Aberystwyth University

Active categorical perception in an evolved anthropomorphic robotic arm

Tuci, Elio; Massera, Gianluca; Nolfi, Stefano

Published in:

Proc. of the IEEE International Conference on Evolutionary Computation (CEC), special session on Evolutionary Robotics

Publication date:

2009

Citation for published version (APA):

Tuci, E., Massera, G., & Nolfi, S. (2009). Active categorical perception in an evolved anthropomorphic robotic arm. In *Proc. of the IEEE International Conference on Evolutionary Computation (CEC), special session on Evolutionary Robotics* (pp. 978-1) <http://hdl.handle.net/2160/5410>

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

Active categorical perception in an evolved anthropomorphic robotic arm.

E. Tuci, G. Massera, S. Nolfi

Abstract—Active perception refers to a theoretical approach to the study of perception grounded on the idea that perceiving is a way of acting, rather than a cognitive process whereby the brain constructs an internal representation of the world. The operational principles of active perception can be effectively tested by building robot-based models in which the relationship between perceptual categories and the body-environment interactions can be experimentally manipulated. In this paper, we study the mechanisms of tactile perception in a task in which a neuro-controlled anthropomorphic robotic arm, equipped with coarse-grained tactile sensors, is required to perceptually discriminate between spherical and ellipsoid objects. The results of this work demonstrate that evolved continuous time non-linear neural controllers can bring forth strategies to allow the arm to effectively solve the discrimination task.

I. INTRODUCTION

An important consequence of being situated in an environment consists in the fact that the sensory stimuli experienced by a robot are co-determined by the action performed by the robot itself. That is, the actions and the behaviour exhibited by the robot later influence the stimuli sensed by the robot, their duration in time, and the sequence with which they are experienced. This implies that: (i) perception (i.e., the ability to categorize objects and events in the environment) is strongly influenced by action [1]; and (ii) sensory-motor coordination (i.e., the ability to act in order to sense stimuli or sequence of stimuli which enable and/or favour the ability of the robot to perform its task) is a crucial aspect of perception and more generally of situated intelligence [2], [3], [4], [5], [6].

Although the importance of the topic is now widely recognized, our understanding of how natural organisms perceive actively their environment is still limited to few specific cases (e.g., [7], [8]). Similarly, our ability to build artificial systems which are able to exploit sensory-motor coordination is still very limited. The first type of limitation can be explained by considering that experimental research rarely takes into account detailed data encoding how organisms interact with their environment over time. The second type of limitation can be explained by considering that, from the point of view of the designer of the robot, identifying the way in which the robot should interact with the environment in order to sense the favourable sensory states is extremely difficult. One promising approach, in this respect, is constituted by adaptive methods in which the robots are left free to determine how

they interact with environment (i.e., how they behave, in order to solve their task).

This paper illustrates how a further elaboration of adaptive methods proposed in related studies can be successfully applied to a significant more complex scenario [3], [4], [9], [10]. In particular, we demonstrate how a non-trivial problem which consists in perceptually categorizing objects with different shapes can be solved in an effective and robust way through an evolutionary adaptive method. With this method, free parameters (i.e., those that are modified during the adaptive process) encode features that regulate the fine-grained interaction between the robot and the environment. The adaptive process consists in retaining or discarding the free parameters on the basis of their effects at the level of the overall behaviour exhibited by the robot (see [11], [12] for an illustration of the methodological approach employed).

The proposed scenario involve a simulated anthropomorphic robotic arm, equipped with coarse-grained tactile sensors and proprioceptors which encode the position of the arm and of the hand (see Figure 1). The robot is asked to perceptually categorize spherical and ellipsoid objects. The two objects are rather similar (i.e., the longest radius of the ellipsoid is only 20% longer than the radius of the sphere). The robot is allowed to interact in different trials (each lasting 4 seconds) with different objects (one at a time) placed over a table. The objects are placed in the two different initial locations shown in Figure 1(c). Moreover, ellipsoid objects are placed in orientations which are randomly chosen, in each trial, within the four sectors shown in Figure 1(d).

