The "What" problem: the emergence of new goals in a robot

Riccardo Manzotti
KTEL
Institute of Environmental and Human Sciences
IULM University, Via Carlo Bo, 8, 20143, Milan
riccardo.manzotti@iulm.it

Marco Alessandro Villamira
KTEL
Institute of Environmental and Human Sciences
IULM University, Via Carlo Bo, 8, 20143, Milan
marco.villamira@iulm.it

***Abstract* - Biological and cognitive systems have the capability of developing new goals during phylogenesis of species or during ontogenesis of single individuals. On the other hand, current artificial cognitive systems focus on how achieving a given fixed set of hard-wired goals. They search an optimal solution of a problem, given a set of goals and a set of optimization criteria. They look for "how" to achieve a given goal. Natural agents develop new goals in order to cope with partially unknown and ever changing environment. They must find "what" they want to achieve and not only "how". The development of new goals on the basis of the interaction with the environment is here defined the "what" problem. The development of a collection of goals permits to redefine the concept of Umwelt in what could be considered the teleological Umwelt of an agent. The objective of this paper is twofold: i) to outline the "what" problem and ii) to describe a robotic architecture capable of addressing it.**

Index Terms - agent, goal, teleological system, robot, ontogenesis, ontogenesis.

I. WHAT AND HOW

Current implementations of artificial systems focus mainly on intelligent algorithms to achieve a fixed goal (or a fixed set of goals) with optimal performance. They look for "how" to achieve a given goal. Many biological systems do not only achieve an optimal performance with respect to a given set of hardwired goals, they develop new goals. They must find "what" they want to achieve and not only "how". The development of new goals on the basis of the interaction with the environment is here defined the "what" problem. The goals can be totally or partially dependent on a limited set of phylogenetically determined goals. However, the new ones are the result of the interaction with the environment and can overcome the original ones.

A goal is an event that is more likely to happen again because of the structure of an agent. This definition is equally applicable both to artificial and natural agents.

Usually artificial systems are designed with an already fixed set of goals that has to be reached. Therefore, designers focus their efforts to find "how" those goals can be achieved. One of the most sophisticated way to do that is learning. "Learning is usually defined as a modification in our behavior: a modification driven by a goal" [1]. The various learning paradigms focus mostly on this modification of the behavior.

In the neural network field, Supervised and Unsupervised Learning and Reinforcement Learning are examples of the learning paradigms. For instance, Sutton and Barto claim that "reinforcement learning is learning what to do – how to map situations to actions – so as to maximize a numerical reward signal [the goal]" [2], p.3. In the following they claim that "the basic idea is simply to capture the most important aspects of the real problem facing a learning agent interacting with its environment to achieve a goal. […] All reinforcement learning agents have explicit goals" [2], p.4-5. In other words, Reinforcement Learning deals with situations in which the agent seeks "how" to achieve a goal despite uncertainty about its environment: a goal fixed at design time. In this paper I want to address the process by which "what" (the goal) the agent has to achieve is generated. This is the "what" problem.

Using the Reinforcement Learning terminology, the "what" problem is equivalent to looking for new reward functions. In Reinforcement Learning systems "the reward function must necessarily be unalterable by the agent" [2], p.8, that is the goal is fixed. On the contrary, many biological systems are capable of developing partially or totally unpredictable goals. We propose to design systems that develop their own reward functions.

To develop new goals is important since the environment cannot be completely predicted at design time. Therefore many adaptive systems are be able not only to modify their behavior in order to perform optimally on the basis of some fixed criteria, but they are able to add new goals.

In nature there are basically two situations in which new goals are added: phylogenesis and ontogenesis. During phylogenesis subsequent generations of individuals develop new goals that arise from the emergence of new ecological balances. During ontogenesis each individual is capable of developing new goals.

