning (Figs. 8 and 9). The reasons for this comparisons are two-fold: (i) there are strong similarities; (ii) there is evidence that many cognitive learning processes could be reduced to Pavlov's associationism [5,32]. In Pavlov's case, the focus was on the capability of modifying the relation between a given stimulus and a given response. Although, Pavlov's dog was able to select a different stimulus (the ring of the bell), the focus was more on the fact that the dog was capable of linking the stimulus to a behaviour (the salivary response) rather than to the capability of selecting a given stimulus from the continuum of the environment.

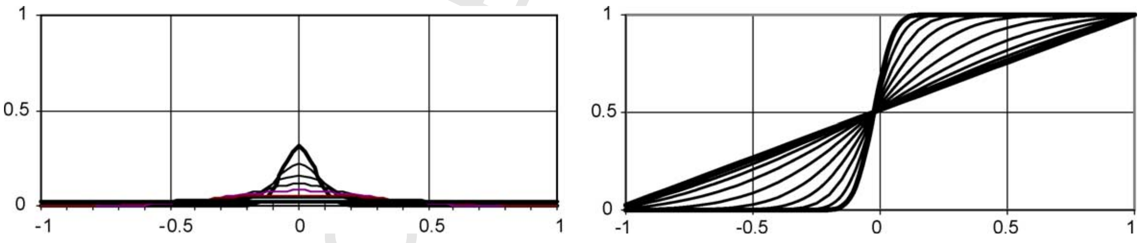

Fig. 8. The three stages of conditioning in the classical Pavlov experiment.

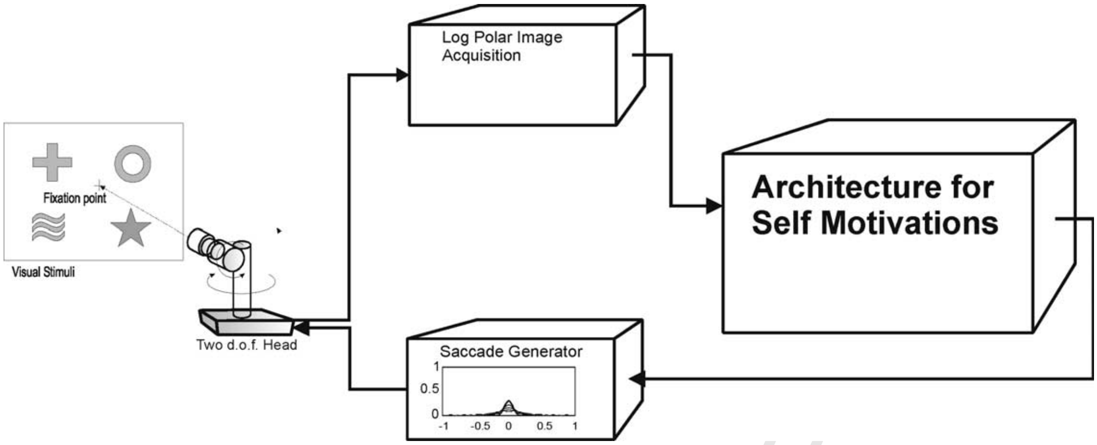

Fig. 9. The three stages of ontogenetic development.

In Pavlov's experiment, there are two hardwired receptors for two different kinds of stimuli (sound of a bell and meat powder): one is a neural structure capable of recognizing the presence of food and another is a neural structure capable of recognizing the ring of a bell. Before the conditioning process, the behavioural response (the salivation) was only connected with the presence of food. During the training, the conditioned response became stronger, more drops of saliva were secreted. The learning consisted in the creation of a connection between the conditioned stimulus and the response.

In our case, the conditioned stimulus does not exist before the conditioning process. The machine is not capable of recognizing the unconditioned stimulus (the shape of an object). It only recognizes coloured objects. At first sight, our experiment might recall Pavlov's experiment. It could be argued that the Phylogenetic Stimulus corresponds to the Unconditioned Stimulus, and the Ontogenetic Stimulus corresponds to the Conditioned Stimulus; and the Developmental Signal might correspond to the Response (first Unconditioned and then Conditioned). This is not the case. In the described circumstances, since the colour was presented conjointly with the shape of an object, a new ontogenetic stimulus (the shape) is added to the machine's repertoire of stimuli.

