Accepted on Cognitive Science, November 2013

Learning to Manipulate and Categorize in Human and Artificial Agents

Giuseppe Morlino[1, 3], Claudia Gianelli[2], Anna M. Borghi[2, 1], Stefano Nolfi[1]

[1] Institute of Cognitive Sciences and Technologies (ISTC) – CNR, Rome, Italy

[2] Department of Psychology, University of Bologna, Italy

[3] Department of Computer Science, University Sapienza of Rome, Italy

{giuseppe.morlino, stefano.nolfi}@istc.cnr.it, {claudia.gianelli3, annamaria.borghi}@unibo.it

Keywords: embodied cognition – categorization – category and action learning – shape bias – neural networks.

Classifications:

Fields: Artificial Intelligence, Psychology.

Topics: Intelligent agents, Situated cognition, Skill acquisition and learning.

Methods: Computer simulation, Neural Networks, Human experimentation.

Abstract. This study investigates the acquisition of integrated object manipulation and categorization abilities through a series of experiments in which human adults and artificial agents were asked to learn to manipulate two-dimensional objects that varied in shape, color, weight, and color intensity. The analysis of the obtained results and the comparison of the behavior displayed by human and artificial agents allowed us to identify the key role played by features affecting the agent/environment interaction, the relation between category and action development, and the role of cognitive biases originating from previous knowledge.

1. Introduction

The aim of this study was to investigate how humans acquire and use categories while manipulating novel objects. Given the fundamental role played by categorization in behavioral and cognitive systems, a huge literature has already investigated the perceptual, cognitive, and neurobiological processes that mediate this crucial skill. In this paper we focus on *category learning*, i.e. on the acquisition of the ability to categorize unfamiliar stimuli (Ashby & Maddox, 2005), and on *category learning and use*, i.e. on the acquisition of categories to accomplish a specific function (object manipulation). For a discussion of the implications of category use in category learning see Markman and Ross (2003).

In this study, human and artificial agents learned to categorize objects indirectly, through interactions with these objects, rather than directly, as it happens in classification tasks in which agents are explicitly required to associate items to categories. The reason behind this choice came from considering that living organisms must frequently learn object categories that are relevant for specific goal-directed actions. For example, in most contexts humans categorize a bottle as an object to drink from but in other contexts the same bottle can be categorized as a tool for keeping a door open (Crajé, Lukos, Ansuini, Gordon, & Santello, 2011). Along this line, numerous authors have argued that categorization is grounded in the sensorimotor system and action based (Glenberg, 1997; Wilson, 2002; Borghi, 2005; Gallese & Lakoff, 2005; Barsalou, 2002). One notable example of the influence of action on categorization is provided by Smith (2005). Other authors have claimed, in more general terms, a central role of action in cognition (e.g., Nöe, 2004; Pfeifer & Schieier, 1999). We framed our experiments within this theoretical view and explored the relationship between action and category learning by having agents perform specific goal-directed behaviors with objects and categorize these objects through these behaviors.

This study also addressed the issue of how categorization enables agents to display different behaviors in functionally different agent/environment contexts and to use the same objects for different purposes. Indeed, the role played by context has become increasingly relevant in recent literature on affordances. While previous studies mainly focused on the automatic activation of affordances, independently from the task (e.g., Tucker & Ellis, 2001), more recent studies have pointed out the flexibility of affordances showing that their activation is modulated by both the task and the physical and social context (e.g., Borghi, Flumini, Natraj, & Wheaton, 2012; Costantini, Ambrosini, Scorolli, & Borghi, 2011; Ellis, Swabey, Bridgeman, May, Tucker, & Hyne, 2011; Tipper, Paul, & Hayes, 2006; Mizelle & Wheaton, 2010; Natraj, Poole, Mizelle, Flumini, Borghi, & Wheaton, 2013; Yoon, Humphreys, & Riddoch, 2010). Along the same line, kinematics studies have revealed that humans vary the way they grasp an object not only depending on its physical features, such as shape and weight, but also on its content, on the context and on the end-goal of the action (e.g., Ansuini, Tognin, Turella, & Castiello, 2007; Ansuini, Giosa, Turella, Altoè, & Castiello, 2008; Crajé et al., 2011; Lederman & Wing, 2003; Sartori, Straulino, & Castiello, 2011).