The free parameters that are varied during the adaptive process consists in the synaptic weights and in the time constant of the neurons of a continuous time neural controller shown in Figure 2. Variation of the free parameters are retained or discarded on the basis of the ability of the robot to: (i) categorize the shape of the objects at the end of each trial (i.e., to label objects with different shapes with non-overlapping outputs in a two-dimensional categorization space); and (ii) keep touching the object with the palm of the hand. The robots are thus left free to determine how to interact with the object (providing that they keep touching the object with the palm) and how to label each category (provided that the labels for the two objects do not overlap in the categorization space). In the next four sections, we describe in details the characteristics of the body of the robot, of the sensors and of the actuators, of the neural controller, of the evolutionary algorithm, and of the fitness function. In section VI, we describe the obtained results. Finally, in section VII, we draw our conclusions and we illustrate our

E. Tuci, G. Massera, and S. Nolfi are with the Institute of Cognitive Sciences and Technologies, CNR, Via S. Martino della Battaglia 44, 00185 Rome, IT (phone: +39-06-44595255; fax: +39-06-44595243; email: [elio.tuci, gmassera, stefano.nolfi]@istc.cnr.it).

Fig. 1. The kinematic chain (a) of the arm, and (b) of the hand. Cylinders represent rotational DOFs. The axes of cylinders indicate the corresponding axis of rotation. The links among cylinders represents the rigid connections that make up the arm structure. The numbers from 1 to 10 refer to the parts of the hand equipped with tactile sensors whose readings are inputs of the robot controller. See the text for details on the notation. (c) The two initial positions. Angle of joints $J_1, \dots, J_7$ are $\{-50^\circ, -20^\circ, -20^\circ, -100^\circ, -30^\circ, 0^\circ, -10^\circ\}$ for position A, and $\{-100^\circ, 0^\circ, 10^\circ, -30^\circ, 0^\circ, 0^\circ, -10^\circ\}$ for position B. The sphere and the ellipsoid viewed (d) from above; (e) from the left. The radius of the sphere is 2.5 cm. The radii of the ellipsoid are 2.5, 3.0 and 2.5 cm. In (d) the arrows indicate the intervals within which the initial rotation of the ellipsoid is set.

future plans.

II. THE ROBOT'S STRUCTURE

The simulated robot consists of an anthropomorphic robotic arm with 7 actuated degrees of freedoms (hereafter DOFs) and a hand with 20 actuated DOFs. Proprioceptive and tactile sensors are distributed on the arm and the hand. The robot and the robot/environmental interactions are simulated using Newton Game Dynamics (NGD), a library for accurately simulating rigid body dynamics and collisions (more details at www.newtondynamics.com). The arm consists mainly of three elements: the arm, the forearm, and the wrist (see Figure 1(a)). These elements are connected through articulations displaced into the shoulder (joint J_1 for the extension/flexion, J_2 for the abduction/adduction, and J_3 for the supination/pronation movements), the elbow (joint J_4 for the extension/flexion movements), and the wrist (joints J_5, J_6, J_7 for the pitch/roll/yaw movements).

The robotic hand is composed of a palm and fourteen phalangeal segments that make up the digits (two for the thumb and three for each of the other four fingers) connected through 15 joints with 20 DOFs (see Figure 1(b)). The joints in the hand belong to three different types: metacarpophalangeal (MP), proximal interphalangeal (DIP), and distal interphalangeal (PIP). All of them bring forth the extension/flexion movements of each finger while only the MP joints are for the abduction/adduction movements (Figure 1(b)). The thumb has an extra DOF in MP joints

which is for the axial rotation. This rotation makes possible to move the thumb towards the other fingers (see [13] for a detailed description of the structural properties of the arm). The joints of the arm are actuated by two simulated antagonist muscles implemented accordingly to the Hill's muscle model, as detailed in the next Section.

III. THE ROBOT'S SENSORS, CONTROLLER, AND ACTUATORS

The agent controller consists of a continuous time recurrent non-linear network (CTRNN) with 22 sensory neurons, 8 internal neurons, and 18 output neurons (see Figure 2 and also [14]). At each time step, the activation values y_i of sensory neurons $i = 1, \dots, 7$ is updated on the basis of the state of the proprioceptive sensors of the arm and of the wrist which encode the current angles, linearly scaled in the range $[-1, 1]$, of the seven corresponding joints located on the arm and on the wrist (i.e., joints $J_1, J_2, J_3, J_4, J_5, J_6$, and J_7 in Figure 1(a)). The activation values y_i of sensory neurons $i = 8, \dots, 17$ is updated on the basis of the state of tactile sensors distributed over the hand. These sensors are located on the palm, on the second phalange of the thumb, and on the first and third phalange of each finger (see Figure 1(b)). These sensors return 1 if the corresponding part of the hand is in contact with any another body (e.g., the table, the sphere, the ellipsoid, or other parts of the arm), otherwise 0. The activation values y_i of sensory neurons $i = 18, \dots, 22$ is updated on the basis of the state of the hand proprioceptive

sensors which encode the current extension/flexion state of the five corresponding fingers (i.e., the state of the MP-B joint for the thumb and the MP joints of the other fingers). The readings of the hand proprioceptors are linearly scaled in the range $[0, 1]$ (with 0 for fully extended and 1 for fully flexed finger). To take into account the fact that sensors are noisy, tactile sensors return, with 5% probability, a value different from the computed one, and 5% uniform noise is added to proprioceptive sensors.

