

# Active categorical perception of object shapes in a simulated anthropomorphic robotic arm

Elio Tuci, Gianluca Massera, and Stefano 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 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 categorise spherical and ellipsoid objects. We show that best individuals, synthesised by artificial evolution techniques, develop a close to optimal ability to discriminate the shape of the objects as well as an ability to generalise their skill in new circumstances. The results show that the agents solve the categorisation task in an effective and robust way by self-selecting the required information through action and by integrating experienced sensory-motor states over time.

**Index Terms**—Categorical perception, evolutionary robotics, artificial neural networks.

## I. INTRODUCTION

Categorical perception can be considered the ability to divide continuous signals received by sense organs into discrete categories whose members resemble more one another than members of other categories. Categorical perception represents one of the most fundamental cognitive capacities displayed by natural organisms, and it is an important prerequisite for the exhibition of several other cognitive skills [see 1]. Not surprisingly, categorical perception has been extensively studied both in natural sciences such as Psychology, Philosophy, Ethology, Linguistics, and Neuroscience, and in artificial sciences such as Artificial Intelligence, Neural Networks, and Robotics [see 2, for a comprehensive review of this research field]. However, in the large majority of the cases, researchers have focused their attention on categorisation processes that are passive and instantaneous. Passive categorisation processes take place in those experimental setups in which the agents can not influence the experienced sensory states through their actions. Instantaneous categorisation processes are those in which the agents are demanded to categorise the current experienced sensory state rather than a sequence of sensory states distributed over a certain time period.

In this paper, instead, we study categorisation processes that are active and eventually distributed over time [3, 4]. This task is achieved by exploiting the properties of autonomous embodied and situated agents. An important consequence of

being situated in an environment consists in the fact that the sensory stimuli experienced by an agent are co-determined by the action performed by the agent itself. That is, the actions and the behaviour exhibited by the agent later influence the stimuli it senses, their duration in time, and the sequence with which they are experienced. This implies that: (i) categorical perception is strongly influenced by an agent’s action [see also 5, 6, on this issue]; and (ii) sensory-motor coordination (i.e., the ability to act in order to sense stimuli or sequence of stimuli that allow an agent to perform its task) is a crucial aspect of perception and more generally of situated intelligence [see 7].

Although the significance of embodiment and situatedness for the study of the underlying mechanisms of behaviour and cognition is widely recognised, building artificial systems that are able to actively perceive and categorise sensory experiences is a challenging task. This can be explained by considering that, from the point of view of the designer, identifying the way in which an agent 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 evolutionary methods in which the agents are left free to determine how they interact with the environment (i.e., how they behave, in order to solve their task). With these methods, free parameters (i.e., those that are modified during the evolutionary process) encode features that regulate the fine-grained interactions between the agent and the environment. The evolutionary 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 agent [see 8, 9, 10, for a detailed illustration of the methodological approach employed].

In this paper, we describe an experiment in which evolutionary methods are used to investigate the perceptual skills of an autonomous agent demanded to actively categorise unanchored spherical and ellipsoid objects placed in different positions and orientations over a planar surface. The agent is a simulated anthropomorphic robotic arm with 27 actuated Degrees of Freedom (hereafter, DoFs). The arm is equipped with coarse-grained tactile sensors and with proprioceptive sensors encoding the position of the joints of the arm and of the hand. The task requires the agent to produce different categorisation outputs for objects with different shapes and similar categorisation outputs for objects with the same shape. The aim of this study is to prove that, in spite of the complexity of the experimental scenario, the evolutionary approach can be successfully employed to design neural mechanisms to allow the robotic arm to perform the perceptual categorisation task. Moreover, we unveil the operational principles of successful

E. Tuci, G. Massera, and S. Nolfi are with the ISTC-CNR, Via San Martino della Battaglia, n. 44, 00185 Rome, Italy, e-mail: {elio.tuci, gianluca.massera, stefano.nolfi}@istc.cnr.it (see <http://laral.istc.cnr.it/elio.tuci/>).

agents. In particular, we look at (i) how the robot acts in order to bring fourth the sensory stimuli which provide the regularities necessary for categorising the objects in spite of the fact that sensation itself may be extremely ambiguous, incomplete, and noisy; (ii) the dynamical nature of sensory flow (i.e., how sensory stimulation varies over time and the time rate at which significant variations occur); (iii) the dynamical nature of the categorisation process (i.e., whether the categorisation process occur over time while the robot interacts with the environment); (iv) the role of qualitatively different sensation originated by different sensory channels in the accomplishment of the categorisation task.

We prove that a further elaboration of evolutionary methods proposed in related studies can be successfully applied to problems that are non-trivial and significantly more complex with respect to the state of the art reviewed in Section II. In particular, we show that the best evolved robots develop a close to optimal ability to discriminate the shape of the objects as well as an ability to generalise their skill in new circumstances. These results prove that the problem can be solved in an effective and robust way by self-selecting the required information through action and by integrating experienced sensory-motor states over time.

## II. STATE OF THE ART

There is a growing body of literature in robotics which is devoting increasingly more efforts in obtaining discrimination of material properties (e.g., hardness, texture) and object shape using touch in artificial arms. Many of these works, like the one described in [11], draw inspiration from human perceptual capability to develop highly elaborated touch sensors. In [11], the authors describe a tendon driven robotic hand covered with artificial skin made of strain gauges sensors and polyvinylidene films. The strain gauges sensors mimic the functional properties of Merkel cells in human skin and detect the strain. Polyvinylidene films mimic the functional properties of the Meissner corpuscles and detect the velocity of the strain. The artificial hand, through the execution of squeezing and tapping procedures, manages to discriminate objects based on their hardness. In a similar vein, a research group at the Lund University has developed three progressively complex versions of a robotic hand (LUCS Haptic Hand I, II, and III) designed for haptic perception tasks [12, 13, 14]. The perceptual capabilities of the three versions of LUCS, which differ in their morphology and in their sensory capabilities, have been tested during the execution of a grasping procedure on objects made of different material (e.g., plastic and wood). The authors showed that the sensory patterns generated in interactions with the objects are rich enough to be used as a basis for haptic object categorisation [15]. Other robotics systems combine visual and tactile perception to carry out fairly complex object discrimination tasks [see 16, 17, 18].

Generally speaking, we can say that, in spite of the heterogeneity in hardware and control design, the research works mentioned above focus on the characteristics of the tactile sensory apparatus and/or on the categorisation algorithms. In these works, the way in which the sensory feedback affects

the movement of the hand is determined by the experimenter on the basis of her intuition. Moreover, the discrimination phase follows the exploration phase and it is performed by elaborating sensory data gathered during manipulation of the objects (i.e., the data collected during the exploration phase cannot influence the agents successive behaviour).

The work described in this paper differs significantly from the above mentioned literature since the way in which the agent interacts with the environment is not designed by the experimenter but is adapted in order to facilitate the categorisation task and since the agent is left free to shape its motor behaviour on the basis of previously experienced sensory states. Rather than studying the performances of particularly effective tactile sensors or of specific categorisation algorithms, we focus on the development of autonomous actions for the discrimination of objects shape through coarse-grained binary tactile sensors and proprioceptive sensors. The issue of how a robot can actively develop categorisation skills has been already investigated in few recent research works. In general terms, these works demonstrate how adapted robots exploit their action to self-select stimuli which enable and/or simplify the categorisation process and how this leads to solutions which are parsimonious and robust [see 19, 20, 21, 22].

Particularly relevant for this study is the work described in [23]. The authors studied the case of a simulated robotic “finger” which has been evolved for discriminating the shape of spherical versus cubic objects (anchored to a fixed point) of different sizes and orientations. The robotic finger is constituted by an articulated structure made by three segments connected through motorised joints with six DoFs, six corresponding actuators, six proprioceptive sensors encoding the current position of the joints, and three tactile sensors placed on the three corresponding segments of the finger. The authors observed that the adapted robots solve their problem through simple control rules that makes the robot scan for the object by moving horizontally from the left to the right side and by moving slightly up as a result of collisions between the finger and the object. These simple control rules lead to the exhibition of two different behaviours. With spherical objects, the robotic finger fully extends itself on the left side of the object after following the object surface. With cubic objects, the robotic finger remains fully bended close to one of the corners of the cube. These two behaviours corresponds to well differentiated activations of the proprioceptive sensors. These differences are used by the finger to distinguish the two types of object. Note that, although the discriminating cue necessary to categorise is available in each single sensory pattern experienced after the exhibition of the appropriate behaviour, this cue results from the dynamical process arising as a result of several robot/environmental interactions. In [24], the authors show that a visually guided robot arm whose neuro-controller is evolved for reaching and tracking, can exploit its actions to self-select stimuli which facilitates the accomplishment of spatial and temporal coordination.