The study was conducted by carrying out parallel experiments on human and artificial agents (Vassie & Morlino, 2012). This direct comparison was made to verify the extent to which observed data can be accounted for by a simple model implemented in experiments with artificial agents and to clear up the role played by humans' previous knowledge. In particular, we investigated the degree to which humans rely on online affordances (Gibson, 1979) and the extent to which perception of affordances re-activates previously learned associations between visual properties and motor responses acquired through experience (Ellis & Tucker, 2000).

To target these issues, we designed a software application that presented an agent, human or artificial, with a set of 2D objects. The agent could move the objects, one at a time, and received a numerical feedback that indicated how well the object was manipulated. The objects varied with respect to four features: two features were directly perceivable, irrespectively of whether and how the object was manipulated, while the other two were action-dependent, i.e. they could be sensed only through direct manipulation (see section 2 for more details). Furthermore, two features, one directly perceivable (shape) and one action-dependent (weight), defined the desired manipulation, while the other two were distractors. Objects were divided into *families*, requiring similar manipulations, and *categories*, requiring specific manipulations (e.g. the family of square objects had to be shaken but the categories of light-square and heavy-square objects had to be shaken vertically and horizontally, respectively). Therefore, in order to master the task, the agent had to learn the required manipulations and the way in which objects were grouped.

Our predictions were as follows:

Embodiment. We hypothesized that the physical characteristics of the objects and the effects of interaction between the agent and the environment could facilitate the development of action-oriented categorization, i.e. of integrated manipulation and categorization skills (Poirier, Hardy-Vallee, & De Pasquale, 2008; Tuci, Massera, & Nolfi, 2010). More specifically, we predicted that the discriminative features that affect the agent/environment interaction, such as weight, might facilitate the acquisition of differentiated manipulations with respect to discriminative features, such as shape, that do not influence the effect of the actions performed by the agents.[1]

Previous knowledge. We hypothesized that previous knowledge might affect learning processes, i.e. biases the type of manipulations/categorization performed by the learning agents. An example of this process is constituted by the *shape-bias* effect, i.e., the tendency to extend count-nouns on the basis of shape similarity, a phenomenon present in Western societies and becoming rather stable quite early, in 2 year-old children (Landau, Smith, & Jones, 1988, 1998; Samuelson & Smith, 1999). In addition, affordances have been proposed to be the product of the conjunction, in the brain, of repeated visuomotor experiences (Ellis & Tucker, 2000). Thus, we expected that for humans a given shape should elicit a specific movement (e.g. circular movement for circles), while this should not be the case for artificial agents, which lack any previous knowledge.

[1] Whether a feature is interactive or not also depends from the task/scenario. In our scenario the shape of the object does not influence the effect of the agent's actions. In an alternative 3D manipulation scenario, the shape of the objects could influence actions since, for example, spherical objects will tend to roll. Indeed, recent research on 2-year-old children has revealed that actions performed on objects modify the way their shape is perceived and lead them to categorize differently (Smith, 2005). We intentionally selected a task/scenario with an interactive and non-interactive predictive feature (weight versus shape) to disentangle the role of interactivity in category learning and use.

2. Methods

2.1. Experimental scenario

The experimental scenario involved 16 two-dimensional objects that were generated by combining four binary perceptual features: shape (circle/square), color (red/green), weight (light/heavy), and blinking (varying or not-varying color intensity during motion). These 16 objects were grouped into four categories that required four corresponding manipulations. Categories were determined on the basis of the combination of the shape and weight features: circle-light, circle-heavy, square-light and square-heavy. Each category included four objects that varied with respect to the other two perceptual features; more precisely, each category included a red-blinking, a red-non-blinking, a green-blinking, and a green-non-blinking object. However, variations with respect to these features were functionally irrelevant. Two of the perceptual features (shape and color) were directly perceivable, while two (weight and blinking) were action-dependent (meaning that they could be perceived only by manipulating the object and by integrating perceived information over time). In this sense, light/heavy objects varied with respect to the inertia with which they reacted to the movements. On the other hand, blinking/non-blinking objects differed with respect to whether their color varied or remained the same during motion. It is worth noting that weight is the only feature that influenced the agent/environment interaction; indeed, objects of different weight responded differently to the same action, while the other three features only affected the state of the agent's sensors.

Agents were allowed to interact repeatedly with each of the 16 objects and at the end of each period of interaction they received a score in the range [0, 100] that indicated how well the object had been manipulated, i.e. how close the performed manipulation was to the desired one. Each object was associated with one of four desired manipulations (see right part of Fig. 1), which were: place left, place right, shake vertically, and shake horizontally. These four desired manipulations were divided into two families: objects to be placed (independently of where they had to be placed) and objects to be shaken (independently of how they had to be shaken). For a detailed description of how the behaviors exhibited by the agents were scored, see appendix A.