Internal neurons are fully connected. Additionally, each internal neuron receives one incoming synapse from each sensory neuron. Each output neuron receives one incoming synapse from each internal neuron. There are no direct connections between sensory and output neurons. The network neurons are governed by the following equation:

$$\begin{aligned} \tau_i \dot{y}_i &= \begin{cases} -y_i + gI_i; & i = 1, \dots, 22 \\ -y_i + \sum_{j=n}^m \omega_{ji} \sigma(y_j + \beta_j); & i = 23, \dots, 48; \end{cases} \\ n &= 1, m = 30 \text{ for } i = 23, \dots, 30; \\ n &= 23, m = 30 \text{ for } i = 31, \dots, 48; \\ \sigma(x) &= \frac{1}{1 + e^{-x}} \end{aligned} \quad (1)$$

In this equation, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, g is a gain factor, I_i the intensity of the perturbation on sensory neuron i , ω_{ji} the strength of the synaptic connection from neuron j to neuron i , β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate. τ_i with $i = 23, \dots, 30$, β_i with $i = 1, \dots, 48$, all the network connection weights ω_{ij} , and g are genetically specified networks' parameters. τ_i with $i = 1, \dots, 22$ and $i = 31, \dots, 48$ is equal to ΔT . There is one single bias for all the sensory neurons.

The activation values y_i of motor neurons determine the state of the simulated muscles of the arm. In particular, the total force exerted by a muscle is the sum of three forces $T_A(\sigma(y_i + \beta_i), x) + T_P(x) + T_V(\dot{x})$, which are calculated on the basis of the following equations:

$$T_A = \sigma(y_i + \beta_i) \left(-\frac{A_{sh} T_{max} (x - R_L)^2}{R_L^2} + T_{max} \right) \quad (2)$$

$$\begin{aligned} A_{sh} &= \frac{R_L^2}{(L_{max} - R_L^2)} \\ T_P &= T_{max} \frac{\exp \left\{ K_{sh} \frac{x - R_L^2}{L_{max} - R_L} \right\} - 1}{\exp \{ K_{sh} \} - 1} \\ T_V &= b \cdot \dot{x} \end{aligned}$$

where $\sigma(y_i + \beta_i)$ is the firing rate of output neurons $i = 31, \dots, 46$, with $i = 31, 32$ for joint J_1 , $i = 33, 34$ for joint J_2 , $i = 35, 36$ for joint J_3 , $i = 37, 38$ for joint J_4 , $i = 39, 40$ for joint J_5 , $i = 41, 42$ for joint J_6 , $i = 43, 44$ for joint J_7 . x is the current elongation of the muscle; L_{max} and R_L are the maximum and the resting length of the muscle; T_{max} is the maximum force that could be generated; K_{sh} is the passive shape factor and b is the viscosity coefficient. The parameters of the equation are identical for all fourteen

Fig. 2. The architecture of the neural controllers.

muscles controlling the seven DOFs of the arm and have been set to the following values: $K_{sh} = 3.0$, $R_L = 2.5$, $L_{max} = 3.7$, $b = 0.9$, $A_{sh} = 4.34$ with the exception of parameter T_{max} which is set to 3000N for joint J_2 , to 300N for joints J_1 , J_3 , J_4 , and J_5 , and to 200N for joints J_6 and J_7 . Muscle elongation is simulated by linearly mapping within specific angular ranges the current angular position of each DOF (see [13] for details).

The joints of the hand are actuated by a limited number of independent variables through a velocity-proportional controller. That is, for the extension/flexion, the force exerted by the MP, PIP, and DIP joints (MP-A, MP-B, and PIP in the case of the thumb) are controlled by a two steps process: first, the θ is set equal to the firing rate $\sigma(y_i + \beta_i)$ of the output neuron i , linearly mapped into the range $[-90^\circ, 0^\circ]$; second, the desired angular positions of the finger joints MP, PIP, DIP are set to θ , θ , and $(2.0/3.0) \cdot \theta$ respectively. For the thumb, its movement towards the other fingers (i.e., the extra DOF in MP joints) corresponds to the desired angle of $-(2.0/3.0)/\theta$. The DOFs that regulate the abduction/adduction movements of the fingers are not actuated.

