Development of categorization abilities in evolving embodied agents : a study of internal representations with external social inputs.

Francesco Pugliese
Natural and Artificial Cognition Laboratory, University of Naples, Via Porta di Massa 1, Naples, Italy - francesco.pugliese@unina.it

Abstract . The present paper investigates how embodied and situated agents perform tasks that require skills of categorization. The agents are artificially created by an adaptation process. The task is to categorize different shapes of objects using sensory-motor and linguistic input. Results show that the autonomous agents are able to solve the categorization task by integrating the sensory-motor experienced states and employing “linguistic” input from the environment. This shows that autonomous agents are able to develop some "emerging" abilities by exploiting the information present in the environment in order to recognize and discriminate objects. Autonomous agents also exhibit a "social" behavior, because they are able to categorize the objects in the environment, even when external inputs are unavailable. The purpose of this work is to prove the theoretical hypothesis that the “social” information (external labels), deriving from another agent or from the trainer, facilitates individual capacity to categorize, by the creation of internal representations.

Keywords: behavioral categorization, categorical perception, active vision system, internal representations, embodied and situated agents,cognition

1 Introduction

Categorization is the ability to discriminate different environmental situations (and different states of interaction between the agent and the environment) by producing different behaviors [1,2]. The way in which categorization may occur, depends on the interaction between the control system of the agent, his body and the environment where it is situated. The categorization may emerge by actively exploiting sensory-motor experience and by recording it in some internal states of the agent’s control system.

Generally, the process of categorization is performed individually by the agents, but in some cases it may be obtained socially, for example, with the help of a form of communication that can improve the process of categorization itself, by providing to the agent the capacity of making use of other individuals' information.

It is possible to classify the categorization ability into : a) active categorization, which is the ability of an agent to categorize, by using the actions in order to affect its own sensory state. For example, a robot with a camera equipped with zoom, mounted on a moving head, would be able to "manipulate" what it sees, simply by moving his head; b) passive categorization, which is the opposite of active categorization, since the sensory state is less affected by the actions, because, for example, the robot is stationary or it has a limited number of degrees of freedom, and it can not move in directions that allow the robot to converge towards an optimum perception. An example of passive categorization, could be a robot with a fixed camera, which must solve the task of recognize a particular shape of objects inside an image. It can not move and it is not able to alter the image perceived by the optical sensors of the camera.

Another classification of categorization ability is based on the ways in which the categorization might emerge : a) behavioral categorization, which is the ability to produce different behaviors in order to discriminate different environmental situations, and usually it is the result of a sequence of interactions between the control system with the agent's body and the environment. For example, in an environment with edible and not edible food, an agent should be able to reach the edible one, but stay away from not-edible one. b) categorical perception, which is the agent's ability to produce different behaviors with respect to different environmental conditions, using the "memory" of internal states. In this case, the control system of the agent maps groups of sensory stimuli belonging to a category into the same internal states.

The behavioral categorization can emerge from a careful and parsimonious use of local properties present during the interaction between the body and the environment. For this reason b) and c) may occur purely reactive agents. Reactive behavior is defined as behavior that is configured by a direct association between the sensory input of the robot and its actuators, that is without the help of representations or internal states. A study is made on behavioral categorization by implementing an experimental setup [1] where a robot is placed in a circular room whose perimeter is marked by separating it into 40 cells. The task of the robot is to go and remain in the left half of the environment, starting from any point on the circumference. The control system consists of a neural network with 20 sensory neurons, that encode the number of cell where the robot is located. The neural network has no internal neurons. By evolving genetically the robot, with a fitness to compute the number of cycles (of life) in which the robot remains in the left part of the environment, a solution is found that only exploits the sensory-motor coordination : the robot reacts the same way, at the same sensory pattern since the evolutionary process selects certain "behavioral attractors". In this case attractors are represented by pairs of adjacent cells present in the left semicircle. The robot, by reacting rapidly to this attractors, continuously moves clockwise and counterclockwise. Finally, robot's fluctuations in correspondence of attractors produce the robot is stationary on the left side of the environment.