----------------------------- Figure 1 about here ------------------------------

Across trials in the experiment, the assignment of desired manipulations to object categories was fixed. This assignment was manipulated between subjects: human participants and artificial agents were randomly assigned to one of the conditions and performed the experiment in this condition only. We provided four experimental conditions systematically varying the relation between objects and desired manipulations (see Fig. 1), thus we had: place-circles–shake-squares, place-squares–shakes-circles, place-heavy–shake-light, and unstructured, which consisted in placing light circles and heavy squares, and shaking light squares and heavy circles. Notice how in the first three conditions the type of desired manipulation, namely whether the objects should be shaken or placed, could be determined on the basis of a single perceptual feature. In the unstructured condition, instead, both the type of manipulation, placing or shaking, and the specific manipulation. In other words, horizontally and vertically shaking could only be determined by taking into account the combination of two perceptual features.

Data analysis of performance in these four experimental conditions allowed us to: i) verify the role of cognitive biases in human participants; ii) investigate the role of action-dependent versus directly perceivable object features; iii) verify the role of features that affect versus those that do not affect the agent/environment interaction; iv) study how agents can develop integrated action and categorization capabilities.

2.2. Experiments with human participants

20 participants took part in the experiment (7 males, 13 females) and were equally divided into the four experimental conditions described above, 5 participants for each experimental condition. The average age of the participants was 21.34 (SD = 3.71) and they all self-reported to be right-handed.

Participants carried out the experiment at a computer terminal and were asked (see appendix B) to manipulate the objects with the mouse trying to maximize the received scores. Each object could be manipulated for a maximum of 30 s. The instructions informed the participants that after each manipulation feedback would be displayed rating the extent to which she/he approximated the correct desired manipulation for the current object. The participants were also told that they could skip to the next object by clicking on the left button of the mouse or could re-try with the same object by clicking the right button. Moreover, the participants were informed that the experiment consisted of a 25 minutes training phase in which she/he was asked to learn to properly manipulate the objects and of a 5 minutes testing phase in which she/he was asked to maximize the scores. Objects were presented randomly and participants were not aware of the number and type of desired manipulations. During the experiment we recorded the position of the mouse and of the object over time, the state of the mouse buttons, and the scores. At the end of the two phases the experimenter interviewed the participants (questionnaire reported in the appendix B).

2.3. Experiments with artificial agents

The artificial agents were constituted by three-layer recurrent neural networks (see Fig. 2) that, every 100 ms, received the perceptual properties and the position of the object as input and calculated the movement of the mouse pointer as output. The input layer had four sensory neurons that encoded the current position of the mouse pointer and of the centroid of the object, two sensory neurons that encoded (as binary features) whether the current object had a circular or squared shape, and three sensory neurons that encoded the current RGB color of the object. The sensory neurons projected connections to the internal neurons with recurrent connections. The internal neurons, which included three dynamical neurons (see Gigliotta & Nolfi, 2008), projected connections to the motor neurons. The two motor neurons were standard logistic neurons that encoded the displacement (mapped into the range [0, 100] pixels) of the mouse pointer over the horizontal and vertical dimension.

------------------------------ Figure 2 about here -------------------------------

The agents were trained using a trial and error process in which the free parameters were varied randomly; variations were retained or discarded depending on whether or not they led to improved average performance over the whole set of objects. This was done by using an evolutionary method (Nolfi & Floreano, 2000). More specifically, the initial population consisted of 100 randomly generated genotypes that encoded the connection weights, the biases, and the time constants of 100 corresponding neural controllers. Each parameter was encoded by eight bits and normalized in the range [-5.0, +5.0] in the case of the connection weights and biases and in the range [0.0, 1.0] in the case of the time constants of the internal neurons. The 30 best genotypes of each generation were allowed to reproduce by generating five copies each, with 2% of their bits replaced with a new randomly selected value. The evolutionary process lasted 500 generations and was repeated 30 times, starting from different randomly generated parameters, for each experimental condition. The performance of the individuals was computed by averaging the score obtained during 64 trials during which they were allowed to interact four times with each of the 16 objects for 100s. The object position was randomly set at the beginning of each trial. The reason behind this choice is that evolutionary algorithm is one of the simplest yet effective ways to train a neural network through a trial and error process on the basis of a distal reward.