The activation values y_i of output neurons $i = 47, 48$ are used to categorize the shape of the object (i.e., to produce different output patterns for different object types (see Section V).

IV. THE EVOLUTIONARY ALGORITHM

A simple generational genetic algorithm is employed to set the parameters of the networks (see [15]). The initial population contains 100 genotypes. Generations following the first one are produced by a combination of selection with elitism, and mutation. For each new generation, the 20 highest scoring individuals ("the elite") from the previous generation are retained unchanged. The remainder of the new population is generated by making 4 mutated copies of each of the 20 highest scoring individuals. Each genotype is a vector comprising 420 parameters. Each parameter is encoded with 16 bits. Initially, a random population of vectors is generated. New genotypes, except "the elite", are produced by applying mutation. Mutation entails that each bit of the genotype can be flipped with a 1.5% probability. Genotype parameters are linearly mapped to produce network

parameters with the following ranges: biases $\beta_i \in [-4, -2]$, weights $\omega_{ij} \in [-6, 6]$, gain factor $g \in [1, 10]$ for all the sensory neurons; decay constants τ_i with $i = 23, \dots, 30$ are exponentially mapped into $[10^{-2}, 10^{0.3}]$ with the lower bound corresponding to the integration step-size used to update the controller and the upper bound, arbitrarily chosen, corresponds to about $\frac{1}{2}$ of the maximum length of a trial (i.e., 2 s). Cell potentials are set to 0 when the network is initialised or reset, and circuits are integrated using the forward Euler method with an integration step-size $\Delta T = 0.01$ (see [16]).

V. THE FITNESS FUNCTION

During evolution, each genotype is translated into an arm controller and evaluated 8 times in position A and 8 times in position B (see Figure 1(c)). For each position, the arm experiences 4 times the ellipsoid and 4 times the sphere for a total of $E = 16$ trials. In each position, the rotation of the ellipsoid with respect to the z-axis is randomly set in the range $[350^\circ, 10^\circ]$ in the first presentation, $[35^\circ, 55^\circ]$ in the second presentation, $[80^\circ, 100^\circ]$ in the third presentation, and $[125^\circ, 145^\circ]$ in the fourth presentation (see also Figure 1(d)). At the beginning of each trial, the arm is located in the corresponding initial position (i.e., A or B), and the state of the neural controller is reset. A trial lasts 4 simulated seconds ($T=400$ time step). A trial is terminated earlier in case the object falls off the table.

In each trial e , an agent is rewarded by an evaluation function which seeks to assess its ability to recognise and distinguish the ellipsoid from the sphere. This requires an agent to be able to categorize the objects; that is, to place them in non-overlapping regions of a two-dimensional categorization space $C \in [0, 1] \times [0, 1]$. The categorization and the evaluation of the agent's discrimination capabilities is done in the following way:

- in each trial e , the agent represents the experienced object (i.e., the sphere S or the ellipsoid D) by associating to it a rectangle R_{Se} or R_{De} whose vertices are:

the bottom left vertex:

$$\left(\min_{0.95T < t < T} \sigma(y_{47}(t) + \beta_{47}), \min_{0.95T < t < T} \sigma(y_{48}(t) + \beta_{48}) \right)$$

the top right vertex:

$$\left(\max_{0.95T < t < T} \sigma(y_{47}(t) + \beta_{47}), \max_{0.95T < t < T} \sigma(y_{48}(t) + \beta_{48}) \right)$$

- the sphere category, referred to as C_S , corresponds to the minimum bounding box of all R_{Se} ; the ellipsoid category, referred to as C_D , corresponds to the minimum bounding box of all R_{De} .

The final fitness FF attributed to an agent is the average

score over a set of 16 trials and it is computed as follows:

$$\begin{aligned} FF &= F_1 + F_2 \\ F_1 &= \frac{1}{E} \sum_{e=1}^E \left(1 - \frac{d_e}{d_{max}} \right) \\ F_2 &= \begin{cases} 0 & \text{if } F_1 < 1; \\ 1 - \frac{\text{area}(C_S \cap C_D)}{\min\{\text{area}(C_S), \text{area}(C_D)\}} & \text{otherwise} \end{cases} \end{aligned} \quad (3)$$

with d_e the euclidean distance between the object and the centre of the palm at the end of the trial e ; d_{max} the maximum distance between the palm and the object when located on the table. F_1 rewards the robots for touching the objects. F_2 corresponds to the inverse of a quantity which indicates how much the categorization spaces C_S and C_D overlap. $F_2 = 1$ if C_S and C_D do not overlap (i.e., if $C_S \cap C_D = \emptyset$). The fact that, for each individual, F_1 must be 1 to be rewarded with F_2 , constrains evolution to work on strategies in which the palm is constantly touching the object. This condition has been introduced because we thought it represents a pre-requisite for the ability to perceptually discriminate the shape of the objects. However, alternative formalisms which encode different evolutionary selective pressures may work as well.