The categorical perception is required when the autonomous agent needs internal states for mapping a sequence of stimuli received in the relevant category: for this reason, categorical perception requires an ability to discover the categories and relations between the different categories and behaviors that should be exhibited when the event, related to the category, happens. The opportunity to develop skills in perceptual categorization is investigated into a series of experiments that simulate the operation of two e-puck robots [3]. The robots are placed in a square bounded by walls. Two target areas are located inside the square and each one has a different color. The purpose of the robots is to find and stay in the same target area, together. The control system of the robot consists of a neural network with 14 sensory neurons, which encode the state of activation of 8 infrared sensors, 1 ground sensor (which is activated when the robot is on a target area), 4 directional communication sensors and an input-output communication connection. Moreover there are two hidden neurons and 2 output neurons that control each one of the two engines, and 1 output neuron that controls the emitter of the communication signal. During the evolutionary process, it is calculated the reward that individuals obtain when remaining together in the same target area, by means of a group fitness function. The reward is the sum of 0.l whenever a robot is positioned within a target area alone, and the score 0.5, if both robots are placed together. At the end of the evolutionary process it is observed that the robots can solve the task by developing some strategies based on a rudimentary form of communication. During the first generation, the robots develop some remarkable exploratory skills consisting of attempts to avoid the walls of the environment. Then, each robot develop more purely social skills such as emitting signals that have the effect of altering the trajectory of other robots. For example, if both robots stay in the same target area, they begin to produce a different sequence of signals that integrate both information from the infrared sensors and from communications, with the ultimate effect of producing the two stationary robots in the same target area. So, communicative behaviors allow two robots to "feel" when they are outside the target area, and to approach and remain relatively close.

A special case of categorical perception is Active Categorical Perception which may arise when the use of internal states is not enough to discover the regularities needed to discriminate between different categories of stimuli, because regularities are hidden or absent. In fact, in these situations, the sensory stimuli experienced by an agent are co-determined by the action performed by the agent itself. Active categorical perception are recently analyzed by an experiment [4] where a simulated anthropomorphic robotic arm, equipped with tactile sensors, is required to perceptually categorize spherical and ellipsoid objects. The anthropomorphic robotic arm is equipped with proprioceptive and tactile sensors distributed on the external part of the arm. The whole arm mainly consists of three elements: the arm, the forearm, and the wrist. The arm's controller consists of a continuous time recurrent non-linear network (CTRNN) with 22 sensory neurons, 8 internal neurons, and 18 motor neurons. The neural controller has been trained by an evolutionary process in which the free parameters of dynamical neural networks are varied randomly and in which variations are retained or discarded on the basis on their effects on the overall ability of the robots to carry out their task. Results indicate that robots are capable of developing an ability to effectively categorize the shape of the objects despite the high similarities between the two types of objects, the difficulty of effectively controlling the arm, and the need to reduce the effects produced by gravity, inertia, collisions, etc. More specifically, the best individuals display an ability to correctly categorize the objects located in different positions and orientations, as well as an ability to generalize their skill to objects positions and orientations never experienced during evolution.

2 Materials and methods

The purpose of this experimental setup is to evolve an agent in order to solve a non-trivial categorization problem where an embodied and situated agent can discriminate objects of different shape and size. The agent is provided with simple sensory-motor system and the capability of integrating information, over the time. In particular, the intent of the experiment is to understand how external labels can facilitate agent's categorization abilities. In this way, when the robot discriminate between two or more pictures, it receive a "linguistic" external input, corresponding to the correct category.

2.1 Experimental setup n.1

A simulated robot is placed into an environment whose ground contains a geometric picture which is drawn on. Environment is essentially a square arena of 80 x 80 cm, surrounded by walls. In order to make the task more complex (respect to the classical problems of the literature), noise presence is simulated on ground sensors. Noise is randomly distributed on the arena floor with a probability of 30%, and "granularity" of noise with 4:1 ratio (each noise pixel is made of 4 screen pixels). Conventionally, only two kinds of pictures can be showed on the environment ground (theoretically they might be more than two). In particular, those shapes are an equilateral triangle and a square.