A useful concept is that of the *Umwelt* of a subject [33,34]: the set of all events which can interact with a subject given its sensory/motor/cognitive capabilities. In the case of the ontogenetic development of new motivations, the Umwelt of the machine is increased and enlarged to a new kind of event. Two things have happened: (i) the machine has learned to recognize something which was previously unknown to it; (ii) the machine has linked such new stimulus to a given motor behaviour. Pavlov's experiment highlighted the fact that the dog had learnt a relation between an already assessed stimulus to a motor response. The goal of our experiment is to create the capability of recognizing new stimuli.

4.4. Experimental results

We presented different sets of visual stimuli to the system. A first set consisted in a series of colourless geometrical figures as shown in Fig. 10a on the left. The frequency with which the system was looking at different points was measured. The system spent more time on stimuli corresponding to its motivations by reducing the amplitude of its saccades. At the beginning the system was looking around completely randomly with large saccades since its Ontogenetic Module was unable to catch anything relevant and the Phylogenetic Module was programmed to look for very saturated coloured objects, which were absent in the first set.

Subsequently we presented a different stimulus: a series of coloured figures (Fig. 10b on the left). The difference is shown in Fig. 10b. The head spent more time on the coloured shapes instead on the white background because of the phylogenetically implanted rule.

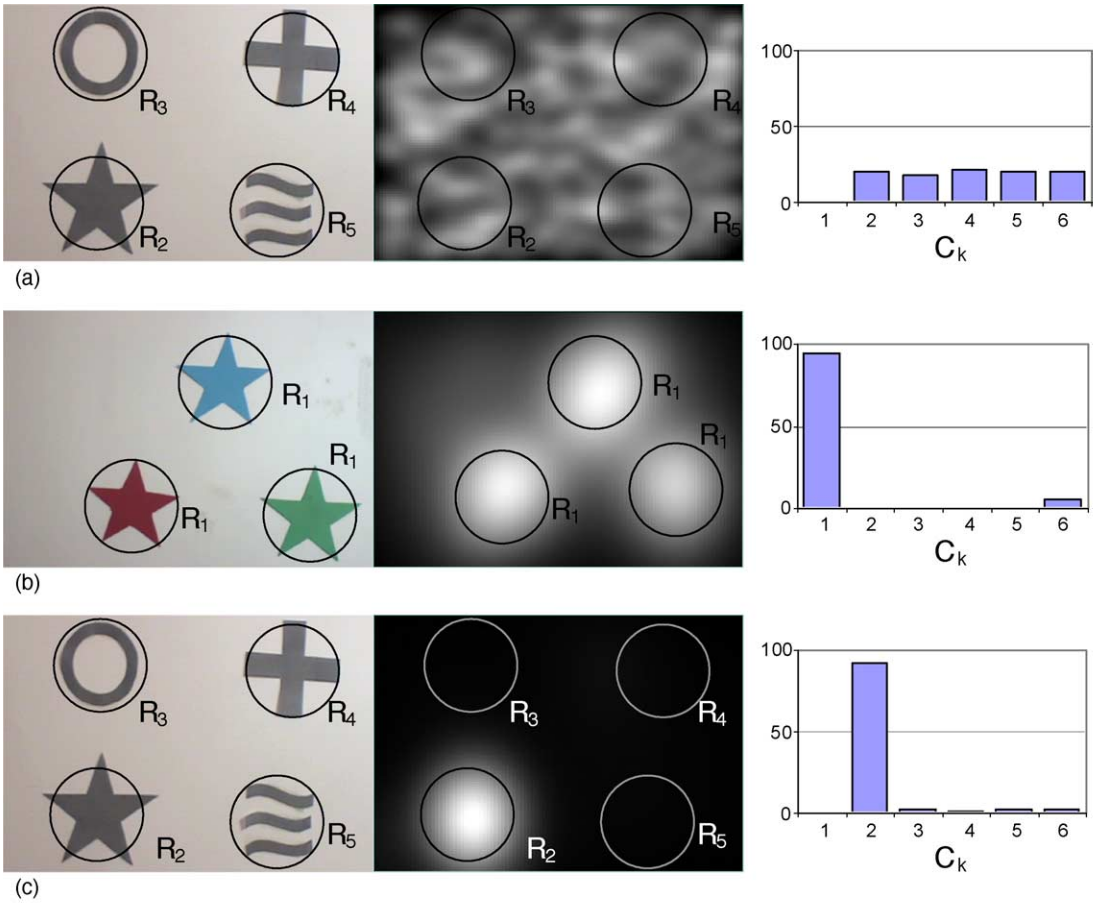

Fig. 10. Experimental results.