Unlike in the experiments described in [23, 24], sensory-motor coordination does not always guarantee the perception of well differentiated sensory states in different contexts corresponding to different categories. Under these circumstances,

the agent can actively categorise their perceptual experiences by integrating ambiguous sensory information over time. Few studies have already shown that evolved wheeled robots compensate for unreliable sensory patterns due to coarse sensory apparatus by acting and re-acting to temporally distributed sensory experiences, in a way to bring forth the necessary regularities that allow them to associate a stimulus with its category [see 25, 26].

The experiment presented in this paper focuses on a non-trivial task that is significantly more complex to that investigated in previous studies due to the high similarity between the objects to be discriminated, the difficulty of controlling a system with many degree of freedom, and the need to master the effects produced by gravity, inertia, collisions, etc. As shown in Section VII, the analysis of the strategy displayed by best evolved robots demonstrates that, also in this case, sensory-motor coordination plays a crucial role, as in [23, 24]. Indeed, the best robots manipulate the objects so to experience the regularities which allow them to appropriately categorise the shape of the objects. However, sensory-motor coordination does not seem to guarantee the perception of fully differentiated sensory states corresponding to different categories. The problem caused by the lack of clear categorical evidences is solved through the development of an ability to integrate ambiguous information over time through a process of evidences accumulation.

### III. THE ROBOT'S STRUCTURE

The simulated robot consists of an anthropomorphic robotic arm with 7 actuated 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](http://www.newtondynamics.com)). The arm consists mainly of three elements: the arm, the forearm, and the wrist. 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 roll/pitch/yaw movements, see Figure 1a).

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 1b). There are three different types of hand joints: metacarpophalangeal (MP), proximal interphalangeal (PIP), and distal interphalangeal (DIP). All of them bring forth the extension/flexion movements of each finger while only the metacarpophalangeal joints are for the abduction/adduction movements. The thumb has an extra DoF in metacarpophalangeal joints which is for the axial rotation. This rotation makes possible to move the thumb towards the other fingers [see 27, for a detailed description of the structural properties of the arm]. The active joints of the robotic arm are actuated by two simulated antagonist muscles implemented according to the Hill's muscle model, as detailed in the next Section.




Fig. 1: The kinematic chain (a) of the arm, and (b) of the hand. (c) The architecture of the arm neural controller. In (a) and (b), 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. In (c) the circles refer to the artificial neurons. Continuous line arrows indicate the efferent connections for the first neuron of each layer. Dashed line arrows indicate the correspondences between joints and tactile sensors and input neurons. The labels on the dashed line arrows refer to the notation used in equation 1a to indicate the readings of the corresponding sensors.

#### IV. 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 motor neurons [see Figure 1c and also 28]. 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 (see joints  $J_1, J_2, J_3, J_4, J_5, J_6$ , and  $J_7$  in Figure 1a). 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 (see label  $T_1$  in Figure 1b), on the second phalange of the thumb (see label  $T_2$  in Figure 1b), and on the first phalange (see labels  $T_4, T_6, T_8, T_{10}$  in Figure 1b) and the third phalange (see labels  $T_3, T_5, T_7, T_9$  in Figure 1b) of each finger. 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 of the five corresponding fingers (see joints  $J_8, J_9, J_{10}, J_{11}$ , and  $J_{12}$  in Figure 1b). The readings of the hand proprioceptive sensors 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 motor neuron receives one incoming synapse from each internal neuron. There are no direct connections between sensory and motor neurons. The values of sensory neurons are updated using equation 1a, the values of internal neurons with equation 1b, and the values of motor neurons with equation 1c.

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

with  $\sigma(x) = (1 + e^{-x})^{-1}$ . In these equations, using terms derived from an analogy with real neurons,  $y_i$  represents the cell potential,  $\tau_i$  the decay constant,  $g$  is a gain factor,  $I_i$  the intensity of the perturbation on sensory neuron  $i$ ,  $\omega_{ji}$  the strength of the synaptic connection from neuron  $j$  to neuron  $i$ ,  $\beta_j$  the bias term,  $\sigma(y_j + \beta_j)$  the firing rate.  $\tau_i$  with  $i = 23, \dots, 30$ ,  $\beta_i$  with  $i = 1, \dots, 48$ , all the network connection weights  $\omega_{ij}$ , and  $g$  are genetically specified networks' parameters.  $\tau_i$  with  $i = 1, \dots, 22$  and  $i = 31, \dots, 48$  is equal to the integration time step  $\Delta T = 0.01$ . 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:

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

where  $\sigma(y_i + \beta_i)$  is the firing rate of output neurons  $i = 31, \dots, 46$ .  $x$  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 can 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 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 27, 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 step process: first,  $\theta$  is set equal to the firing rate  $\sigma(y_i + \beta_i)$  (with  $i = 45$  for the thumb movement, and  $i = 46$  for the other finger movement), linearly mapped into the range  $[-90^\circ, 0^\circ]$ ; second, the desired angular positions of the finger joints MP, PIP, DIP are set to  $\theta$ ,  $\theta$ , 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  $-(2.0/3.0) \cdot \theta$ . 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 categorise the shape of the object (i.e., to produce different output patterns for different object types, see also Section VI).

#### V. THE EVOLUTIONARY ALGORITHM

A simple generational genetic algorithm is employed to set the parameters of the networks [see 29]. 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




Fig. 2: (a) Position A; angle of joints  $J_1, \dots, J_7$  are  $\{-50^\circ, -20^\circ, -20^\circ, -100^\circ, -30^\circ, 0^\circ, -10^\circ\}$ ; (b) Position B; angle of joints  $J_1, \dots, J_7$  are  $\{-100^\circ, 0^\circ, 10^\circ, -30^\circ, 0^\circ, 0^\circ, -10^\circ\}$ ; (c) the sphere and the ellipsoid viewed from above; (d) the sphere and the ellipsoid viewed from the side. The radius of the sphere is 2.5 cm. The radii of the ellipsoid are 2.5, 3.0 and 2.5 cm. In (c) the arrows indicate the intervals within which the initial rotation of the ellipsoid is set.

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. 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  $\tau_i$  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 half of a trial length (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 [see 30].

## VI. 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, for a total of  $K = 16$  trials (see Figure 2a and 2b). For each position, the arm experiences 4 times the ellipsoid and 4 times the sphere. Moreover, 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 2c).

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  $k$ , an agent is rewarded by an evaluation function which seeks to assess its ability to recognise and distinguish the ellipsoid from the sphere. Note that, rather than imposing a representation scheme in which different categories are associated with *a priori* determined state/s of the categorisation neurons (i.e., neurons 47 and 48), we leave the robot free to determine how to communicate the result of its decision. That is, the agents can develop whatever representation scheme as long as each object category is clearly identified by a unique state/s of the categorisation neurons. This system has also the advantage that it scales up to categorisation tasks with objects of more than two categories, without having to introduce structural modifications to the agent's controller. More precisely, we score agents on the basis of the extent to which the categorisation outputs produced for objects of different categories are located in non-overlapping regions of a two dimensional categorisation space  $C \in [0, 1] \times [0, 1]$ . The categorisation and the evaluation of the agent's discrimination capabilities is done in the following way:

- in each trial  $k$ , the agent represents the experienced object (i.e., the sphere  $S$  or the ellipsoid  $E$ ) by associating to it a rectangle  $R_k^S$  or  $R_k^E$  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  $R_k^S$ ; the ellipsoid category, referred to as  $C^E$ , corresponds to the minimum bounding box of all  $R_k^E$ .

The final fitness  $FF$  attributed to an agent is the sum of two fitness components  $F_1$  and  $F_2$ .  $F_1$  rewards the robots for touching the objects, and corresponds to the average distance over a set of 16 trials between the centre of the palm and the experienced objects.  $F_2$  rewards the robots for developing an unambiguous category representation scheme on the basis of the position in a two-dimensional space of  $C^S$  and  $C^E$ .  $F_1$  and  $F_2$  are computed as follows:

$$F_1 = \frac{1}{K} \sum_{k=1}^K \left( 1 - \frac{d_k}{d_{max}} \right), \text{ with } K = 16; \quad (3)$$

$$F_2 = \begin{cases} 0 & \text{if } F_1 \neq 1; \\ 1 - \frac{\text{area}(C^S \cap C^E)}{\min\{\text{area}(C^S), \text{area}(C^E)\}} & \text{otherwise;} \end{cases} \quad (4)$$

with  $d_k$  the euclidean distance between the object and the centre of the palm at the end of the trial  $k$ ;  $d_{max}$  the maximum distance the centre of the palm can reach from the object when located on the table.  $F_2 = 1$  if  $C^S$  and  $C^E$  do not overlap




Fig. 3: Graph showing the fitness of the best agents at each generation of the five evolutionary runs that managed to generate highest score individuals for at least 10 consecutive generations:  $run_1$ ,  $run_2$ ,  $run_3$ ,  $run_4$ ,  $run_5$ .