The robot has a circular body with a radius of 11 cm and it's equipped with two motors which control the movements of two wheels, respectively. Moreover, the robot is provided with eight infrared sensors placed around the body perimeter. Infrared sensors are able to detect the presence of obstacles up to a distance of 15 cm. On the lower surface of the robot chassis (body) are installed eight ground sensors, which are able to detect color tones of the floor. Ground sensors return a value normalized between 0 and 1 according to the tone detected : so, if one of the ground sensors is situated on the internal area of the geometric picture, then the sensor will return a value proximal to 1. Otherwise, if a ground sensor is placed outside the picture, then it will return 0 (since the arena floor color is white). Finally, the robot is equipped with two more sensors (radio receivers) that receive "linguistic" inputs from the outside environment. The robot returns a categorization output pattern (a specific actuator) which is the category selected by the robot.

The robot's control system consists of a neural network with recurrent hidden neurons. The neural network is made of 24 neurons that are organized into 3 layers : input, output and hidden. Input layer contains 16 neurons that encode the activation state of the corresponding input sensors, which are 8 neurons that receive the reading of infrared sensors and 8 neurons that receive the reading from the ground sensors. Furthermore, 1 neuron receives at time t, the activation of the categorization output at time t-1. Internal layer consists of 2 hidden neurons which are connected with each other by 2 recurrent connections. The encoding of internal neurons is the "leaky" mode. Output layer is made of 2 neurons which control respectively the speed of two motors, and 1 neuron which encode the categorization output. The categorization output is binary encoded : if the neuron returns a value less than 0.5 (and more than 0.0) then it's representing the "square" category, otherwise, if it returns a value more than 0.5 (and less than 1.0) then it's representing a "triangle". Neural network topology (see Figure 1) is structured into complete connections between input layer and output layer : there are complete direct connections from input to hidden and from hidden to output, then 1 recurrent connection join categorization output (at t-1 time) to input.

The free parameters of the robot's neural controller, i.e. synaptic weights of the connections, biases, and time constants, are encoded inside the genotypes of the individuals of the genetic algorithm. Connection weights are encoded into 8 bit strings, and then they are normalized into [-5.0, +5.0] interval. Instead, time constants are normalized into the range [-1.0,1.0]. The initial population consists of 100 randomly generated genotypes (individuals) which represent synaptic weights of the connection, biases and time constants of 100 corresponding neural networks. At the start of the generation, each genotype is converted into the numeric description of the respective neural network. Then, the numeric string associated to a neural network is recorded into the robot's control system, which perform his actions (using the current genotype) during the entire life (2000 cycles). The 20 best genotypes (with respect to the fitness) are allowed to reproduce by generating an offspring of 5 individuals (totally, 20 x 5 = 100 individuals which represent next generation) with a mutation probability of 2% (2% of the bits are replaced randomly by a new value). The whole evolutionary process takes 300 generations.

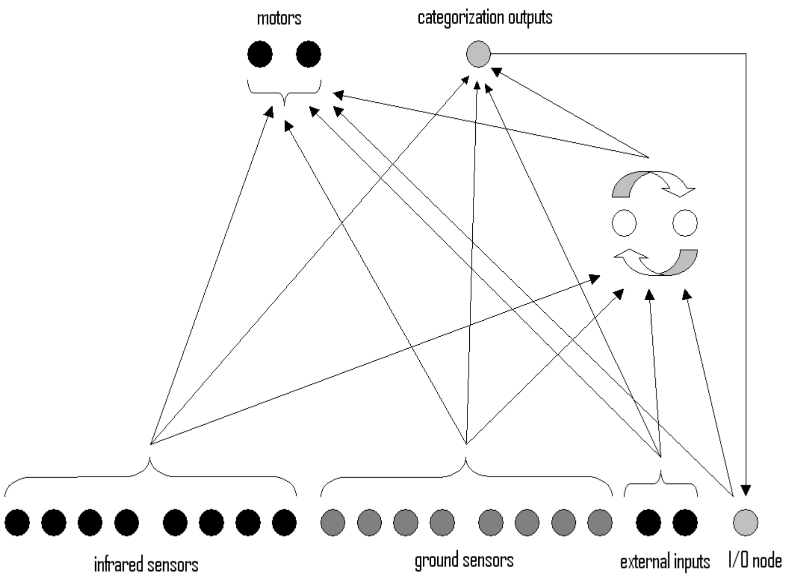

Fig. 1. The figure shows the architecture of robot's neural controller.