In this paper, the topic of the generation of the goals of the agent is addressed. One of the main issues in Artificial Intelligence has been to search for optimal solutions to problems whose goals were already defined. This produced impressive results. However, the other side of the problem must not be underestimated: the generation of goals.

II. A TAXONOMY OF AGENT-ARCHITECTURES

Not all the goals of biological systems are fixed at birth: they only possess a limited, survival driven, built-in set of goals. As they grow and develop, biological systems continuously generate new goals on the basis of two separate factors: their genetic background and their past experience. Both are necessary in order to generate a particular goal.

In behaviour-based artificial structures, goals are defined elsewhere at design time [3, 4] but the behaviour changes according to the interaction with the environment. We refer here only to reactive systems or "traditional behavior-based"..In complex biological systems, goals are not fixed. Goals are the result of the interaction between experience and a limited number of innate instincts (the ones provided by genes). In many biological systems, it is possible to distinguish between phylogenetic aspects and ontogenetic ones, nature versus nurture [5-7].

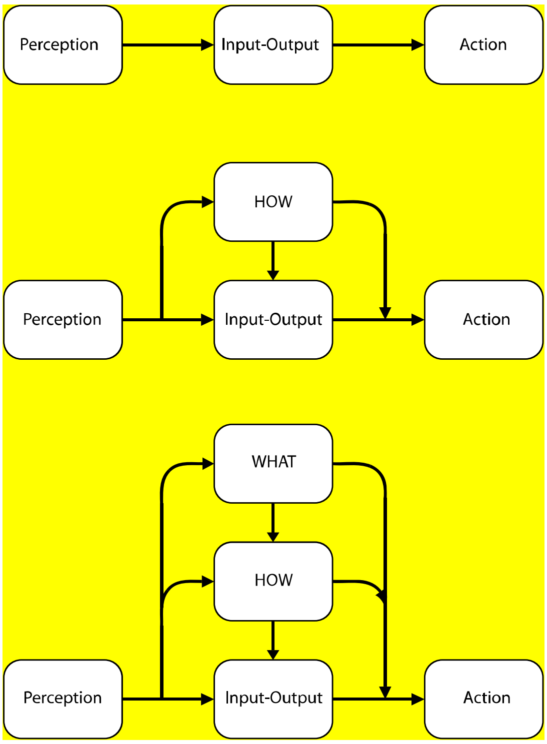

Fig. 1 Three possible architectures in relation to the environment.

A tentative set of different architectures based on the "what" problem can be devised. The taxonomy of these architectures is the following: a fixed control architecture, a "how generating" architecture and a "what generating" architecture.

In the first case, the system has no capability of modifying how it does what it does. There is an input-output mapping module, which takes 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, machine automata, and deliberative robots .

In the second case, the system is capable of modifying its behavior to fulfill an a priori target. The system is capable of modifying how it behaves. The input-output mapping module is flanked by a module capable of generating how to accomplish a goal. The "how" module modifies the a priori rules contained in the input-output mapping 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, the system is capable of modifying not only how it does what it does, but also to define what it does. The "what" module sets the goals that have to be pursued by the "how" module.

On the basis of these three levels a taxonomy of three kinds of architectures is proposed (Fig. 1): fixed control architectures, learning architectures, goal generating architectures.

A. *Fixed control architecture*

In this case, the causal structure of the system is fixed. There is no ontogenesis whatsoever. Notwithstanding the behavioral 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 modified by their experiences. Suitable examples of this category are Tolam's artificial sow bug [8], Braitenberg's thinking vehicles [9], Brooks' artificial insects [10, 11].

B. *"HOW" generating architecture*

A different level of structural dependency with the environment is provided by the architectures that generate how to perform a task. 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 goal. The goal 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) [2, 12]. In both cases some kind of learning is applied. These systems lack the capability of creating new goal. A behaviour-based robot generate how to navigate avoiding static and dynamic obstacles. However the goal behind this task is defined by the a priori design of the system. There are a few implementations of this kind of architectures: Babybot at LIRA-Lab, [13, 14], Cog at MIT [15, 16].