(i.e., if  $C^S \cap C^E = \emptyset$ ). 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.

## VII. RESULTS

Ten evolutionary simulations, each using a different random initialisation, were run for 500 generations. Figure 3 shows the fitness of the best agent at each generation for the five evolutionary runs that managed to generate highest score individuals for at least 10 consecutive generations. The other five runs failed to achieve this first objective. A quick glance at these curves indicates that  $run_1$  reaches very quickly (in about 100 generations) a plateau on the highest fitness score and keeps on generating highest score agents until the end of evolution.  $run_2$ ,  $run_3$ ,  $run_4$ ,  $run_5$  also generate highest score agents but they need more generations and the solutions seem to be more sensitive to the effect produced by those parameters of the task randomly initialised and/or by noise. Although all the agents with the highest fitness are potentially capable of accomplishing the task, the effectiveness and the robustness of their collective strategies have to be further estimated with more severe post-evaluation tests. In the next Section, we show the results of a series of post-evaluation tests aimed at estimating the robustness of the best evolved discrimination strategies chosen from  $run_1$ ,  $run_2$ ,  $run_3$ ,  $run_4$ , and  $run_5$ . In Section VII-B, we show the results of post-evaluation tests aimed at estimating the role of different sensory channels for categorisation. Finally, in Section VII-C, we analyse the dynamics of the best evolved agents categorisation strategy. It is important to note that, although all the post-evaluation analyses have been carried out on all the best evolved agents, for the sake of space, for several tests we include only the results concerning the performances of one of these agents<sup>1</sup>

<sup>1</sup>An exhaustive description of the analyses carried out on all the best evolved agents, results of tests not shown in the paper, further simulations as well as movies of the bests evolved strategies can be found at [http://lral.istc.cnr.it/esm/active\\_perception](http://lral.istc.cnr.it/esm/active_perception).

### A. Robustness

To verify to what extent the robots are able to discriminate between the two types of object regardless the initial orientation of the ellipsoid object, we run post-evaluation tests (referred to as test P) in which we systematically vary the ellipsoid initial orientation. More precisely in test P, an agent is demanded to distinguish for 360 times the two objects placed in position A, and for 360 times placed in position B. In each position, the agent 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 orientation of the ellipsoid around the z-axis changes of  $1^\circ$ , from  $0^\circ$  in the first trial to  $179^\circ$  in the last trial. For each run, we selected and post-evaluated 10 agents chosen among those with the highest fitness. It is important to note that these agents are selected from evolutionary phases in which the run managed to generate highest score individuals for at least 10 consecutive generations. Table I shows the results of the agent  $A_j$  generated by  $run_j$ , with  $A_j$  being the best agent among those selected from  $run_j$ .

Note that, compared to the evolutionary conditions, in which the agents are allowed to perceive the ellipsoid only 4 times with 4 different initial orientations, P is a severe test. The results unambiguously tell us whether or not the five selected highest fitness agents are capable of distinguishing and categorising the ellipsoid from the sphere in a much wider range of initial orientations of the former object. For each selected agent, test P 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. Table I shows, for each selected agent  $A_j$ , the results of all the five tests  $P_i$ .

The performance of the agent  $A_j$  at test  $P_i$  is quantitatively established by considering all the responses given by  $A_j$  over 3600 trials (i.e., 720 trials per test  $P_i$ , repeated 5 times, with  $i, j = 1, \dots, 5$ ). In each post-evaluation trial, the response of the agent is based on the firing rates of neurons 47 and 48 during the last 20 time steps (i.e.,  $0.95 \cdot T < t < T$ ) of each trail  $k$ . 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 VI. At the end of each test  $P_i$ , we have 360 rectangles associated to trials in which the agent experienced the sphere (i.e., rectangles  $R_k^S$  with  $k = 1, \dots, 360$ ), and 360 rectangles associated to trials in which the agent experienced the ellipsoid (i.e., rectangles  $R_k^E$  with  $k = 1, \dots, 360$ ). At the end of the five post-evaluation tests  $P_i$ , we build five pairs of non-overlapping minimal bounding boxes (i.e.,  $C_i^S$  and  $C_i^E$ ), a pair for each test  $i$ , as explained in Section VI. At this point, we take as a quantitative estimate of the robustness of an agent categorisation strategy, the highest number of  $R_k^S$  and  $R_k^E$  rectangles that can be included in  $C_i^S$  and  $C_i^E$  respectively, by fulfilling the condition that none of the  $C_i^S$  overlaps with any of the  $C_i^E$ . Table I shows, for each selected agent and for each test  $P_i$ , the number of rectangles ( $R_k^S$  and  $R_k^E$ ) for post-evaluated agent, and for post-evaluation test, that can be included in  $C_i^S$  and  $C_i^E$  by fulfilling the condition that none of the  $C_i^S$  overlaps with any of the  $C_i^E$ . The last row of this

Table tells us that, for agent  $A_1$ ,  $A_3$ ,  $A_4$ , and  $A_5$ , the total number of rectangles that can be included by the minimal bounding boxes without breaking the non-overlapping rule is extremely high, with a percentage of success over 97%. These four agents are quite good in discriminating and categorising the sphere and the ellipsoid in a much wider range of initial orientations of the ellipsoid. Agent  $A_2$ , whose performance is slightly worse, is excluded from all further post-evaluation tests.

The agents with a performance at the first test P above 95% (i.e.,  $A_1$ ,  $A_3$ ,  $A_4$ , and  $A_5$ ) undergo a further series of tests P in circumstances in which (i) the length of the longest radius of the ellipsoid progressively increases/decreases (see Figure 4a); (ii) the length of the radius of the sphere progressively increases/decreases (see Figure 4b); (iii) the initial position of the object and of the hand varies (see Figure 4c). In these as well as in all the other post-evaluation tests we describe from now on concerning  $A_1$ ,  $A_3$ ,  $A_4$ , and  $A_5$ , a trial  $k$  can: (i) successfully terminate if the  $R_k^E$ , built as illustrated above, completely falls within the agent's two-dimensional space delimited by the five bounding boxes  $C_i^E$  built during the first test P; (ii) unsuccessfully terminate with a sphere response if the  $R_k^E$  completely falls within the agent's two-dimensional space delimited by the five bounding boxes  $C_i^S$  built during the first test P; (iii) unsuccessfully terminate with a none response, if the  $R_k^E$ , completely falls outside the agent's two-dimensional space delimited by the ten bounding boxes  $C_i^S \cap C_i^E$  built during the first test P.

As far as it concerns tests in which the length of the longest

TABLE I: The table shows, for post-evaluated agent ( $A_j$  with  $j = 1, \dots, 5$ ), and for post-evaluation test ( $P_i$  with  $i = 1, \dots, 5$ ), the number of rectangles  $R_k^E$  and  $R_k^S$  that can be included in bounding boxes  $C_i^E$  and  $C_i^S$ , respectively, by fulfilling the condition that none of the  $C_i^E$  overlaps with any of the  $C_i^S$ . The last row indicates the total number of correct categorisation choices and percentage of success over 3600 evaluation trials. See the text for further details.

	$A_1$		$A_2$		$A_3$	
	$R_k^E$	$R_k^S$	$R_k^E$	$R_k^S$	$R_k^E$	$R_k^S$
$P_1$	357	360	310	351	340	358
$P_2$	359	360	311	347	342	358
$P_3$	356	360	312	349	343	356
$P_4$	357	360	304	353	341	355
$P_5$	358	360	303	348	349	356
Tot./(%)	3587 / 99.6%		3288 / 91.3%		3498 / 97.2%	

	$A_4$		$A_5$	
	$R_k^E$	$R_k^S$	$R_k^E$	$R_k^S$
$P_1$	347	356	355	354
$P_2$	356	358	356	355
$P_3$	348	355	356	354
$P_4$	342	354	354	355
$P_5$	349	355	353	353
Tot./(%)	3520 / 97.8%		3545 / 98.5%	

radius of the ellipsoid progressively increases/decreases, we notice that distortions that further increase the longest ellipsoid radius up to 1 cm, are rather well tolerated by the agents, with  $A_1$  and  $A_5$  that manage to reliably differentiate the two objects with a success rate higher than 90%. Distortions that tend to reduce the longest radius of the ellipsoid are clearly disruptive for all the agents, with an expected 50% success rate when the ellipsoid is reduced to a sphere. In tests in which the ellipsoids have a radius progressively shorter than the radius of the sphere, the performance of all the agents are quite disrupted (see Figure 4a).