During the generations, each individual is evaluated for 20 trials lasting 2000 cycles, with each cycle lasting 100ms. At the start of each epoch (trial) robot moves from a random position and orientation inside the environment but it always starts outside the target area. Position and size of pictures are selected randomly. The experiment is replicated 10 times by using different initial populations.

The fitness function of the genetic algorithm is calculated by performing the sum (for each life cycle) of +1.0, if the robot is placed on a target area and if it is categorizing the shape correctly (according the categorization output signaling). Otherwise, the robot receives -1.0 if it is placed on a target area and it is not categorizing the shape correctly.

2.2 Experimental setup n.2

This experimental setup is inspired by previous experiments [5,6] in order to test the hypothesis of this work by using another paradigm. The experimental setup n.2 is the same as n.1, except that it used an evolutionary active vision system (simulated in this setup) consisting of a simulated "fovea" (a square visual area), instead of the robot. The fovea is able to move only horizontally or vertically with respect to the environment. For each cycle, the retina can not move more than 3 pixels.

The fovea's control system consists of a neural network with recurrent hidden neurons. The neural network (see Figure 2) is made of 27 neurons. The input layer contains 21 neurons that encode the activation state of the corresponding input sensors, which are 16 image neurons, 2 external inputs, two of which receives fovea position, and 1 I/O node. The internal layer consists of 3 leaky recurrent hidden neurons. Finally, the output layer consists of 2 neurons which define fovea movements, and 1 categorization output.

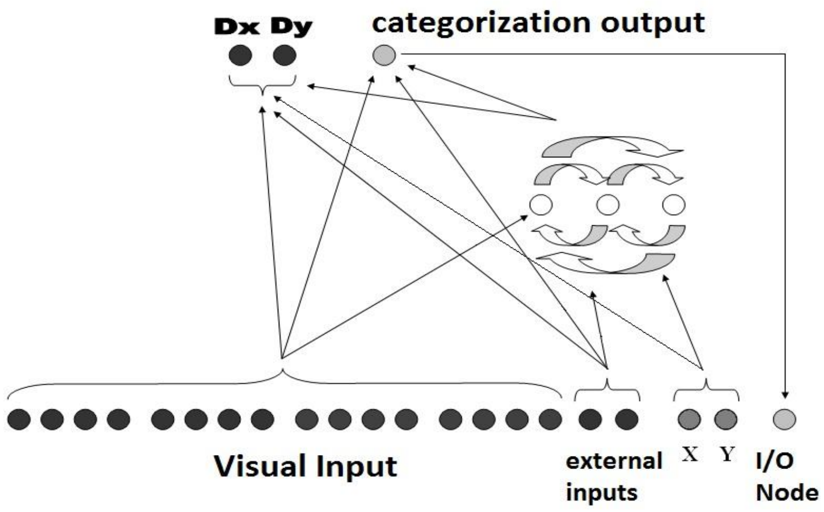

Fig. 2. The figure shows the architecture of fovea's neural controller.

3 Results

3.1 Results of experimental setup n.1

By evolving the experiment n.1 with 20 replications, it can be observed that the robot seems to solve the categorization task correctly, by exploiting the external linguistic input. Initially, the categorization output, returned by the robot, does not have a constant trend, but it's variable over the time. The output takes the right configuration only when the robot is on the figure and it receives the right social input (on the 1000^{th} cycle). At the end of evolution, the emerging strategy, to solve the task, is the same in all the best replications. Initially, the evolved robot exhibits a purely exploratory behavior, by jumping from one wall to another one and drawing curvilinear trajectories, in the environment. This first strategy allows the robot to improve the chances of identifying the target area position, since it's more likely to be located inside the environment.

When the target area is detected, the robot exhibits a "line following" behavior, since it moves following the perimeter of the picture, in order to find some critical features of the geometric shape (see Figure 4). Then, when the robot seems positioned on a feature of the picture, such as a corner, it exploits an "active perceptual categorization" in order to distinguish the corners and so recognizing the picture shape. Essentially, when the robot reaches the corner, it performs a "measurement" of the angle by exploiting the integration of information from ground sensors with representations evolved into internal states. Depending on the category that the robot establishes, it emits a categorization output that is always close to the value 0.0 for a square and it is close to 1.0 for a triangle. When the external inputs arrive (from 1000th), the categorization is reinforced in the direction of the right categorization. This fact shows that the robot (from the 1000 cycle) uses primarily external inputs in order to establish the right category.