Finally we presented again the initial stimulus (the set of colourless shapes). The system spent more time on the colourless shapes than on the background (Fig. 10c). The behaviour of the system changed since the system added a new motivation (shapes) to the previous ones (saturated colours).

In order to measure the different behaviour of the system, the time spent by the system on each shape was measured in two different ways: a qualitative one (the middle column) and a quantitative one (the right column).

To get a qualitative visual description of how much time was spent by the system on each point of its field of view, we assigned to each point of the visual field an intensity value proportional to the normalized time the system gaze spent on it. The images in the centre of Fig. 10a–c were generated after 10^3 saccades (equivalent to about 500 s). The field of view of the head was divided in a 64×64 array. For each point (i, j) in the visual field, the amount of time the gaze of the head was directed on it was computed:

$$t_{i,j} = \text{total time spent looking at point } (i, j).$$

The intensity of the point was then set proportional to a normalized value of $t_{i,j}$. With the first set of visual stimuli, the resulting image is in Fig. 10a (middle). The system does not show any polarization towards a specific part of the field of view. The behaviour of the system is completely different to its response to the coloured stimulus: there are three definite centres of interest (Fig. 10b, middle). However, after the interaction with the star has shaped a new goal which becomes part of its behaviour. In Fig. 10c (middle) the original grey stimuli produces a completely different response: the grey star became a centre of interest.

To get the quantitative measure (right column), we measured the time spent by the head inside the circular areas shown in the left column surrounding the stimuli: a rough indicator of the time spent looking at a certain shape. The region of interest were named according to the following notation: the coloured figures (R_1), grey star (R_2), grey cross (R_3), grey waves (R_4), grey circle (R_5). In the graphs of Fig. 10 on the right, for each region, the normalized time was computed according to the following formula:

$$c_k = 100\frac{\sum_{(i,j)\in A_k} t_{i,j}}{\sum t_{i,j}},$$

where $t_{i,j}$ is the same of the previous formula, and A_k corresponds to a set of region: A_1 corresponds—for each group of stimuli—to what is not occupied by the stimuli; A_2 corresponds to the union of the three areas occupied by the three coloured stars (R_1); while $A_{2,3,4,5}$ correspond, respectively, to the four regions occupied by the grey shapes ($R_{2,3,4,5}$).

In order to test the efficacy of the architecture presented, the experiment was repeated in a simulated environment. In this way it was possible to check its soundness and generalize its software implementation. In the simulated version of the experiment, similar stimuli were presented and a simulated gaze was directed towards different points of the image. The images used were 1024×768 pixels; the artificial retina had a 64 pixels diameter. In Fig. 11, the experimental results are visible. All the other parameters exactly match the Eurohead. Instead of computing a frequency density value to each point of the field of view, a collection of 10^3 is displayed for each of the presented stimuli. From a qualitative point of view, the relevant changes in the behavioural and motivational property of the system are clearly visible.

In future, we are planning to implement this architecture in more complex robotics setup and in more realistic environment. However, we believe that the general principle is already clearly illustrated by these simplified experiments.