As far as it concerns tests in which the length of the radius of the sphere progressively increases/decreases, we notice that these distortions are particularly disruptive for all the agents except for  $A_5$ . This agent is not as disrupted as the other agents in those tests in which the sphere becomes progressively smaller, and it is very successful in tests in which the radius of the sphere is at least 7 millimetres longer than the longest radius of the ellipsoid (see Figure 4b).

Finally, in a further series of post-evaluation tests we estimated the robustness of the best evolved strategies in tests in which the initial positions of object and of the arm change. To simplify our analysis, we focused only on those circumstances in which the movement of the arm respect to the initial positions experienced during evolution are determined by displacements of only one joint at time (see Figure 4c). Although the results are quite heterogeneous, there are some features which are shared by all the agents. First, displacements of joint  $J_1$  for position A are tolerated quite well. Second, the wider the displacement, the bigger the performance drop, with the exception of  $J_4$  for agents  $A_1$   $A_3$   $A_4$ , in which displacements that tend to progressively bring the hand/object closer to the body result in a better performance for both positions. It is important to note that,  $A_4$  is particularly sensitive to disruptions to joint  $J_1$  and  $J_2$  for position B, and joint  $J_6$  for position A.

### B. The role of different sensory channels for categorisation

To understand the mechanisms which allow agents  $A_1$ ,  $A_3$ ,  $A_4$ , and  $A_5$ , to solve their task, we first established the relative importance of the different types of sensory information available through arm proprioceptive sensors (i.e.,  $I_i$  with  $i = 1, \dots, 7$ , see also Figure 1c), tactile sensors (i.e.,  $I_i$  with  $i = 8, \dots, 17$ , see also Figure 1c), and hand proprioceptive sensors (i.e.,  $I_i$  with  $i = 18, \dots, 22$ , see also Figure 1c). This has been accomplished by measuring the performance displayed by the agents in a series of *substitution tests* in which one type of sensory information experienced by each agent during the interaction with an ellipsoid has been replaced with the corresponding type of sensory information previously recorded in trials in which the agent was interacting with a sphere. In these tests, each agent experiences the ellipsoid in all the initial rotations (i.e., from  $0^\circ$  to  $179^\circ$ ) excluding those for which, given the randomly chosen seed for the tests, its responses turned out to be wrong in the absence of any type of substitution (i.e., the rectangle  $R_k^E$  did not fall within any of the five bounding boxes  $C_i^E$  resulted from the

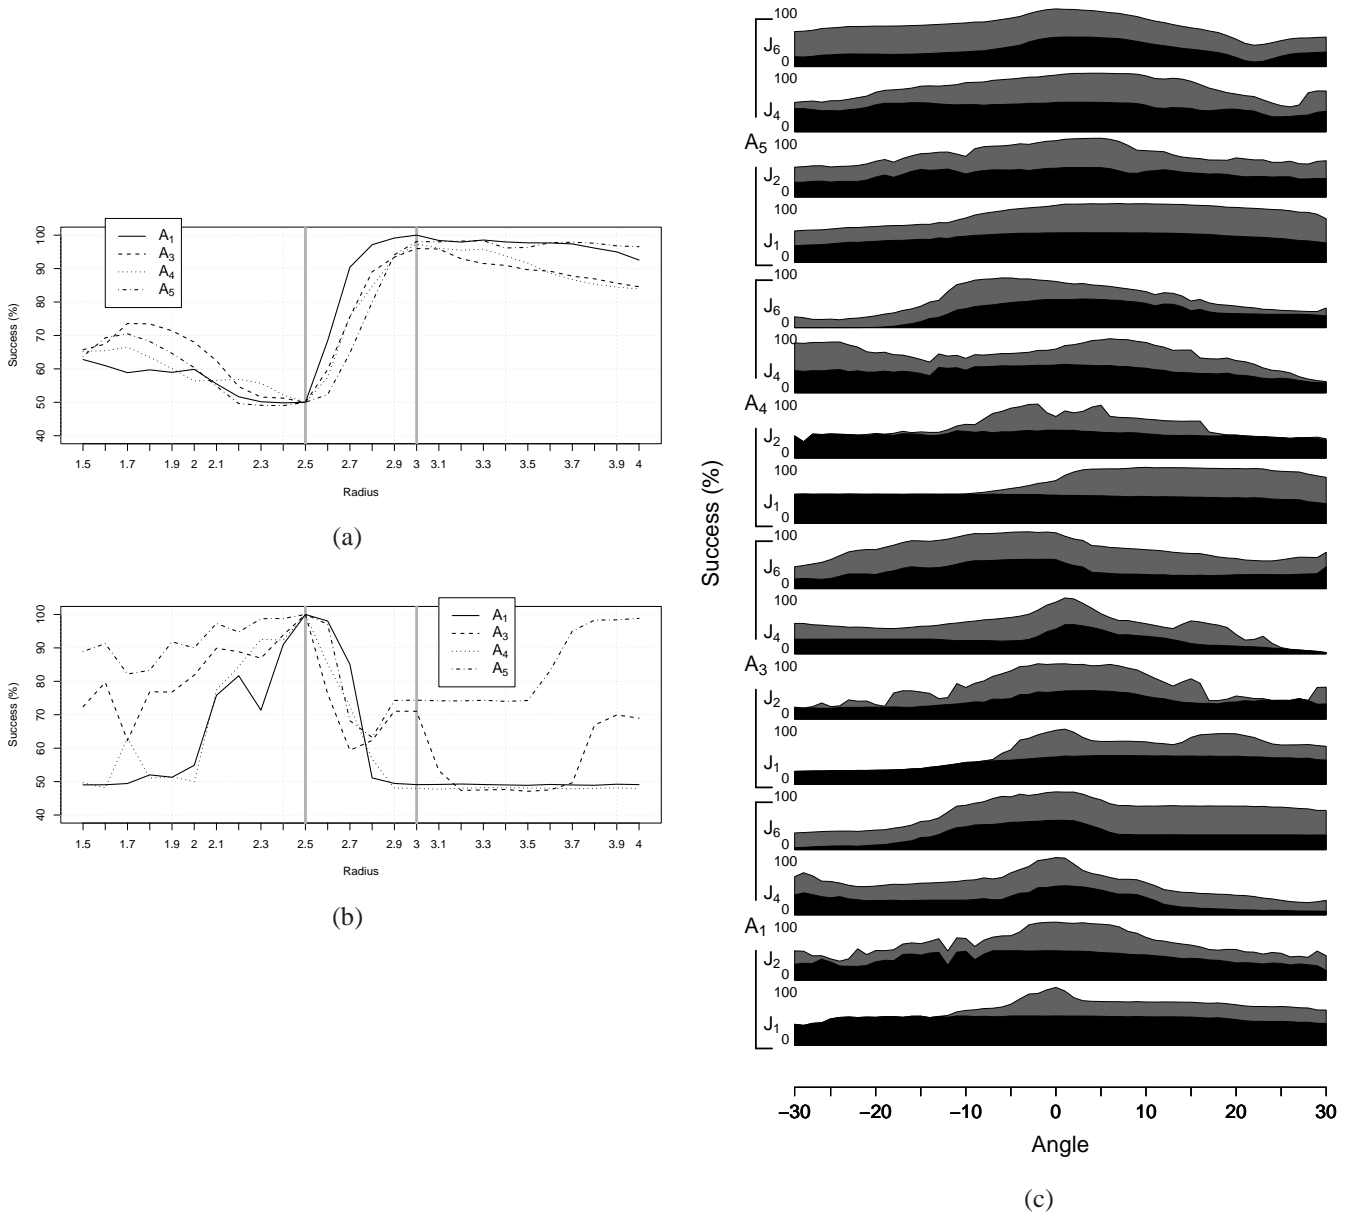


Fig. 4: Graphs showing the percentage of success in post-evaluation tests in which (a) the length of the longest radius of the ellipsoid progressively increases/decreases; (b) the length of the radius of the sphere progressively increases/decreases; (c) the initial position of the object and of the hand varies. Black is for position A, and grey for position B. Note that grey areas extend upward over the black areas. Black and grey never overlap. See also the text for further details. Recall that the original radius of the sphere is 2.5 cm, and the radii of the ellipsoid are 2.5, 3.0, and 2.5 cm.

test P described in Section VII-A). For each ellipsoid initial orientation, each *substitution tests* is repeated 180 times. The rationale behind these tests is that any performance drop caused by the replacement of different type of sensory information provides an indication of the relative importance of that sensory channel on the categorisation process.