As already stated, the purpose of the experiments is to investigate the possibility that the receiving of external labels (from the outside environment), designating categories, can influence the ability to categorize individually. For this reason, the experimental setup n.1 has been also evolved again, with 20 replications, in the condition of total absence of external labels.

It is observed that fitness curves of the experiment, evolved without external inputs, have values of fitness curves lower than the fitness of the experiment evolved with external inputs. The result is not surprising since the robot, in the original condition (with external inputs) have an advantage because it receives "help" on the right category.

In order to show the effect of external inputs on the individual categorization ability, the robots are tested into 3 different conditions for a performance measure : a) Robot evolved with external inputs and tested with external inputs; b) Robot evolved with external inputs and tested without external inputs; c) Robot evolved without external inputs and tested without external inputs. Performance curves are obtained by testing all 3 conditions (see Figure 4), by calculating the total number of correct categorization during the lifetime. Practically, the measure of performance is obtained by adding +1, during a cycle, if the robot is on a target area and it provide the right categorization output for that particular area's shape.

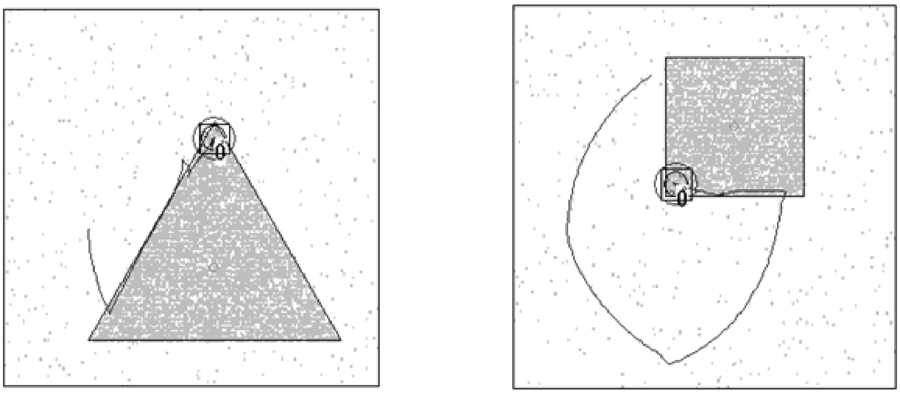

Fig. 3. The robot follows the perimeter of picture looking for parts of identification.

The performance measure is determined over 100 trials for each condition. Then the performance curves are sorted, so that the differences are better highlighted (see Figure 4).

Moreover, a statistical test is performed, an “heteroscedastic Student's T test two-tailed" (see Figure 5). The test is carried out on pairs among distributions of 2000 performance measures for each condition. The p-value of the T-tests, is less than a significance of 1% in all pairings between distributions.

Performance curves and statistics confirm the result already shown by fitness, i.e. the robots evolved with external inputs get higher performance respect to the robots evolved in condition of total absence of external labels. However, the performance curves and the statistics highlight a result more important than the fitness curves :

looking at the second and the third curve (in descending order) it is possible to see that the robot evolved in the presence of external labels (when it is tested without such external inputs) is able to categorize better than the robot evolved without external inputs. It seems that the robot is able, anyway, to "internalize" the social information which he had in the evolving phase, in order to obtain an advantage in terms of performance, in testing phase. This shows that social information available in learning-evolving phase can be exploited by an embolied and situated agents for improving own individual categorization ability, also when this social information is not available anymore. A key role of information acquisition could be played by the hidden internal nodes. To prove this, the previous experiment is replicated by eliminating hidden neurons. Performance curves of the new experiment (see Figure 6) show that, in the absence of hidden neurons, the robot evolved with external inputs and tested without them no longer has a performance advantage in comparison to the robot evolved without external inputs. This fact proves without any doubt the importance of internal hidden units for the formation of internal representations.