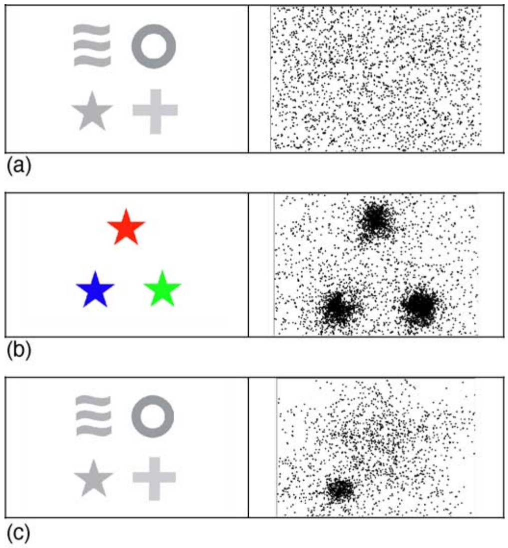

Fig. 11. A simulated version of the experiment: on the left the artificial stimuli, on the right the measured fixation points.

5. Conclusions

Ever since Grey Walters' wrote about his turtles the history of robots has chronicled their efforts to establish a relationship with the environment. The transition from deliberative robots to reactive robots, then to behaviour-based robots bears witness to this trend. The recent appearance on the market of entertainment robots sheds new light on motivation-based robots.

Environment driven motivations provide the internal criteria for the development of artificial beings and supply their means and goals: how they do what they do.

Ontogenetic development allows the artificial being to elaborate the criteria on which it can associate external stimuli. In the aforementioned experiment the visual stimuli were associated first on the basis of a phylogenetic criterion (the colour saturation), then on the basis of an ontogenetic criterion derived from the system experience (shape). The ontogenetic architecture allows to self-associate different stimuli (the colour and the shape) on the basis of the interaction with the environment. A new motivation (looking for a shape) is the product of the individual history of the architecture in a given environment. Recently, self associative learning has been identified as the possible key to the development of consciousness [5]. It follows that an ontogenetic architecture based on environment-derived motivations might provide the basis for the development of an artificial conscious robot.

In this paper we have used the intentionalistic mentalistic vocabulary to introduce intentional concepts such as 'motivations' and 'experience'. A correct definition of these terms applied to artificial beings should be free from any ontological or linguistic commitment. This tenet is evident when instead of biological beings we have to deal with artificial systems since it is not clear whether they possess intentional properties or not. For instance, if we are dealing with human beings, it is safe to use words like 'intentions', 'motivations', 'experience'. However, if we are dealing with robots or other kinds of artificial systems, it is ambiguous to use the same terminology. Edelman and Tononi wrote that: "to understand the mental we may have to invent further ways of looking at brains. We may even have to synthesize artifacts resembling brains connected to bodily functions in order fully to understand those processes. Although the day when we shall be able to create such conscious artifacts is far off we may have to make them before we deeply understand the processes of thought itself" [35].