The results of this first series of *substitution tests* tell us that, for all the agents, the replacement of the sensory information originated by the arm proprioceptive sensors and by the

hand proprioceptive sensors in position A, only marginally interfere with their performance. That is, for position A, the agents undergo a substantial performance drop only due to replacement of tactile sensation (see Figure 5 black columns in correspondence of tactile sensors). The clear performance drop in these *substitution tests* concerning tactile sensation clearly indicates that, for position A, the agents heavily rely on tactile sensation to distinguish the ellipsoid from the sphere and to correctly perform the categorisation task.






Fig. 5: Graphs showing, for agents  $A_1$ ,  $A_3$ ,  $A_4$ , and  $A_5$ , the results of *substitution tests* concerning the readings of arm proprioceptive sensors, tactile sensors, and hand proprioceptive sensors for position A (see black columns) and for position B (see grey columns).

For position B, the results are slightly more heterogeneous. For agent  $A_1$ , the results of *substitution tests* indicate that both the replacement of tactile sensations and of the hand proprioceptive sensor produce about 20% performance drop (see Figure 5 white columns in correspondence of tactile and hand sensors). For the other agents, tactile sensation keeps on being extremely important for the correct categorisation of the objects (see Figure 5 white columns in correspondence of tactile sensors). However, for agent  $A_4$ , the replacement of the arm and of the hand proprioceptive sensor produces a performance drop of about 40% in the case of the arm and 20% in the case of the hand sensors (see Figure 5 white columns in correspondence of arm and hand sensors). Thus, we conclude that, for agent  $A_1$  the categorisation of the ellipsoid in position B is performed by exploiting information distributed over two sensory channels, that is tactile and hand sensors. The information provided by the two sensory channels seems to be fused together in a way that, for several orientations, the lack or the unreliability of information from one channel can be compensated by the availability of reliable information from the other channel (data not shown). The other agents seem to strongly rely on tactile sensation, with agent  $A_4$  that makes also use of arm and hand sensation to discriminate the objects.

Given that, tactile sensation is the major source of discriminating cues in order to distinguish spheres from ellipsoids in position A, for all the selected agents, and in position B for  $A_3$ , and  $A_5$ , we pursue further investigations, to see whether among the tactile sensors, there are any whose activations play a predominant role in the categorisation task. We begin by running *substitution tests* in which we applied the kind of replacements described above only to single tactile sensors. It turned out that the categorisation abilities of the agents are not hindered by replacements which selectively hit the functioning of single tactile sensors. The performance of all the agents remain largely above 90% success rate (data not shown).

Thus, we proceeded by running *substitution tests* in which




Fig. 6: Graphs showing the results of *substitution tests* concerning the readings  $I_i$  with  $i = 8, \dots, 17$  of all the possible combinations of two elements of the tactile sensors for position A. Each square is coloured in shades of grey. The grey scale is proportional to the percentage of success, with white indicating combinations in which the agent is 100% successful, and black combinations in which the agent is 100% unsuccessful.

we applied replacements to all the possible combinations of two elements of the tactile sensors. Although this analysis have been carried out on all the agents for position A, and on agents  $A_3$ , and  $A_5$ , for position B, in the following we illustrate in details only the results of  $A_1$  (i.e., the best performing agent, see Table I) for position A<sup>1</sup>. The results are shown in Figure 6, in which, the grey scale of the small squares is proportional to the percentage of success, with white indicating combinations in which the agent is 100% successful, and black combinations in which the agent is 100% unsuccessful. These *substitution tests* did not produce clear cut results. However, by looking at Figure 6 we can see that there are specific sensors which, when disrupted in combination with any other sensor, produce a clear performance drop. In particular, disruptions applied to the reading of the tactile sensors placed on the third phalange of the middles finger (i.e.,  $I_{12}$ ), and in minor terms, disruption applied to the reading of the tactile sensors placed on the first phalange of the ring finger (i.e.,  $I_{15}$ ) induce the agent to mistake the ellipsoid for the sphere. We conclude that, agent  $A_1$  heavily relies on the patterns of activation of tactile sensors in which the reading of  $I_{12}$  and  $I_{15}$  are particularly important to distinguish the ellipsoid from the sphere. For what concerns the other agents, the performance of agent  $A_3$  drops in position A when substitutions concern the reading of  $I_{10}$  in combination with any other tactile sensor. In position B, a performance drop is recorded when substitutions concern the reading of  $I_8$  or  $I_{12}$  in combination with any other sensor. Agent  $A_4$  in position A is particularly disrupted by substitutions concerning the reading of  $I_{11}$  or  $I_{12}$  in combination with any other sensor. Agent  $A_5$  in position A is disrupted by substitution concerning the reading of  $I_{12}$  with any other sensor, and of  $I_{12}$  or  $I_{17}$  with any other sensor in position B. In conclusion, in those circumstances in which we observed a predominance of tactile sensation to carry out the categorisation task, the agents tend to rely on combinations

of tactile sensors, with the tactile sensor placed on the third phalange of the middle finger basically more relevant than the other sensors for all the agents (data not shown).

### C. On the dynamics of the categorisation process

In this section, we focus our attention on the dynamics of the categorisation process. More specifically, we analyse: (i) to what extent the sensory stimuli experienced while the agents interact with the objects provide the regularities required to categorise the objects; (ii) to what extent the agents succeed in self-selecting discriminative stimuli (i.e., stimuli that can be unambiguously associated with either category); (iii) how long the agents need to interact with the object before being able to tell whether they are touching a sphere or an ellipsoid; (iv) whether the categorisation process occurs instantaneously by exploiting the regularities provided by single unambiguous sensory patterns or whether it occurs over time by integrating the regularities provided by several stimuli.

To answer these questions we run qualitative and quantitative tests. The former are observations of the trajectories of the categorisation outputs in the two-dimensional categorisation space  $\{\sigma(y(t)_{47} + \beta_{47}), \sigma(y(t)_{48} + \beta_{48})\}$ , in single trials. The latter are tests that further explore the dynamics of the categorisation processes by taking advantage of the fact that in both positions almost all the best evolved agents exploit tactile sensation to carry out the task. The quantitative tests have been carried out on all the agents for position A, and on agents  $A_3$ , and  $A_5$ , for position B. In the following, we illustrate in details only the analysis concerning  $A_1$  (i.e., the best performing agent, see Table I) for position A. However, it turned out that, successful categorisation strategies are very similar from a behavioural point of view, and in terms of the mechanisms exploited to perform the task. Therefore, the reader should consider the operational description of  $A_1$  representative of the categorisation strategies of  $A_3$ ,  $A_4$ , and  $A_5$  in position A, and of  $A_3$  and  $A_5$  in position B<sup>1</sup>.

The first two tests aim at establishing to what extent the stimuli experienced by  $A_1$  during its interactions with the objects provide the regularities required to categorise the objects. We begin our analysis by computing a slightly modified version of the Geometric Separability Index (hereafter, referred to as  $GSI$ ). The  $GSI$ , originally proposed by Thornton [31], is an estimate of the degree to which tactile sensors readings associated with the sphere or with the ellipsoid are separated in sensory space. We built four hundred data sets, one for each time step with the ellipsoid (i.e.,  $\{\tilde{I}_k^E\}_{k=1}^{180}$ ), and four hundred data sets, one for each time step with the sphere (i.e.,  $\{\tilde{I}_k^S\}_{k=1}^{180}$ ). Where,  $\tilde{I}_k^E$  is the tactile sensors reading experienced by the agent while interacting with the ellipsoid at time step  $t$  of trial  $k$ ; and  $\tilde{I}_k^S$  is the tactile sensors reading experienced by the agent while interacting with the sphere at time step  $t$  of trial  $k$ . Recall that, trial after trial, the initial rotation of the ellipsoid around the z-axis changes of  $1^\circ$ , from  $0^\circ$  in the first trial to  $179^\circ$  in the last trial. Each trial is differently seeded to guaranteed random variations in the noise added to sensors readings. At each time step  $t$ , the  $GSI$

is computed in the following:

$$GSI(t) = \frac{1}{K} \sum_{k=1}^K z_k, \text{ with } K = 180;$$

$$z_k = \begin{cases} 1 & \text{if } m^{EE} < m^{ES}; \\ 0 & \text{if } m^{EE} > m^{ES}; \\ \frac{u}{u+v} & \text{otherwise}; \end{cases} \quad (5)$$