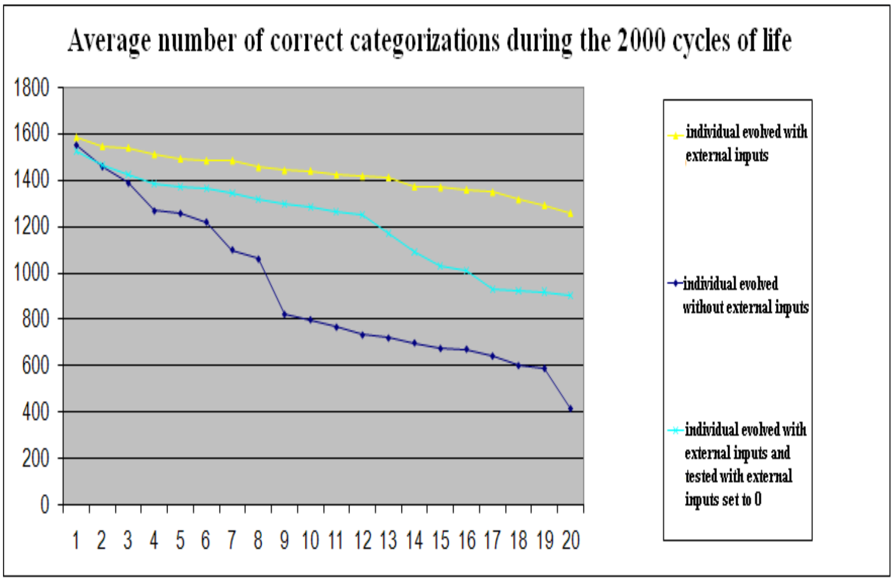

Fig. 4 . Performance curves sorted by number of categorizations in three conditions.

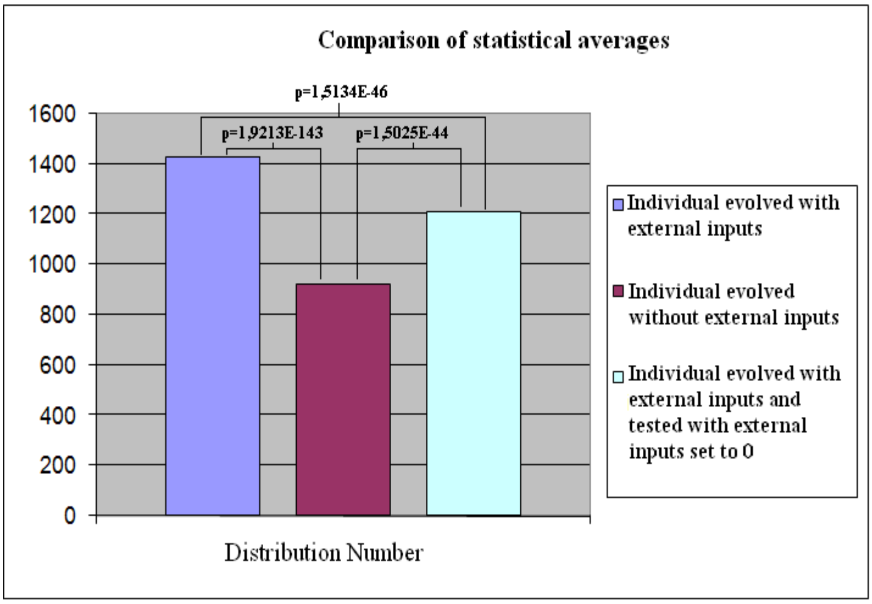

Fig. 5. Averages of the number of categorizations calculated in 20 seed x 100 epochs, for each condition. It 's also possible to observe the p of Student's t test two-tailed heteroscedastic.