VI. RESULTS

Eight evolutionary simulations, each using a different random initialisation, were run for 500 generations. Figure 3 shows the fitness of the best individual at each generation for the best three evolutionary runs. Notice that, after generation 300, the best individuals of all the three runs display optimal or close to optimal performance. This means that these individuals manage to touch the objects with the palm and to distinguish ellipsoid from spherical objects located in the two different spatial locations and regardless of the rotation around the z-axis of the ellipsoid.

In the next parts of this section, we show the results of two series of post-evaluation tests aimed to estimate the robustness of the best evolved categorization strategies: (A) under circumstances in which the effect of favourable conditions linked to the initial rotation of the ellipsoid are

Fig. 3. Fitness of the best individual at each generation of the best three evolutionary runs.

Fig. 4. (a) Test P_1 , individual I_1 . Continuous grey line rectangles correspond to R_{Se} . Continuous black line rectangles correspond to R_{De} . Dashed line rectangles are non-overlapping minimal bounding boxes referred to as C_{Si} and C_{Di} . Pairs of non-overlapping minimal bounding boxes which delimit, for each post-evaluation test P_i with $i = 1, \dots, 5$, the categorization responses of (b) individual I_1 ; (c) individual I_2 ; and (d) individual I_3 .

ruled out, and (B) under circumstances in which the initial position of the object and of the hand varies. Finally, we analyze the dynamics of the robot's categorization behaviour.

A. Robustness with respect to the initial rotation of the ellipsoid

To verify to what extent the robots are able to discriminate between the two type of objects regardless the initial orientation of the ellipsoid object, we tested the evolved robots with objects placed in all possible initial orientations. More precisely in the test P , the three highest fitness individuals (I_j with $j = 1, 2, 3$) taken from run n. 2, are demanded to distinguish for 360 times the two objects placed in position A, and for 360 times placed in position B. In each position, an individual experiences half of the times the sphere (i.e., for 180 trials) and half of the times the ellipsoid (i.e., for 180 trials). Moreover, trial after trial, the initial rotation of the ellipsoid around the z-axis changes of 1° , from 0° in the first trial to 179° in the last trial.

Note that, compared to the evolutionary conditions, in which the individuals is allowed to perceive the ellipsoid only 4 times with 4 different initial rotations, P_i is a severe test. The results unambiguously tell us whether or not

the three selected highest fitness individuals are capable of distinguishing and categorising the ellipsoid from the sphere for whatever rotation of the former object around the z-axis. For each selected individual, test P_i is repeated 5 times (i.e., P_i with $i = 1, \dots, 5$), with each repetition differently seeded to guaranteed random variations in the noise added to sensors readings.

The performance of the individual I_j at test P_i is quantitatively established by considering all the responses given by I_j over 3600 trials (i.e., 720 trials per test P_i , repeated 5 times, with $j = 1, 2, 3$, and $i = 1, \dots, 5$). In each post-evaluation trial, the response of the individual is based on the firing rates of neurons 47 and 48 during the last 4 time steps of each trail e . In particular, the smallest and the highest firing rates recorded by both neurons are used to define the bottom left and the top right vertices of a rectangle, as illustrated in Section V. At the end of each test P_i (i.e., a set of 720 trials), we have 360 rectangles associated to trials in which the individual experienced the sphere (hereafter, rectangles R_{Se}), and 360 rectangles associated to trials in which the individual experienced the ellipsoid (hereafter, rectangles R_{De}). At end of the five post-evaluation tests P_i , we build five pairs of non-

overlapping minimal bounding boxes (i.e., C_{Si} and C_{Di}), a pair for each test i , as explained in Section V.

At this point, we take as a quantitative estimate of the robustness of an agent categorization strategy, the highest number of R_{Se} and R_{De} rectangles that can be included in C_{Si} and C_{Di} respectively, by fulfilling the condition that none of the C_{Si} overlaps with any of the C_{Di} .