C. *"WHAT" generating architecture*

The last architecture of the taxonomy corresponds to a system that generates both how to perform a given task and what task must be performed. 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 goals 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 [17-20] as well as attempts to locate similar structures in the cortical architecture of humans [21]. 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 a priori 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.

III. TELEOLOGICAL UMWELT AND OPEN-TELEOLOGICAL SYSTEMS

All the previous systems react to external stimuli; they extract the elements necessary for their actions from the environment. The sensory capability of a system defines an accessible environment made of all those events the system is sensitive.

Given an accessible environment, the system reacts only to a limited subset of it: the perceivable environment. For instance a system reacts to configurations of visual stimuli like faces, or characters, or gestures. In general an artificial system is designed to react only to a few of them. Biological systems show the same kind of limitations. The Austrian biologist Von Uexküll, called the perceivable environment the Umwelt. Each animal lives, according to this biologist, in what he defined its Umwelt. His ideas derived from his work in the field of zoology where it is possible to observe that, given the same environment, two different specimens of two different species occupy the same physical space but have a completely different experience of the same physical world [22, 23]. The same rationale can be applied to artificial agents [24]. For instance, a Grey Walter's vehicle has a limited Umwelt constituted by the presence or the absence of light: two elementary events to which the vehicle reacts.

Finally the agent could be able to use some of the events of its environment as goals. Some events belonging to the perceived environment can be further selected as the goals of the system. The system chooses some events as goal to be pursued. In another words, some of this event will be selected by the system in order to be repeated.

With respect to a given agent there are four different kind of environment: 1) the environment in itself made of all possible physical events in the spatial and temporal neighborhood of the agent; 2) the part of environment which is available to the agent's sensorial apparatus; 3) the part of the environment which the agent is capable of perceiving; 4) the part of the environment which is a goal for the agent. The latter is what we define here as a teleological Umwelt: that is the set of all those events that are used as goals by a given agent.

Both the Umwelt and the teleological Umwelt can be fixed or expandable. If the Umwelt is fixed, the agent in unable to learn; otherwise the agent is capable of learning. If the teleological Umwelt is fixed, the agent's goals are hardwired; otherwise the agent is able to produce new goals and is here called a teleologically open system or agent.

Teleologically open systems are systems capable of dealing with the "what" problem defined above and the objective of the rest of the paper is to sketch their structure.

The previous description is here formalized (Fig. 2). The environment is defined as the set of all possible events in the surrounding of a given agent (World 1). The sensory and motor capabilities define a second subset of the first one. Only those events that can produce an effect (directly or indirectly) in the agent by means of its body structure are part of the second subset (World 2). However this does not entail that the agent is causally related with all of them. This would the same as claiming that just because someone has a pair of ears s/he should be able to have a reaction for all possible different languages. This is not the case. Only a limited number of them (World 3 or Umwelt) have been selected during the past phylogenetic and ontogenetic history of the agent and only those constitutes the relevant environment for it. Finally the most limited set of all is made only of those events which have been selected as goals by the agent (World 4 or teleological Umwelt).

With regard to the previous taxonomy of agents there is the following relation with the defined worlds. An agent with a fixed control architecture will have a fixed World 3 and a fixed World 4. An agent with a "how" generating architecture has an expandable World 3 but a fixed World 4. Finally, an agent with a "what" generating architecture has an expandable World 3 and an expandable World 4.