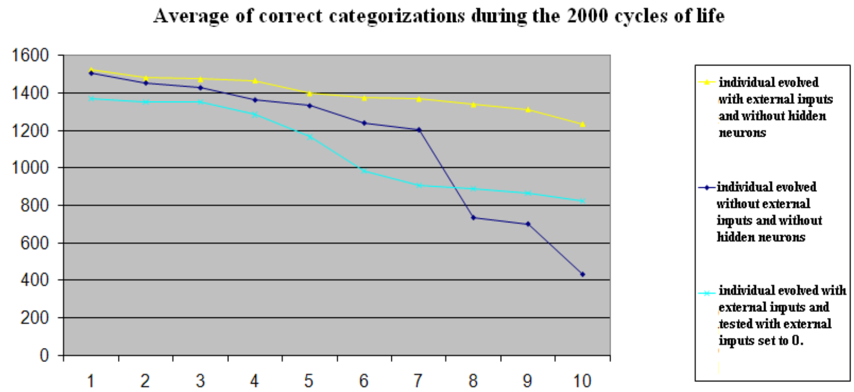

Fig. 6. Performance curves sorted in three conditions, without hidden neurons

The hidden neurons also play an essential role for the ability of integrating sensory-motor information with the social input. In order to better illustrate the role of hidden neurons in the definition of internal representations, the activity of the two hidden neurons, during a sample execution of experiment 1, is plotted (see Figure 7). Sampled data, for each replication, consist of 20,000 values (2000 activation values relating to the 2000 life cycles x 10) belonging to the interval [0.0, 1.0]. Activation data are captured in the three conditions : a) individual evolved and tested with external inputs, b) individual evolved without external inputs, c) individual evolved with external inputs and tested with external inputs set to 0. The x-axis shows the values of the first hidden node, H1, the y-axis shows the values of the second hidden node, i.e. H2. Looking at the charts it's possible to see that data in a) and c) are grouped in two clusters representing two different internal states associated with two different categories. In b) "evolved without external inputs" condition, the distinction of internal representations is less clear. In conclusion, the plots of internal neurons activities confirm the assumptions made previously, about the possibility that the external labels enable the neural network able to create well defined internal representations of categories, which can be exploited later, even when these labels are no longer available.

3.2 Results of experimental setup n.2

Applying the paradigm of experimental setup n.2, leads to obtain the same results as experimental setup n.1. However, the strategies used by fovea for categorization are more than setup n.1. The mobile fovea, starting from a random position on the visual area (environment), exhibit initially a purely exploratory behavior, since it is able to perceive only a very limited portion of the displayed image. In general, the research strategy of the picture is to move along a spiral path from the initial position, so it can maximize the probability of detecting the position of the geometric shape. . Once it is positioned on geometric shape, the retina (fovea) can use two different strategies to categorize the object: a) try to identify the degree of inclination of the side of the geometric shape, which in the case of the square is parallel to the horizontal or vertical axes of the visual field, while the triangle is inclined at an angle of 60 °; b) measure the angle of the corners of a geometric shape, which in the case of the square are right angles, while in the case of the triangle are acute angles (<90 °) (see Figure 8).

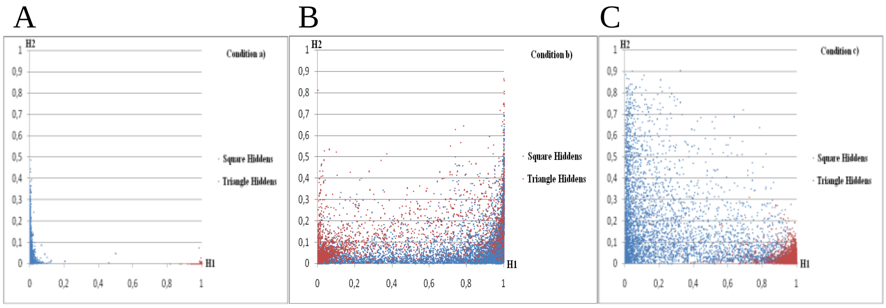

Fig. 7. Graph of activations of the two internal neurons in the 3 conditions.