$$m^{EE} = \min_{\forall j \neq k} (H(\tilde{I}_k^E, \tilde{I}_j^E))$$

$$m^{ES} = \min_{\forall j} (H(\tilde{I}_k^E, \tilde{I}_j^S))$$

$$u = |\{\tilde{I}_j^E : H(\tilde{I}_k^E, \tilde{I}_j^E) = m^{EE}\}_{\forall j \neq k}|$$

$$v = |\{\tilde{I}_j^S : H(\tilde{I}_k^E, \tilde{I}_j^S) = m^{ES}\}_{\forall j}|$$

where  $H(x, y)$  is the Hamming distance between tactile sensors readings.  $|x|$  means the cardinality of the set  $x$ .  $GSI$  equal to 1 means that at time step  $t$  the closest neighbourhood of each  $\tilde{I}_k^E$  is one or more elements of the set  $\tilde{I}_k^E$ .  $GSI$  equal to 0 means that at time step  $t$  the closest neighbourhood of each  $\tilde{I}_k^E$  is one or more elements from the set  $\tilde{I}_k^S$ . As shown in Figure 7a, for agent  $A_1$  position A, the  $GSI(t)$  tends to increase from about 0.5 at time step 1 to about 0.9 at time step 200, and remains around 0.9 until time step 400. This trend suggests that during the first 200 time steps, the agent acts in a way to bring forth those tactile sensors readings which facilitate the object identification and classification task. In other words, the behaviour exhibited by the agent allows it to experience two classes of sensory states which tend to become progressively more separated in the sensory space. However, the fact that the  $GSI$  does not reach the value of 1.0 indicates that the two groups of sensory patterns belonging to the two objects are not fully separated in the sensory space. In other words, some of the sensory patterns experienced during the interactions with an ellipsoid are very similar or identical to sensory patterns experienced during interactions with the sphere and vice versa.

To analyse in more details to what extent the stimuli experienced by the agent could be associated to the correct or the wrong category we calculated the  $E$ -representativeness. The latter refers to the probability with which a single tactile sensors pattern is associated to the category ellipsoid. The  $E$ -representativeness is computed on a set of 32.400 trials, given by repeating 180 times each the 180 trials corresponding to 180 different ellipsoid initial orientations, from  $0^\circ$  to  $179^\circ$ . During these trials, for each single tactile sensors pattern, we recorded the number of times each pattern appears during interaction with the ellipsoid ( $N$ ) and during interactions with the sphere ( $M$ ). The  $E$ -representativeness of a single pattern is given by  $(\frac{N}{N+M})$ . It is important to notice that an  $E$ -representativeness of 1.0 or 0.0 corresponds to fully discriminative stimuli that can be unambiguously associated with the ellipsoid or the sphere category, respectively, while 0.5  $E$ -representativeness corresponds to fully ambiguous stimuli. The graph in Figure 7b refers to the  $E$ -representativeness of the last 20 patterns (i.e, patterns recorded from time step 380 to time step 400) of single successful trials of test P described in Section VII-A. Each trial refers to a different initial orientation of the ellipsoid. A quick glance at Figure 7b indicates that

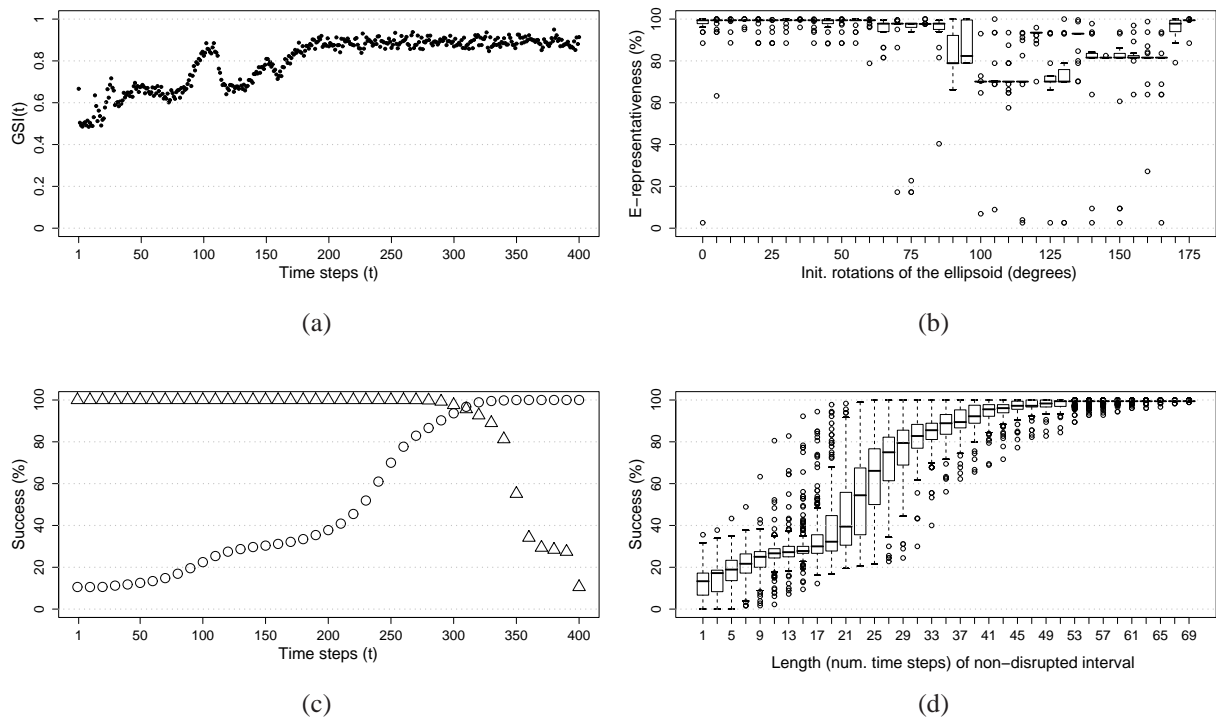


Fig. 7: Graphs showing: (a) the Geometric Separability Index (GSI); (b) the  $E$ -representativeness of the tactile sensors patterns recorded in the last 20 time steps of 180 different trials with the ellipsoid; (c) the percentage of success in *pre-substitution tests* (see triangles) and *post-substitution tests* (see empty circles); (d) the percentage of success at the *window-substitution tests*.

there are trials in which the agent has to deal with tactile sensors patterns that have very low  $E$ -representativeness. That is, they are very weakly associated with the ellipsoid. Patterns with very low  $E$ -representativeness tend to appear in trials in which the initial orientation of the ellipsoid is chosen in the interval  $75^\circ, \dots, 175^\circ$ . These patterns may have at least two not mutually excluding origins: (i) they may come from the fact that the agent is not able to effectively position the object in a way to unequivocally say whether is a sphere or an ellipsoid; (ii) they may be determined by the noise injected into the system. The fact that agent  $A_1$  succeeds in correctly discriminating the category of the objects also during trials in which it does not experience fully discriminating stimuli indicates that the problem is solved by integrating over time the partially conflicting evidences provided by sequences of stimuli. In fact, if the agent employs a reactive strategy (i.e., no need of memory structure), it would be deceived by those sensor patterns, very strongly associated with the sphere, that appear in interaction with the ellipsoid. Under this circumstance an agent that employs a reactive strategy would mistake the ellipsoid for a sphere. Since, in spite of the deceiving patterns, the agent is 100% successful, it looks like the agent is employing a discrimination strategy which uses the dynamic properties of its controller.

Other evidence that supports the integration over time hypothesis come from additional analyses conducted employing further types of *substitution tests*. In particular, we substitute, for a certain time interval, tactile sensors patterns experienced

by  $A_1$  in interaction with the ellipsoid with those experienced in interaction with the sphere. In a first series of tests, referred to as *pre-substitution tests*, substitutions have been applied from the beginning of each trial up to time step  $t$  where  $t = 1, \dots, 400$ . In a second series of tests, referred to as *post-substitution tests*, substitutions have been applied from time step  $t$ , where  $t = 1, \dots, 400$ , to the end of a trial  $t=400$ . Each test has been repeated at intervals of 10 time steps. For agent  $A_1$  position A, the results of *pre-substitution tests* and *post-substitution tests* are illustrated in Figure 7c. This graph shows that, regardless of the rotation of the ellipsoid, pre-substitutions which do not affect the last 100 time steps do not cause any performance drop. For *pre-substitution tests* that involve more than 300 time steps the amount of performance drop is higher for longer substitution periods (see triangles in Figure 7c). Similarly, the agent does not incur in any performance drop if post-substitutions affect less than 100 time steps. For *post-substitution tests* that affect more than the last 100 time steps the amount of performance drop is higher for longer substitution periods (see empty circles in Figure 7c).