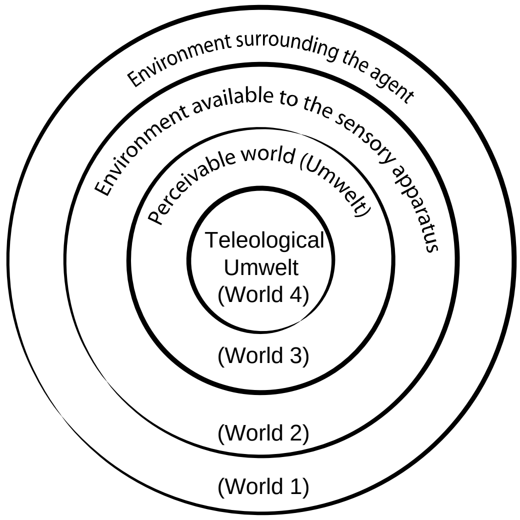

Fig. 2. The environment in relation to the different kind of agents.

An event belongs to the teleological Umwelt if the system must is able to use that event as a goal. This is the crucial difference. When something happens, it must not only produce some effect in the agent, but the effect it produces must be usable as a goal for the system as a whole.

A final consideration regards the role of ontogenesis. As we have mentioned at the beginning, it is crucial that the agent has the capability of expanding world 3 and 4 during its interaction with the environment. This entails learning and development: in short ontogenesis. Natural beings are capable of this kind of change. Artificial agents could.

IV. An ontogenetic goal generating architecture

The architecture we present in this paragraph (Fig. 3) is a potential candidate to endorse the ontogenetic development of new goals.

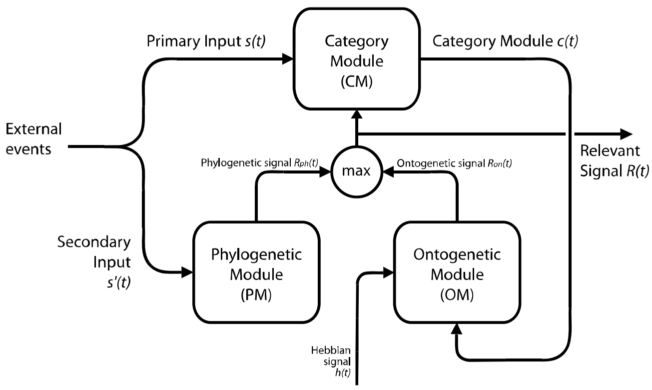

Fig. 3 The structure of an Ontogenetic Goal Generator.

We identified a module with general capacity of adaptability: a module which is not devoted to any specific sensory modality. Since this module aims at the generation of new goals for the agent during ontogenetic development, the module has been called the Ontogenetic Goal Generator.

Each module will receive data from at least one sensory modality. In reality the module could receive data from many different sensory modalities in parallel. In general we could say that the module is receiving two signals: a vector signal and a scalar signal. Similarly the module is producing two signals: a vector signal and a scalar signal.

A. Phylogenetic functions

Every real system doesn't have necessarily to start from scratch. Some information can be embedded in the system in such a way as to permit the system to bootstrap itself and to take advantage more quickly of the environment. However the specificity of this information is also its biggest disadvantage. Small or large differences in the environment as well as in the sensory-motor apparatus of the agent con completely destroy its utility. One way to overcome this nuisance is to implement some kind of meta-functions: a function that in a given environment will provide a specific function. In nature, an example of this meta-function is provided by imprinting. The bird has no specific knowledge of the shape of its mother. However it has some kind of meta-function (look for a moving object, in a fixed window frame of time). In this way the meta-function will be able to build a new function.

Since these functions must be ready before the development begins we called them 'phylogenetic' since in biological beings this kind of information is collected during phylogenesis. They could have been as well been called 'hard-wired' functions.

Phylogenetic functions serve as ready-to-use repository of goals. They tell to the agent what is relevant and what is not. Phylogenetic functions are extremely important. Since the potential complexity of the environment is unlimited, it makes sense to have some kind of mechanism in order to reduce the span of events to categorize and to explore. This mechanism corresponds to the phylogenetic functions. For instance, face recognition is triggered in human by some phylogenetic functions which are more interested to those events which are roughly similar to a face (or that are related to facial expression, emotion expression and social communication). In an agent there are n phylogenetic functions. Each function is normalized and it provides a value of 1 for those patterns that correspond to the category and a smaller value for all the other patterns. Of course it would be nice if the phylogenetic function would show some kind of graceful behavior.