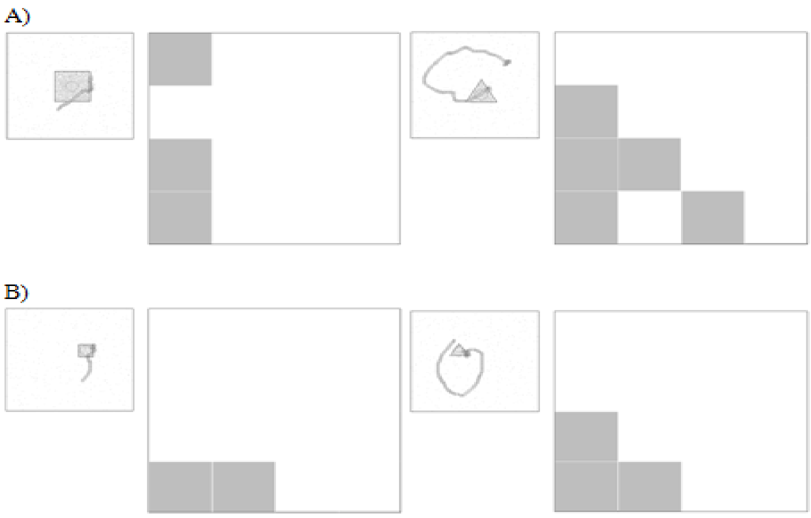

Fig. 8. Different strategies of pattern recognition of the mobile retina

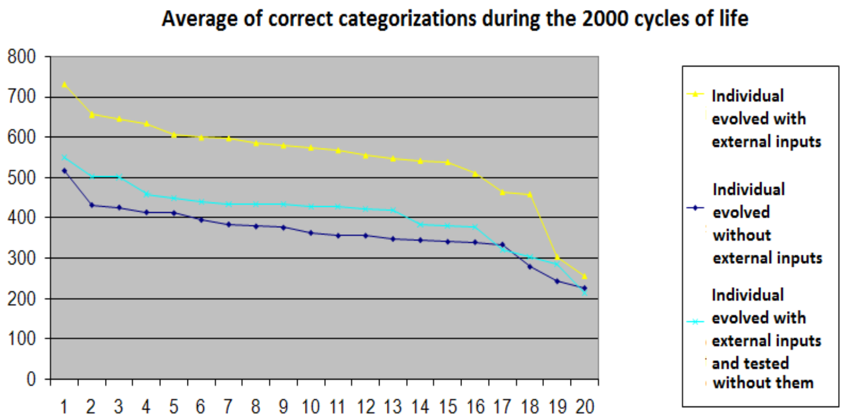

Fig. 9. Performance curves sorted in three conditions, of active vision system.

4 Conclusions

This work prove how a robot, evolved by genetic algorithms, is able to solve tasks that require complex cognitive skills like categorization. The robot was able to develop some emerging skills, such as the ability to exploit the information from the environment for recognizing objects of different shapes and sizes. Also, by using social information, the robot was able to exhibit a "social" behavior, learning to categorize the objects of the environment that surrounds it, even when social information is no longer available. This fact show rigorously an important theoretical result : the social information which is communicated from the external environment (by another individual or by training) allows an agent to solve the task of understanding the surrounding environment, in an easier way. This is true because external information allow creation of good internal representations better than the case in which these external information are not available.

Once these good internal representations are created, they allow the agent to achieve an advantage in single categorization even when the agent is deprived of such external information.

5 Acknowledgements

I would like to thank Stefano Nolfi, Davide Marocco and Marco Mirolli for the support they have given me to carry out this work.

6 Rereferences

1. Nolfi S. (2005a). Categories formation in self-organizing embodied agents. In H. Cohen & C. Lefebvre (Eds), Handbook of Categorization in Cognitive Science. Elsevier. pp. 869-889. S. Harnad, Ed., Categorical Perception: The Groundwork

2. Nolfi, S. (2005b). Behaviour as a complex adaptive system: On the role of self-organization in the development of individual and collective behaviour. ComplexUs, 2 (3-4), 195-203.

3. Marocco, D., Nolfi, S., Ikegami, T (2011).; Emergence of Protosentences in Artificial Communicating Systems, Autonomous Mental Development, IEEE Transactions on , vol.3, no.2, pp.146-153, June 2011 doi: 10.1109/TAMD.2011.2120608

4. Tuci E., Massera G., Nolfi S. (2009). Active categorical perception in an evolved anthropomorphic robotic arm. IEEE International Conference on Evolutionary Computation (CEC), special session on Evolutionary Robotics.

5. Kato, Toshifumi., Floreano, D., (2001) – An evolutionary active-vision system. In Proceedings of the congress on evolutionary computation (CEC01). Seoul, Korea, October 2001. IEEE Press. Piscataway, NJ.

6. Mirolli M., Ferrauto T., Nolfi S. (2009). Categorisation through Evidence Accumulation in an Active Vision System.