2.4. Differences between human and artificial agents

Although the two experimental settings were designed to be as similar as possible, some important differences should be considered in comparing the results. Even though compared to these artificial agents humans are extremely complex in their sensory, motor, and cognitive apparatus, here we want to stress two specific differences that we have to consider to correctly interpret the data.

The first difference is that the artificial agents had no previous knowledge since the free parameters that determined their behavior at the beginning of the training process were set randomly. By contrast, humans possess highly structured capacities and knowledge. More specifically, as mentioned above, human participants have cognitive biases, deriving from their previous experience, which might affect the way in which they interacted with different objects at the beginning and/or during the course of the learning process. Incidentally, the comparison of obtained results can be used to verify the strength of these biases and the positive or negative role that they have in the development of new skills that are not correlated with previously developed abilities.

A second difference concerns the nature of learning processes. The artificial agents learned through a stochastic trial and error process conducted over an extended period of time (by stochastic we refer to the fact that variations were introduced randomly). On the contrary, humans instead were exposed to a minor number of trials but rely on a variety of learning strategies developed during their life.

The aim of the experiments with artificial agents was not to model human behavior but rather to: i) identify commonalities in behavior that can be attributed to general properties of the agent/task and that arise independently of the differences between human and artificial agents, and ii) investigate the differences in the behavior observed in the two cases that in turn provides a way to detect the role played by aspects that are simply modeled in the artificial agents.

3. Results

Overall the analysis of the performance showed that both human and artificial agents managed to master the task. They both learned to manipulate the objects in ways that allowed them to achieve good performance, defined as a score above 70 (maximum score being 100). More specifically, human agents reached an average score of 71.08 at the end of the training section (lasting 25’, as described above), while artificial agents instead reached an average score of 91.21 at generation 50, 96.82 at generation 200, and close-to-optimal score (98.57) at generation 500 (see Fig. 3 and 4). Artificial agents displayed better performance in the condition in which the predictive feature was the weight (place-heavy–shake-light) than in the conditions in which it was the shape (place-circles–shake-squares and place-squares–shake-circles) or in which there was no predictive feature at all (unstructured). Human data also displayed a similar trend; however participants had relatively better performance when the coupling between the object shape and the required movements was congruent with biases deriving from their previous experiences (the place-squares–shake-circles condition).

The behavioral analysis performed on trained agents showed that both human participants and artificial agents adopted overgeneralized strategies to solve the task: they manipulated objects of two or three different categories or all the objects in the same way. However, as we will see, this did not necessarily lead to sub-optimal performance.

Although the way in which scores were calculated did not explicitly encourage the production of minimal trajectories, humans (in contrast with artificial agents) tended to reduce the number of sub-movements performed for objects to be placed during the course of the training process.

3.1. Statistical analysis

ANOVAs were carried out on the scores received by human and artificial agents.

We first performed an ANOVA on human performance in the testing phase and artificial agents at Generation 50 using Participant Type (human or artificial) and Condition (1: place-circles–shake-squares, 2: place-squares–shake-circles, 3: place-heavy–shake-light, 4: unstructured) as between-participants factors and Manipulation (place left, place right, shake vertically, shake horizontally) as within-participants factor. Effect sizes were computed as values of eta-squared (η^2).

We then performed two more ANOVAs, one on human performance in testing phase and the other on artificial agents data. The ANOVA on human performance had Condition as between-participants factor and Manipulation as within-participants factor. The ANOVA on artificial agents had Generation (training stages): 1, 50, 200, 500, in addition to Condition and Manipulation, as between-agents factor.

To facilitate the comparison with human performance we reported in Fig. 3 (right) only the data of artificial agents for Condition and Manipulation at Generation 50, with the data on Condition and Generations being reported in Fig. 4.

In order to investigate significant main effects and interactions, Newman-Keuls post-hoc tests were used.

------------------------------ Figure 3 and 4 about here -------------------------------

3.1.1. Human *vs* Artificial

The ANOVA on human performance in testing phase and artificial agent performance at Generation 50 showed a main effect of Participant Type ($F(1, 132) = 57.16$, $MSE = 2.78$, $p < .001$, $\eta^2 = .245$) with artificial agents performing relatively better than humans ($p < .001$).

Condition was also a main effect ($F(3, 132) = 11.06$, $MSE = .54$, $p < .001$, $\eta^2 = .142$). Post-hoc test showed that place-heavy–shake-light and unstructured were overall easier than the placing conditions ($p_s < .001$), which did not significantly differed each other; place-heavy–shake-light was also easier that unstructured ($p = .037$).