B. Ontogenetic functions

If an agent would have only phylogenetic functions it would be an unsupervised learning or reinforcement learning system. It will have a set of goals fixed at design time and it will try to achieve them with different degrees of success. However it will be a very limited capacity for development: World 4 would be a priori fixed. The ontogenetic functions are completely equal to the phylogenetic ones. The only difference is that they are based on the categories developed during ontogenetic development.

By using the categories developed during development the system is partially unpredictable. It is partially predictable because given a certain environment and certain phylogenetic functions the agent will be able to pick up only certain events. However, in practice, given an unconstrained environment, the development will be unconstrained as well.

C. Relevant Signal

A system must be able to associate to what happens a value proportional to how much that particular event is important on the basis of its past history. This is important since this signal will be subsequently used as a reinforce signal in order to implement the policy of action learning.

The relevant signal is the maximum among all the phylogenetic functions and the ontogenetic ones. In this way it represents how much the current input is representative of the past history of the system (phylogenetic and ontogenetic).

At the beginning this signal will only be the result of the phylogenetic functions but at time goes by the new values coming in from the ontogenetic functions will start to have their role.

D. Categories

They implement the world 3 out of the world 2. They represent those events that are relevant for the agent. All other events are de facto invisible to the agent. Their number is dependent on the agent internal storage capability and on the algorithm used to implement them. Their role is that of representing a set of possible pattern as belonging to a given categories. They can be Euclidean or threshold. Different solution will provide different behavior but the overall logic will be the same. The basic idea is that given a incoming flow of stimuli only a limited subclass of this stimuli will be accepted as belonging to a given category.

V. ROBOTIC IMPLEMENTATION

To test the architecture a robotic implementation was used in a very simple case. A series of different shapes associated with colours were presented to the robot. The system is equipped with a phylogenetic goal 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 goals 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 goal (by directing its gaze towards the stimulus) that was not envisaged at design time and that is the result of the ontogenetic development [25, 26]. A robotic head with two degrees of freedom has been adopted as robotic set-up.

A. Sensor Module

The robotic head was equipped with a videocamera capable of acquiring logpolar images [27, 28].

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.

B. 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

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

where r is the random variable for the amplitude. 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.

C. 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 $d(\vec{v},\vec{w})$: $d:(\Re^n\times\Re^n)\mapsto\Re$, 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. 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*.

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

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.

D. 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:

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

where R_{ph} is the Relevant Signal; *Saturation*(η, ξ) is the colour saturation at the pixel (η, ξ) in log polar coordinates, and *N* is 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.

E. 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 Figure 6a on the left. The frequency with which the system was looking at different points was measured and displayed in Figure 6a on the right. 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 (Figure6b on the left). As is shown in Figure 6b on the right, the head spent more time on the coloured shapes instead on the white background because of the phylogenetically implanted rule.

Finally we presented again the initial stimulus (the set of colourless shapes). This time, the system spent more time on the colourless shapes than on the background (Figure 6c on the right). The behaviour of the system changed since the system added a new goal (the star) to the previous one (saturated colours).

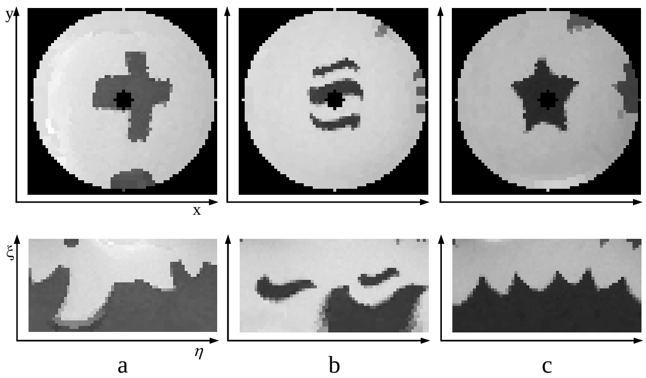

Fig. 4. The Cartesian (upper row) and log-polar (lower row) images for a cross a), a wave b), and a star c).