By looking at the results of *pre-substitution tests* and *post-substitution tests*, we suppose that the agent is integrating sensory states over time for a certain amount of time around time step 310. In particular, the results shown in Figure 7c seem to indicate that, for what concerns agent  $A_1$  position A, the interactions between the agent and the objects can be divided into three temporal phases that are qualitatively different from the point of view of the categorisation process:




Fig. 8: Graphs showing: (a) trajectories of the decision outputs in the two-dimensional categorisation space ( $\sigma(y(t)_{47} + \beta_{47})$ ,  $\sigma(y(t)_{48} + \beta_{48})$ ), with (a)  $t = 50, \dots, 400$ , recorded in a successful trial with the ellipsoid initially orientated at  $115^\circ$ . Big and small rectangles at 100, 200, 300, and 400 time steps indicate the bounding box of the ellipsoid and sphere category, respectively; (b) the *E-representativeness* of the tactile sensory patterns recorded in a successful trial with the ellipsoid initially orientated at  $115^\circ$ .

(i) an initial phase whose upper bound can be approximately fixed at time step 250, in which the categorisation process begins but in which the categorisation answer produced by the agent is still reversible; (ii) an intermediate phase whose upper bound can be approximately fixed at time step 350, in which very often a categorisation decision is taken on the basis of all previously experienced evidences; and (iii) a final phase in which the previous decision (which is now irreversible) is maintained. The fact that the categorisation decision formed by  $A_1$  during the initial phase is not definitive yet is demonstrated by the fact that substitutions of the critical sensory stimuli performed during this phase do not cause any performance drop (see Figure 7c, triangles). The fact that the intermediate phase corresponds to a critical period is demonstrated by the fact that *pre-substitution tests* and *post-substitution tests* affecting this phase produce a significant performance drop (see Figure 7c). The fact that  $A_1$  takes an ultimate decision during the intermediate phase is demonstrated by the fact that *post-substitution tests* affecting the last 80 time steps, approximately, do not produce any drop in performance (see Figure 7c, empty circles).

In a further series of tests, we looked at whether there is and eventually how big it is the hypothesised temporal phase in which the agent is supposed to integrate tactile sensors states. To look at this issue, we employ the *window-substitution tests*. In these tests, substitutions are applied before and after a temporal window centred around time step 310. The length of the temporal window with no substitutions can vary from 1 time step (i.e., no substitution at time step 310) to 69 time steps (i.e., no substitution from time step 276 to 344). As shown in Figure 7d, the wider the window with no substitutions the higher the performance of the agent, with 100% success rate when no substitutions are applied to a temporal phase of about 50 time steps or longer. Although the graph in Figure 7d does not exclude the possibility that the agent employs an instantaneous categorisation process, the

graph seems to suggest that the performance of the agent is in a way correlated to the amount of empirical evidences it manages to gather over time starting from about time step 270 until time step 340.

Finally, additional evidence in support of a dynamic categorisation process based on the integration of tactile sensation over time come from a qualitative analysis of the trajectories of the categorisation outputs in the two-dimensional categorisation space  $\{\sigma(y(t)_{47} + \beta_{47}), \sigma(y(t)_{48} + \beta_{48})\}$ , in single trials. Figure 8a shows the trajectory recorded by  $A_1$  in a trial in which the initial orientation of the ellipsoid was  $115^\circ$ . As we can see,  $A_1$  moves rather smoothly in the categorisation space by reaching in slightly less than 2 s (200 time steps) the corresponding bounding box. If we now look at Figure 8b, we see that during the interaction with the ellipsoid  $A_1$  experiences: (i) few stimuli with a high percentage of *E-representativeness* (i.e., stimuli that are experienced in interaction with an ellipsoid object most of the times); (ii) several stimuli with an intermediate level of *E-representativeness* (i.e., stimuli that are experienced in interaction with the ellipsoid and the sphere in about the 3/4 and 1/4 of the cases, respectively); and (iii) few stimuli with a low percentage of *E-representativeness* (i.e., stimuli that are experienced in interaction with a sphere object most of the times). If we visually compare Figure 8a with Figure 8b, we notice that the experienced sensory patterns with different percentage of *E-representativeness* appear to drive the categorisation output in different regions of the the categorisation space, corresponding to the ellipsoid and the sphere bounding box, respectively. The final position of the categorisation output (i.e., the categorisation decision) therefore is not determined by a single or few selected patterns but is rather the result of a process extended over time in which partially conflicting evidence provided by the experienced tactile sensation is integrated over time. Similar dynamics have been observed by inspecting all other trials. Given this evidence, we conclude that the performance of all best evolved

agents in position A, and of agent  $A_3$  and  $A_5$  in position B, is the result of a dynamic categorisation process based on the integration of tactile sensation over time.

### VIII. DISCUSSION AND CONCLUSIONS

In this paper, we described an experiment in which a simulated anthropomorphic robotic arm acquires an ability to categorise un-anchored spherical and ellipsoid objects placed in different positions and orientations over a planar surface. The agents neural controller has been trained through an evolutionary process in which the free parameters of the 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. This implies that the robots are left free to determine (i) how to interact with the external environment (by eventually modifying the environment itself); (ii) how the experienced sensory stimuli are used to discriminate the two categories; and (iii) how to represent in the categorisation space each object category.

The analysis of the obtained results indicates that the agents are indeed capable of developing an ability to effectively categorise the shape of the objects despite the high similarities between the two types of objects, the difficulty of effectively controlling a body with many DoFs, and the need to master the effects produced by gravity, inertia, collisions etc. More specifically, the best individuals display an ability to correctly categorise the objects located in different positions and orientations already experienced during evolution, as well as an ability to generalise their skill to objects positions and orientations never experienced during evolution. Moreover, the agents are robust enough to deal with categorisation tasks in which the longest radius of the ellipsoid is progressively increased. Other distortions on the original objects dimensions result more disruptive. These results prove that the method proposed can be successfully applied to scenarios which appear to be more complex than those investigated in previous works based on similar methodologies.

The analysis of the best evolved agents indicates that one fundamental skill that allows them to solve the categorisation problem consists in the ability to interact with the external environment and to modify the environment itself so to experience sensory states which are progressively more different for different categorical contexts. This result represents a confirmation of the importance of sensory-motor coordination, and more specifically of the active nature of situated categorisation, already highlighted in previous studies [e.g., 20, 23]. On the other hand, the fact that sensory-motor coordination does not allow the agents to experience fully discriminative stimuli demonstrates how in some cases sensory-motor coordination should be complemented by additional mechanisms. Such mechanism, in the case of the best evolved individuals, consists in an ability to integrate the information provided by sequences of sensory stimuli over time. More specifically, we brought evidence showing that agent  $A_1$  categorise the current object as soon as it experiences useful regularities and that the categorisation process is realised during a significant period of time (i.e., about 50 time steps) in which the agent keeps

using the experienced evidence to confirm and reinforce the current tentative decision or to change it. Similar strategies have been observed in the other three best evolved agents (data not shown<sup>1</sup>). On this aspect see also [22, 33, 34].

The importance of the ability to integrate the regularities provided by sequences of stimuli is also confirmed by the results obtained in a control experiment, replicated 10 times, in which the agents were provided with reactive neural controllers (i.e., neural networks without recurrent connections, with simple logistic internal neurons, and in which all other parameters were kept equal to those described in Section IV). Indeed the performance displayed by the best evolved individuals in this control experiment were significantly worse than those observed in the basic experiment in which the agents were allowed to keep information about previously experienced sensory states (data not shown<sup>1</sup>). Although we cannot exclude that different experimental scenarios (e.g., scenarios involving agents provided with different neural architecture and/or different physical characteristics of the agents) could lead to qualitatively different results, the analysis of the results obtained in this specific scenario overall indicates that the task does not admit pure reactive solutions or alternatively that such solutions are hard to synthesise through an evolutionary process. This may also be due to functional constraints which limit the movements of the robotic arm (e.g., the fact that the fingers can not be extended/flexed separately, or that there was no adduction/abduction movement of the fingers), as well as other implementation details (e.g., the dimensions of the objects with respect to the hand). This issue will be definitely investigated in future works.