References

[1] R.C. Arkin, Behavior-based Robotics, MIT Press, Cambridge, MA, 1999.
[2] D. McFarland, T. Bosser, Intelligent Behavior in Animals and Robots, MIT Press, Cambridge, MA, 1993.
[3] A. Lansner, Motivation based reinforcement learning in autonomous robot, SANS (1994).
[4] P.M. Merikle, T.A.W. Visser, Conscious and unconscious processes: the effects of motivation, Consciousness and Cognition 8 (1999) 94–113.
[5] P. Perruchet, A. Vinter, The self-organizing consciousness, Behavioral and Brain Sciences (2002).
[6] J. Zlatev, The epigenesis of meaning in human beings, and possibly in robots, Minds and Machines 11 (2001) 155–195.
[7] R.A. Brooks, et al., The Cog project: building a humanoid robot, in: Nehaniv (Ed.), Computation for Metaphors, Analogy, and Agents, Springer-Verlag, Berlin, 1999.
[8] R.A. Brooks, A.M. Flynn, Robot beings, in: Proceedings of the IEEE Intelligent Robots and Systems Conference, Tsukuba, 1989.
[9] E.M.A. Ronald, R. Poli, Surprise versus unsurprise: implications of emergence in robotics, Robotics and Autonomous Systems 37 (1) (2001) 19–24.
[10] S.J. Gould, Ontogeny and Phylogeny, Harvard University Press, Cambridge, MA, 1977.
[11] J.L. Elman, et al., Rethinking Innateness: A Connectionist Perspective on Development, MIT Press, Cambridge, MA, 2001.
[12] S. Oyama, The Ontogeny of Information, Duke University Press, Durham, 1985/2000.
[13] E.C. Tolman, Prediction of the vicarious trial and error by means of the schematic sowbug, Psychological Review 23 (1939) 318–336.
[14] V. Braitenberg, Vehicles: Experiments in Synthetic Psychology, MIT Press, Cambridge, MA, 1984.
[15] R.A. Brooks, Intelligence without representations, Artificial Intelligence Journal 47 (1991) 139–159.
[16] R.A. Brooks, New approaches to robotics, Science 253 (1991) 1227–1232.
[17] R.S. Sutton, A.G. Barto, Reinforcement Learning, MIT Press, Cambridge, MA, 1998.
[18] M. Kuperestein, INFANT neural controller for adaptive sensory-motor coordination, Neural Networks 4 (1991) 131–145.
[19] G. Metta, G. Sandini, J. Konczak, A developmental approach to visually guided reaching in artificial systems, Neural Networks 12 (1999) 1413–1427.
[20] R. Manzotti, et al., Disparity estimation in log polar images and vergence control, Computer Vision and Image Understanding 83 (2001) 97–117.
[21] C. Ferrell, C. Kemp, An ontogenetic perspective to scaling sensorimotor intelligence, in: Proceedings of the 1996 AAAI Fall Symposium on Embodied Cognition and Action, 1996.
[22] D. Gachet, et al., A software architecture for behavioural control strategies of autonomous systems, in: Proceedings of the Conference, IECON'92, San Diego, CA, 1992.
[23] J.J. Weng, Cresceptron, SHOSLIF: toward comprehensive visual learning, in: S.K. Nayar, T. Poggio (Eds.), Early Visual Learning, Oxford University Press, New York, 1996.
[24] K. Fukushima, S. Miyake, T. Ito, Neocognitron: a neural network model for a mechanism of visual pattern recognition, IEEE Transactions on Systems, Man, and Cybernetics 13 (1983) 826–834.
[25] K. Fukushima, M. Okada, K. Hiroshige, Neocognitron with dual C-cell layers, Neural Networks 7 (1) (1994) 41–47.
[26] T. Togawa, K. Otsuka, A model for cortical neural network structure, Biocybernetics and Biomedical Engineering 20 (3) (2000) 5–20.
[27] S. Grossberg, The link between brain learning, attention, and consciousness, Consciousness and Cognition 8 (1999) 1–44.
[28] M. Vincze, et al., A system to navigate a robot into a ship structure, Machine Vision and Applications (2003) 14.
[29] G. Sandini, V. Tagliasco, An anthropomorphic retina-like structure for scene analysis, Computer Vision Graphics and Image Processing 14 (1980) 365–372.
[30] G. Metta, et al., Development: is it the right way towards humanoid robotics? in: Proceedings of the Conference, ISA-6, IOS Press, Venezia, 2000.
[31] R.O. Duda, P.E. Hart, D.G. Stork, Pattern Classification, 2nd ed., Wiley, New York, 2001, xx, 654.
[32] I.P. Pavlov, Selected Works, University Press of the Pacific, Honolulu, Hawaii, 1955/2001.

[33] J.v. Uexküll, A Stroll through the Worlds on Animals and Men: A Picture Book of Invisible Worlds, International University Press, New York, 1934.
[34] T. Ziemke, The construction of ’reality’ in the robot: constructivist perspectives on situated artificial intelligence and adaptive robotics, Foundations of Science 6 (1–3) (2001) 163–233.
[35] G.M. Edelman, G. Tononi, A Universe of Consciousness. How Matter Becomes Imagination, Allen Lane, London, 2000.

Riccardo Manzotti was born in Parma, Italy, in 1969. He received his Laurea Degree in Electronic Engineering in 1994, in Philosophy in 2004, his Ph.D. in Robotics in 2000 from University of Genoa. He is currently Research Fellow for the Institute of Human and Environmental Sciences (IULM University, Milan). His present research activities focus on the areas of artificial intelligence, consciousness and robotics.

Vincenzo Tagliasco was born in Savona, Italy, in 1941. He is currently a Full Professor of Bioengineering at the University of Genoa, Italy. His main field of work includes vision and motor system in human beings. At present he teaches and carries out research work in the area of artificial consciousness.