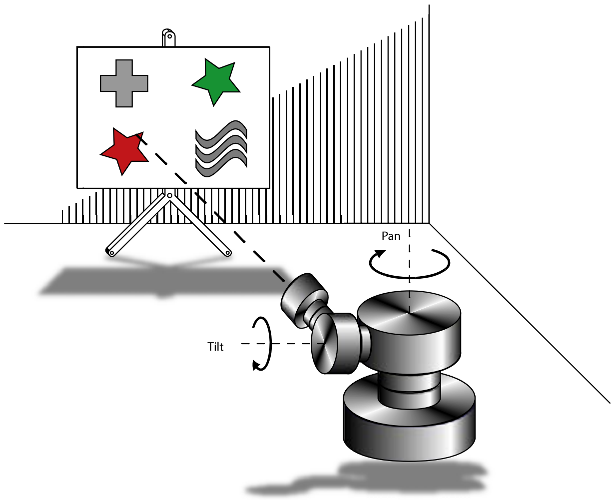

Fig. 5 Sensory and motor set-up.

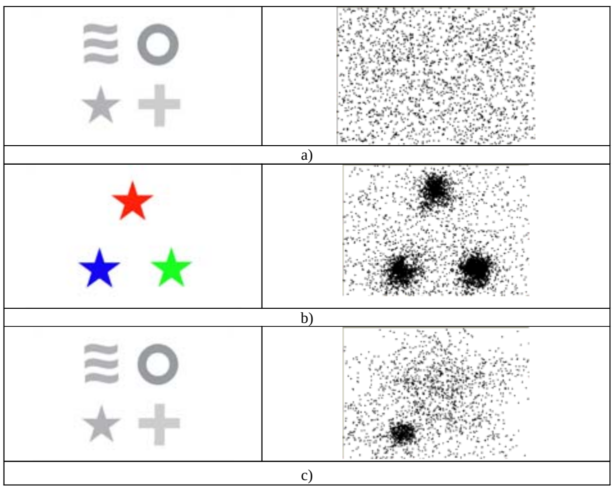

Fig. 6. Experimental results

References

[1] C. R. Gallistel, The Organization of Learning. Cambridge (Mass): MIT Press, 1990.

[2] R. S. Sutton and A. G. Barto, Reinforcement Learning. Cambridge (Mass): MIT Press, 1998.

[3] D. McFarland and T. Bosser, Intelligent Behavior in Animals and Robots. Cambridge (Mass): MIT Press, 1993.

[4] R. C. Arkin, Behavior-Based Robotics. Cambridge (Mass): MIT Press, 1999.

[5] S. J. Gould, Ontogeny and Phylogeny. Cambridge (Mass): Harvard University Press, 1977.

[6] J. L. Elman, E. A. Bates, M. H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett, Rethinking Innateness: A Connectionist Perspective on Development. Cambridge (Mass): MIT Press, 2001.

[7] S. Oyama, The Ontogeny of Information. Durham: Duke University Press, 1985/2000.

[8] E. C. Tolman, "Prediction of the vicarious trial and error by means of the schematic sowbug," Psychological Review, vol. 23, pp. 318-336, 1939.

[9] V. Braitenberg, Vehicles: Experiments in Synthetic Psychology. Cambridge (Mass): MIT Press, 1984.

[10] R. A. Brooks, "Intelligence Without Representations," Artificial Intelligence Journal, vol. 47, pp. 139-159, 1991.