The analysis of the role played by different sensory channels indicates that the categorisation process in the best evolved individuals is primarily based on tactile sensors and secondarily on hand and arm proprioceptive sensors (with arm proprioceptive sensors playing a role only for agent  $A_4$  position B, see Figure 5). It is interesting to note that at least one of the best evolved agents (i.e.,  $A_1$ ) does not only display an ability to exploit all relevant information but also an ability to fuse information coming from different sensory modalities in order to maximise the chance to take the appropriate categorisation decision [see also 32]. More specifically, the ability to fuse the information provided by the tactile and hand proprioceptive sensors, for objects located in position B, allows the robot to correctly categorise the shape of the object in the majority of the cases even when one of the two sources of information is corrupted (see Figure 5).

For the future, we plan to validate the obtained results by porting the best evolved controller on the I-CUB humanoid robotic platform [see 35]. Note that, the porting may require only few changes. In particular, while structurally the simulated arm described in Section III is identical to the real I-CUB, from the functional point of view, it may not match the dynamics of the tendon actuators moving the arm of the real I-CUB. The simulation-reality gap can be closed by firstly quantitatively estimating the mismatch between simulation model and real robot and by appropriately adjusting the system to undo this mismatch. Moreover, we plan to scale up the experiment to a larger number of object

categories, and to study experimental scenarios in which the robots are rewarded for the ability to perform a manipulation task (e.g., grasping different type of objects) that presumably requires categorisation rather than directly for the ability to perceptually categorise the shape of the objects.

#### ACKNOWLEDGEMENT

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

#### REFERENCES

- [1] S. Harnad, Ed., *Categorical Perception: The Groundwork of Cognition*. Cambridge University Press, 1987.
- [2] H. Cohen and C. Lefebvre, Eds., *Handbook of Categorisation in Cognitive Science*. Elsevier, 2005.
- [3] R. Beer, "Dynamical approaches to cognitive science," *Trends in Cognitive Sciences*, vol. 4, pp. 91–99, 2000.
- [4] 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.
- [5] J. J. Gibson, "The theory of affordances," in *Perceiving, Acting and Knowing. Toward an Ecological Psychology*, R. Shaw and J. Bransford, Eds. Hillsdale, NJ: Lawrence Erlbaum Associates, 1977, ch. 3, pp. 67–82.
- [6] A. Noë, *Action in Perception*. MIT Press, Cambridge, MA, 2004.
- [7] R. Pfeifer and C. Scheier, *Understanding Intelligence*. MIT Press, Cambridge, MA, 1999.
- [8] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA, 2000.
- [9] 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.
- [10] D. Floreano, P. Husband, and S. Nolfi, "Evolutionary robotics," in *Springer Handbook of Robotics*, B. Siciliano and O. Khatib, Eds. Springer Verlag, Berlin, Germany, 2008, pp. 1423–1451.
- [11] S. Takamuku, G. Gomez, K. Hosoda, and R. Pfeifer, "Haptic discrimination of material properties by a robotic hand," in *Proceedings of the IEEE 6th International Conference on Development and Learning (ICDL 2007)*, 2007, paper nr 76.
- [12] M. Johnsson, R. Pallbo, and C. Balkenius, "Experiments with haptic perception in a robotic hand," in *Advances in Artificial Intelligence in Sweden*, P. Funk, T. Rognvaldsson, and N. Xiong, Eds. Västerås, Sweden: Mälardalen University, 2005, pp. 81–86.
- [13] M. Johnsson and C. Balkenius, "A robot hand with tpsom neural networks in a model of the human haptic system," in *Proceedings of the International Conference Towards Autonomous Robotic Systems*, M. Witkowski, U. Nehmzow, C. Melhuish, E. Moxey, and A. Ellery, Eds. Springer Verlag, Berlin, Germany, 2006, pp. 80–87.
- [14] —, "Experiments with proprioception in a self-organizing system for haptic perception," in *Proceedings of the International Conference Towards Autonomous Robotic Systems*, M. Wilson, F. Labrosse, U. Nehmzow, C. Melhuish, and M. Witkowski, Eds. Springer Verlag, Berlin, Germany, 2007, pp. 239–245.
- [15] —, "Neural network models of haptic shape perception," *Robotics and Autonomous Systems*, vol. 55, pp. 720–727, 2007.
- [16] P. Dario, C. Laschi, C. Carrozza, E. Guglielmelli, G. Teti, B. Massa, M. Zecca, D. Taddeucci, and F. Leoni, "An integrated approach for the design of a grasping and manipulation system in humanoid robotics," in *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 1, 2000, pp. 1–7.
- [17] L. Natale and E. Torres-Jara, "A sensitive approach to grasping," in *Proceedings of the 6th International Workshop on Epigenetic Robotics*, F. Kaplan, P. Oudeyer, A. Revel, P. Gaussier, J. N. L. Berthouze, H. Kozima, C. Prince, and C. C. Balkenius, Eds., vol. 128. Lund University Cognitive Studies, Lund, Denmark, 2006, pp. 87–94.
- [18] S. Stansfield, "A haptic system for a multifingered harnd," in *Proceedings of the IEEE International Conference on Robotics and Automation*, 1991, pp. 658–664.
- [19] C. Scheier and D. Lambrinos, "Categorization in a real-world agent using haptic exploration and active perception," in *Proceedings of the 4th International Conference on Simulation of Adaptive Behavior (SAB96)*, P. Maes, M. Mataric, J. Meyer, J. Pollack, and S. Wilson, Eds. MIT Press, Cambridge, MA, 1996, pp. 65–74.
- [20] C. Scheier, R. Pfeifer, and Y. Kuniyoshi, "Embedded neural networks: exploiting constraints," *Neural Networks*, vol. 11, no. 7-8, pp. 1551–1596, 1998.
- [21] S. Nolfi, "Power and limits of reactive agents," *Neurocomputing*, vol. 42, pp. 119–145, 2002.
- [22] R. Beer, "The dynamics of active categorical perception in an evolved model agent," *Adaptive Behavior*, vol. 11, pp. 209–243, 2003.
- [23] S. Nolfi and D. Marocco, "Active perception: A sensorimotor account of object categorisation," in *Proc. of the 7<sup>th</sup> International Conference 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.
- [24] T. Buehrmann and E. D. Paolo, "Closing the loop: Evolving a model-free visually-guided robot arm," in *Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems*, J. Pollack, M. Bedau, P. Husbands, T. Ikegami, and R. Watson, Eds. MIT Press, Cambridge, MA, 2004, pp. 63–68.
- [25] E. Tuci, V. Trianni, and M. Dorigo, "Feeling the flow of time through sensory-motor coordination," *Connection*

- Science*, vol. 16, no. 4, pp. 301–324, 2004.
- [26] O. Gigliotta and S. Nolfi, “On the coupling between agent internal and agent/ environmental dynamics: Development of spatial representations in evolving autonomous robots,” *Adaptive Behavior*, vol. 16, pp. 148–165, 2008.
- [27] G. Massera, A. Cangelosi, and S. Nolfi, “Evolution of prehension ability in an anthropomorphic neurorobotic arm,” *Front. Neurobot.*, vol. 1, pp. 1–9, 2007.
- [28] R. Beer and J. Gallagher, “Evolving dynamical neural networks for adaptive behavior,” *Adaptive Behavior*, vol. 1, no. 1, pp. 91–122, 1992.
- [29] D. Goldberg, *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley, 1989.
- [30] S. Strogatz, *Nonlinear Dynamics and Chaos*. Perseus Books Publishing, 2000.
- [31] C. Thornton, “Separability is a learner’s best friend,” in *Proc. of the 4<sup>th</sup> Neural Computation and Psychology Workshop: Connectionist Representations*, J. Bullinaria, D. Glasspool, and G. Houghton, Eds. Springer Verlag, London, UK, 1997, pp. 40–47.
- [32] A. Waxman, “Sensor fusion,” in *Handbook of brain theory and neural networks*, 2nd ed., M. Arbib, Ed. MIT Press, Cambridge, MA, 2002, pp. 1014–1016.
- [33] J. Townsend and J. Busemeyer, “Dynamic representation of decision-making,” in *Mind as motion: Explorations in the dynamics of cognition*, R. Port and T. van Gelder, Eds. MIT Press, Cambridge, MA, 1995, pp. 101–120.
- [34] M. Platt, “Neural correlates of decisions,” *Current Opinion in Neurobiology*, vol. 12, pp. 141–148, 2002.
- [35] G. Sandini, G. Metta, and D. Vernon, “The icub cognitive humanoid robot: An open-system research platform for enactive cognition,” in *50 Years of Artificial Intelligence*, M. Lungarella, F. Iida, J. Bongard, and R. Pfeifer, Eds. Springer Verlag, Berlin, GE, 2007, pp. 358–369.