Figure 4(a) and Figure 4(b) visually illustrate this evaluation process for individual I_1 . In Figure 4(a), which refers to test P_1 , R_{Se} are the grey continuous line rectangles and R_{De} are the black continuous line rectangles. The grey dashed line rectangle is C_{Si} and the black dashed line rectangle is the C_{Di} . The R_{De} not included in the C_{Di} are those rectangles which needed to be excluded in order to have none of C_{Si} overlapping with any of the C_{Di} . The five pairs of minimal bounding boxes, that correspond to the performance of I_1 during the five test P_i , can be seen in Figure 4(b). Figure 4(c) and Figure 4(d) show the five pairs of bounding boxes corresponding to the performance of I_2 and I_3 respectively.

Table I shows, for each selected individual and for each test P_i the number of rectangles (R_{Se} and R_{De}) for post-evaluated individual (I_j with $j = 1, 2, 3$), and for post-evaluation test P_i with $i = 1, \dots, 5$, that can be included in C_{Si} and C_{Di} by fulfilling the condition that none of the C_{Si} overlaps with any of the C_{Di} . The last row of this Table tells us that the total number of rectangles for each individual, that can be included by the minimal bounding boxes without breaking the non-overlapping rule, is extremely high. From this, we conclude that the selected individuals are extremely good in discriminating and categorising the sphere and the ellipsoid regardless of the rotation of the ellipsoid.

TABLE I

NUMBER OF RECTANGLES, $R_{Se} + R_{De}$, FOR THE FIVE POST-EVALUATION TEST P_i OF THE THREE INDIVIDUAL I_j . THE TOTAL ROW IS THE SUM OF ALL RECTANGLES FOR ALL P_i .

	I_1	I_2	I_3
P_1	717	715	714
P_2	719	712	709
P_3	716	711	708
P_4	717	716	713
P_5	718	718	713
Tot.	3587 of 3600	3572 of 3600	3557 of 3600

B. Robustness with respect to the initial position of the objects

In this section, we show the results of further post-evaluations in which we test the robustness of individual I_1 in trials in which the initial positions of object and of the arm change. Note that, an exhaustive analysis on the capability of the robot to perform the categorization task for objects placed at any spatial position reachable by the hand, it would be computationally expensive due to the large number of arm-objects positions to test and to the fact that each position of the object can be reached through a large number of different postures of the arm. Therefore, we decided to limit the

Fig. 5. These graphs show the percentage of success in post-evaluation tests in which the initial position of one joint is displaced. Dark grey bars refer to % of success (over 1800 trials) with respect to displacements applied to joints in position A. Light grey bars refer to % of success (over 1800 trials) with respect to displacements applied to joints in position B.

analysis only to those circumstances in which the movement of the arm with respect to the position experienced during evolution are determined by displacements of only one joint at time. In particular, joint J_1 , J_2 , J_4 , and J_6 are displaced, one at time, and by intervals of 1° , up to a displacement of $\pm 30^\circ$ from the initial positions experienced during evolution (see Figure 1(c) for details). For each joint and for each 1° displacement, we repeated tests P_i with $i = 1, \dots, 5$ above described. For every displacement, the table and the objects are repositioned to always keep the object below the palm. In these tests, a trial with the sphere/ellipsoid is considered successful if the R_{Se}/R_{De} rectangle falls completely within the region delimited by the minimal bounding box C_{Si}/C_{Di} shown in Figure 4(b). Joint J_3 , J_5 , and J_7 have not been tested because any single displacement of just one of these joints, followed by a repositioning of the table and the object, would disrupt the original spatial relationship (e.g., the palm parallel to the XY plane and to the table) between the object and the hand as shown in Figure 1(c).

The results of these tests are shown in Figure 5. In these graphs, which show the percentage of success per displacement, dark grey bars refer to tests in which the displacements are with respect to position A, while light grey bars refer to tests in which the displacements are with respect to position B. We notice that, in position A, the strategy of I_1 can tolerate quite well displacements which concern joint J_6 , and those concerning J_1 and J_2 up to about an interval of 14° in both directions (see Figure 5(a), and 5(b) dark grey bars). I_1 is extremely sensitive with respect to displacement from position B for all the tested joints (see Figure 5(a), 5(b), and 5(c), light grey bars), and with respect to displacement from position A for joint J_4 (see Figure 5(c), dark grey bars).