[11] R. A. Brooks, "New Approaches to Robotics," Science, vol. 253, pp. 1227-1232, 1991.

[12] M. Kuperestein, "INFANT Neural Controller for Adaptive Sensory-Motor Coordination," Neural Networks, vol. 4, pp. 131-145, 1991.

[13] G. Metta, G. Sandini, and J. Konczak, "A developmental approach to visually guided reaching in artificial systems," Neural Networks, vol. 12, pp. 1413-1427, 1999.

[14] R. Manzotti, A. Gasteratos, G. Metta, and G. Sandini, "Disparity estimation in log polar images and vergence control," Computer Vision and Image Understanding, vol. 83, pp. 97-117, 2001.

[15] C. Ferrell and C. Kemp, "An Ontogenetic Perspective to Scaling Sensorimotor Intelligence.," presented at 1996 AAAI Fall Symposium on Embodied Cognition and Action, 1996.

[16] R. A. Brooks, C. Breazeal, M. Marjanovic, B. Scassellati, and M. Williamson, "The Cog Project: Building a Humanoid Robot," in Computation for Metaphors, Analogy, and Agents, vol. 1562, Springer-Verlag Lecture Notes in Computer Science, Nehaniv, Ed. Berlin: Springer-Verlag, 1999.

[17] D. Gachet, M. A. Salichs, L. Moreno, and J. R. Pimentel, "A software architecture for behavioural control strategies of autonomous systems," presented at IECON'92, San Diego (Cal), 1992.

[18] J. J. Weng, "Cresceptron and SHOSLIF: Toward comprehensive visual learning," in Early visual learning, S. K. Nayar and T. Poggio, Eds. New York: Oxford University Press, 1996.

[19] K. Fukushima, S. Miyake, and T. Ito, "Neocognitron: A Neural Network Model for a Mechanism of Visual Pattern Recognition," IEEE Transactions on Systems, Man, and Cybernetics, vol. 13, pp. 826-834, 1983.

[20] K. Fukushima, M. Okada, and K. Hiroshige, "Neocognitron with Dual C-Cell Layers," Neural Networks, vol. 7, pp. 41-47, 1994.

[21] T. Togawa and K. Otsuka, "A model for Cortical Neural Network Structure," Biocybernetics and Biomedical Engineering, vol. 20, pp. 5-20, 2000.

[22] J. v. Uexküll, A Stroll through the Worlds on Animals and Men: A Picture Book of Invisible Worlds. New York: International University Press, 1934.

[23] A. Clark, Being there: putting brain, body and world together again. Cambridge (Mass): MIT Press, 1997.

[24] C. Emmeche, "Does a robot have an Umwelt? Reflections on the qualitative biosemiotics of Jakob von Uexküll," Semiotica, vol. 1, pp. 653-693, 2001.

[25] R. Manzotti, "A process based architecture for an artificial conscious being," in Process theories, vol. Process Theories: Crossdisciplinary studies in dynamic categories, J. Seibt, Ed.: Kluwer Academic Press, 2003, pp. 285-312.

[26] R. Manzotti and V. Tagliasco, "From "behaviour-based" robots to "motivations-based" robots," Robotics and Autonomous Systems, vol. 51 (2-3), 31 May, 2005, pp. 165-190.

[27] G. Sandini, P. Questa, D. Scheffer, and A. Mannucci, "A Retina-like CMOS Sensor and its Applications," presented at SAM-2000, Cambridge, USA, 2000.

[28] G. Sandini and V. Tagliasco, "An Anthropomorphic Retina-like Structure for Scene Analysis," Computer Vision Graphics and Image Processing, vol. 14, pp. 365-372, 1980.

[29] P. Perruchet and A. Vinter, "The Self-Organizing Consciousness," Behavioral and Brain Sciences, 2002.

[30] I. P. Pavlov, Selected works. Honolulu, (Haway): University Press of the Pacific, 1955/2001.