A main effect of Manipulation emerged as well ($F(3, 396) = 7.71$, $MSE = .33$, $p < .001$, $\eta^2 = .049$) with shaking horizontally turning out to be the easiest manipulation ($p_s = .002$ with respect to place left and shake vertically, and $p < .001$ with respect to place right).

As regards the interactions both Manipulation X Participant Type and Condition X Participant Type were significant ($p < .001$ and $p = .017$, respectively), while Manipulation X Condition and Manipulation X Condition X Participant Type were not ($p = .227$ and $p = .167$, respectively). Being humans and artificial agents highly dissimilar, the significant interactions likely depend on factors not relevant to this study. For this reason we further analyzed humans and artificial agents performing two more separate ANOVAs. However, this first ANOVA provided a general overview of the result and, showing that the Participant Type X Manipulation and Participant Type X Condition interactions were significant, informed us that human participants and artificial agents performed differently. However, as reported below, we observed interesting similarities between human and artificial agents, which allowed us to deeply understand the phenomena in place.

3.1.2. Human participants

The analysis of human performance showed a main effect of Condition ($F(3, 16) = 4.82$, $MSE = .07$, $p = .014$, $\eta^2 = .475$): post-hoc tests showed that Condition 1 (place-circles–shake-squares) had overall the worst performance ($p = .062$ respect to Condition 2, $p = .015$ respect to Condition 3, and $p = .020$ respect to Condition 4). The other conditions did not significantly differ.

A main effect of Manipulation emerged as well ($F(3, 48) = 5.76$, $MSE = .06$, $p = .002$, $\eta^2 = .184$). Post-hoc tests showed that no difference was present between the Manipulation of placing on the left and placing on the right ($p = .934$), and between shaking horizontally and vertically ($p = .273$). Interestingly, shaking manipulations obtained better performance. More specifically, participants performed better with shaking horizontally than placing right ($p = .030$) and with shaking vertically than both the placing manipulations ($p = .007$ respect to placing left and $p = .005$ respect to placing right). The difference between shaking horizontally and placing left was not significant ($p = .063$).

Condition X Manipulation interaction was not significant ($F(9, 48) = .85$, $MSE = .06$, $p = .570$, $\eta^2 = .082$).

3.1.3. Artificial agents

The analysis of artificial agents' performance showed a significant main effect of Condition ($F(3, 464) = 23.20$, $MSE = 236.75$, $p < .001$, $\eta^2 = .013$). Condition 2 (place-squares–shake-circles) was overall the worst performance, significantly different from Condition 3 and 4 (place-heavy–shake-light and unstructured), $p_s < .001$. Similarly, Condition 1 (place-circles–shake-squares) significantly differed from Condition 3 and 4 ($p_s < .001$). Performance was overall the worst in Condition 1 and 2 and did not differ between the two ($p = .118$).

Similarly, Condition 3 and 4 did not significantly differ (p = .666) and exhibited the overall best performance.

A significant main effect of Generation was present as well (F(3, 464) = 1559.49, MSE = 236.75, p < .001, η^2 = .893). Obviously, the overall worst performance was Generation 1, significantly differing from all the other generations (p_s < .001). In addition, Generation 50 was significantly worse than Generations 200 and 500 (p_s < .001), while these last two did not differ from each other.

The main effect of Manipulation was significant as well (F(3, 464) = 67.57, MSE = 336.16, p < .001, η^2 = .081). Shaking actions were worse than placing actions (all p_s < .001). Shaking vertically was significantly worse than shaking horizontally (p = .004). Similarly, placing right was worse than placing left (p < .001).

Condition X Generation interaction was significant as well (F(9, 464) = 2.82, MSE = 236.75, p = .003, η^2 = .005). Post-hoc test showed that Condition 3 (place-heavy–shake-light) in Generation 50, 200, and 500 and all Conditions in Generations 200 and 500 did not significantly differ (p_s > .05). All the other differences were instead significant (p_s < .012). In other words, optimal or close-to-optimal performance was reached earlier in the place-heavy–shake-light compared to the other Conditions.

Condition X Manipulation interaction was significant too (F(9, 1392) = 9.64, MSE = 336.16, p < .001, η^2 = .035). Post-hoc tests showed there to be an advantage of placing over shaking actions, and of shaking vertically over shaking horizontally. Condition 3 (place-heavy–shake-light) and Condition 4 (unstructured) resulted in better performance (p_s < .001) than Condition 1 (place-circles–shake-squares) and Condition 2 (place-squares–shake-circles) in the case of placing actions (place left and place right).