C. Analysis of the dynamics of the robot's categorization behaviour

Looking at the movies of the performances of I_1 (i.e., the best performing individual in Table I), we see that this individual starts grasping the object in the first time steps of each trial, and then handles it by letting it slowly roll in between the table and the palm¹. After about 2 s the arm and the hand reach a posture which remains substantially stable until the end of the trial. In order to understand how the robot's categorization outputs distinguish the two objects we run further analysis. In particular, we looked at the trajectories of the average decision outputs in the two-dimensional categorization space $\{\sigma(y(t)_{47} + \beta_{47}), \sigma(y(t)_{48} + \beta_{48})\}$, recorded in test P_1 , by distinguishing between the average values recorded over a set of 180 trials: (i) with the ellipsoid in position A and B (see Figure 6, dark continuous lines, and dark dashed lines respectively); and (ii) with the sphere in position A and B (see Figure 6, grey continuous lines and grey dashed lines respectively). Filled polygons around each trajectory are the standard deviations.

¹Movies of the performances of the best evolved individuals and other complementary materials can be found at <http://laral.istc.cnr.it/esm/discrimination>.

Figure 6(a) refers to trajectories starting from the first ($t = 1$) and ending to the last ($t = T$) time step. Figure 6(b) refers to trajectories starting from the 20th ($t = 20$) and ending to the last ($t = T$) time step. By comparing the two graphs in Figure 6, we notice that, I_1 moves in the categorization space by reaching in less than 0.2 s (from the beginning of the trial) an area in the proximity of the minimal bounding boxes (see Figure 6(b), dashed rectangles). During the large majority of the arm-object interactions, following the first 0.2 s, I_1 moves relatively slow in the categorization space, with the trajectories that tend to diverge after 3 s. Moreover, we notice that, as expected, the standard deviation is higher for the trials with the ellipsoid (see Figure 6(b)). This is clearly due to the fact that, trial after trial, the ellipsoid rotates around the z axis.

The identification of the mechanism which allow the evolved robots to reliably discriminate between the two type of objects, is a particularly challenging task given the complexity of the robot, with many sensors, many actuators and a recurrent non-linear control structure. Performing this analysis goes beyond the objectives of the paper, and it is left for future work.

VII. CONCLUSIONS

In this paper, we described an experiment in simulation in which an anthropomorphic robotic arm, provided with tactile sensors and proprioceptors develops an ability to perceptually categorize spherical and ellipsoid objects. The acquisition of such capacity has been realized through an evolutionary method in which the free parameters of the robots neural controller have been evolved for the ability to produce different categorization outputs and for touching the object with the palm of the hand. During the adaptive process the robots are left free to determine how they interact with the objects (provided that they keep touching the objects with their palm) and how they represent the objects experienced within the two-dimensional categorization space corresponding to the output of the two categorization units (provided that the areas corresponding to the two type of objects do not overlap).

The coarse-grained sensory apparatus of the robotic arm, the need to control 16 different actuators (which affect the state of 27 DOFs), the need to master the effects arising from the physical interactions between the robot and the environment, and the small differences between the two objects, make the perceptual categorization task particularly challenging. Nevertheless, the best evolved robots are able to accomplish their task robustly by displaying close to optimal performance, regardless of the orientation of the objects. The analysis of the best evolved controllers also indicates that they are able to generalize their ability, within limits, for initial positions of the objects and of the arm which have never been experienced during the evolutionary phase.

The analysis of the motor behaviour indicates that the evolved robots accomplish their task by actively manipulating the objects for about 3 s until the arm and the hand of the robot assume a posture that remains substantially

Fig. 6. Test P_1 , individual I_1 . Trajectories of the average decision outputs in the two-dimensional categorization space ($\sigma(y(t)_{47} + \beta_{47})$, $\sigma(y(t)_{48} + \beta_{48})$), with (a) $1 < t < T$; and (b) $20 < t < T$. Black continuous and dashed lines refers to the average values recorded over sets of 180 trials with the ellipsoid in position A and B, respectively. Grey continuous and dashed lines refers to the average values recorded over sets of 180 trials with the sphere in position A and B, respectively. Filled polygons around each trajectory are the standard deviations.

stable during the rest of the trial. The analysis of the categorization process in evolved robots indicates that they start to perceptually differentiate the two categories already after about 0.2 s. From this time on until the end of the trial, the categorization outputs further differentiate until they reach either areas (C_S or C_D) of the categorization space.

In future work, we intend to analyze in details the mechanisms which allow the robot to discriminate between the two categories of objects. In particular we will analyse whether categorization is accomplished by exploiting the effects of the interaction between the robot and the environment (mediated by the execution of a specific behaviour) on the posture assumed by the hand and/or by the arm (as observed in [17], in a much simpler setup). We will look at the role of the tactile sensation, and in particular, to what extent it contributes to the categorization process. We will analyse whether categorization is affected by how the states of the sensors change over time or simply by their current state, and to what extent the internal states of the controller affect the way in which the robot manipulate the object.

Moreover, we plan to investigate: (i) how the representation of categories varies in the categorization space during the training process and which are the effects of the dimensionality of the categorization space; (ii) whether experiencing a larger variety of positions/orientations of the objects during the evolutionary phase, helps the robots to develop more robust categories; and (iii) to what extent the model scale up to a larger number of categories and object shapes.

ACKNOWLEDGMENT

This research work was supported by the *ITALK* project (EU, ICT, Cognitive Systems and Robotics Integrating Project, grant n° 214668). The authors thank Tomassino Ferrauto and their colleagues at LARAL for stimulating discussions and feedback during the preparation of this paper.

REFERENCES

[1] A. Noë, *Action in Perception*. MIT Press, Cambridge, MA, 2005.

[2] R. Pfeifer and C. Scheier, *Understanding Intelligence*. Cambridge, MA: MIT Press, 1999.

[3] S. Nolfi, "Power and limits of reactive agents," *Neurocomputing*, vol. 42, pp. 119–145, 2002.

[4] R. D. Beer, "The dynamics of active categorical perception in an evolved model agent," *Adaptive Behavior*, vol. 11, pp. 209–243, 2003.

[5] E. Tuci, C. Ampatzis, F. Vicentini, and M. Dorigo, "Evolving homogeneous neuro-controllers for a group of heterogeneous robots: coordinated motion, cooperation, and acoustic communication," *Artificial Life*, vol. 14, no. 2, pp. 157–178, 2008.

[6] S. Nolfi, "Behavior and cognition as a complex adaptive system: Insights from robotic experiments," in *Philosophy of Complex Systems, Handbook on Foundational/Philosophical Issues for Complex Systems in Science*, C. Hooker, Ed. Elsevier, In Press.

[7] N. Franceschini, J. M. Pichon, C. Blanes, and J. M. Brady, "From insect vision to robot vision," *Philosophical Transactions: Biological Sciences*, vol. 337, no. 1281, pp. 283–294, 1992.

[8] M. Dill, "Visual pattern recognition in drosophila involves retinotopic matching," *Nature*, vol. 365, pp. 751–753, 1993.

[9] C. Scheier, R. Pfeifer, and Y. Kuniyoshi, "Embedded neural networks: exploiting constraints," *Neural Networks*, vol. 11, pp. 1551–1596, 1998.

[10] S. Nolfi, "Categories formation in self-organizing embodied agents," in *Handbook of Categorization in Cognitive Science*, H. Cohen and C. Lefebvre, Eds. Elsevier, 2005, pp. 869–889.

[11] S. Nolfi and D. Marocco, "Evolving visually-guided robots able to discriminate between different landmarks," in *From Animals to Animats 6. Proc. of the VI Int. Conf. on Simulation of Adaptive Behavior*, J. A. Meyer, A. Berthoz, D. Floreano, H. L. Roitblat, and S. W. Wilson, Eds. Cambridge, MA: MIT Press, 2000, pp. 413–419.

[12] I. Harvey, E. Di Paolo, R. Wood, M. Quinn, and E. Tuci, "Evolutionary robotics: A new scientific tool for studying cognition," *Artificial Life*, vol. 11, no. 1-2, pp. 79 – 98, 2005.

[13] G. Massera, A. Cangelosi, and S. Nolfi, "Evolution of prehension ability in an anthropomorphic neurorobotic arm," *Front. Neurorobot.*, vol. 1, 2007.

[14] R. D. Beer and J. C. Gallagher, "Evolving dynamic neural networks for adaptive behavior," *Adaptive Behavior*, vol. 1, no. 1, pp. 91–122, 1992.

[15] D. E. Goldberg, *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley, 1989.

[16] S. H. Strogatz, *Nonlinear Dynamics and Chaos*. Perseus Books Publishing, 2000.

[17] S. Nolfi and D. Marocco, "Active perception: A sensorimotor account of object categorisation," in *Proc. of the 7th Int. Conf. on Simulation of Adaptive Behavior (SAB '02)*, B. Hallam, D. Floreano, J. Hallam, G. Hayes, and J.-A. Meyer, Eds. MIT Press, Cambridge, MA, 2002, pp. 266–271.