Generation X Manipulation interaction gave complementary information ($F(9, 1392) = 80.48$, $MSE = 336.16$, $p < .001$, $\eta^2 = .291$). Post-hoc tests showed that only at Generation 1 did shaking Manipulations result in worse performance compared to placing ones ($p_s < .001$), while these differences disappeared in the other Generations.

Finally, Condition X Generation X Manipulations interaction was also significant ($F(27, 1392) = 3.12$, $MSE = 336.16$, $p < .001$, $\eta^2 = .034$).

In addition to the previous results, post-hoc tests showed that Condition 1 and Condition 2 were significantly worse than Condition 3 and 4 only in the placing actions ($p_s < .001$) but not in the shaking ones ($p_s > .05$). As also shown above, shaking actions had an overall performance disadvantage.

3.2. Behavioral analysis

The analysis of the behaviors displayed by trained agents indicated that both human and artificial agents tended to manipulate objects belonging to the same category in the same way. However, as we will see in more detail below, objects of different categories were not always manipulated differently. Both human and artificial agents found and exploited manipulations that could be applied to objects of two or more categories achieving good, or optimal scores. Examples of such overgeneralized strategies were abundant both in human and artificial agents and were observed both during the course and at the end of the training process.

In the case of humans, overgeneralized solutions for objects to be shaken were observed in most participants. In some cases participants constantly moved the mouse pointer along the diagonals of the arena (see Fig. 5a or along trajectories that resembled the shape of the objects (see Fig. 5b and 5c). In other cases, instead, participants simply shook objects in varying directions (see Fig. 5d). Overgeneralized solutions were also observed, although less frequently, for objects to be placed. In these cases all objects were brought roughly at the center of the arena (see Fig. 5e) or at one of the two target areas (see Fig. 5f) gathering a relatively high average score without differentiating the manipulations. In some cases we observed higher-level overgeneralizations, in which all the objects were treated in the same way. In these cases, sub-optimal performance was achieved by exhibiting two manipulations in sequence, by first shaking the objects in varying directions or along the diagonals of the arena and then placing them at the center of the arena. For additional details on the number and frequency of overgeneralized strategies see Tab. 1 and 2.

---------------------- Figure 5, Table 1, and Table 2 about here -----------------------

Overgeneralized strategies also characterized the majority of the solutions found by artificial agents (see Tab. 1 for additional details).

In order to analyze the quality of the behavioral strategy we post-evaluated trained agents by using an *effort* measure. In the case of objects to be placed this was computed by subtracting from 1 the ratio between the length of the shortest trajectory and the length of the actual trajectory. In the case of objects to be shaken this measure was computed by subtracting from 1 the difference between the obtained score and the score that would be obtained if the object had to be shaken along the orthogonal axis, i.e. horizontally instead of vertically, or viceversa. Fig. 6 reports the average *effort* of human and artificial agents computed over successful trials, defined as trials in which the manipulation performance was > .9.

----------------------------- Figure 6 about here ------------------------------

As it can be observed, although the performance criterion did not reward *effort* minimization, humans produced shorter trajectories during training for objects to be placed. Artificial agents, instead, did not decrease but rather increase *effort* during training, for all object types.

3.3. Human self-report

At the end of the experiment human participants were interviewed by mean of a questionnaire (see appendix B).

One of the questions asked which object features they had noticed. Shape was the only feature reported by all 20 participants. Interestingly, color, even if it was not relevant for category membership, was reported by 18 participants, more than weight and blinking, which were reported by 17 and 15 participants, respectively.

Most of participant described shape, color and blinking in the same way; weight instead, was described in different ways (inertia to the movement, delay in reaction time, speed, laziness, stickiness, responsiveness to the participant's willingness), even by the same participant.

Participants were also asked to indicate which of the detected features was relevant to the task. The majority of them reported shape and weight as features relevant or partially relevant to the task and discarded color and blinking. Particularly, shape was considered more relevant with respect to weight in the conditions in which it was predictive. Conversely, weight was considered more relevant than shape in the place-heavy–shake-light condition (see Tab. 3).

Besides, in this condition some participants explicitly stated to have exploited the different reactions of light/heavy objects to accomplish the task.

Data related to the remaining questions were incomplete or not particularly relevant here; we thus decided not to report them.

------------------------------ Table 3 about here -------------------------------

4. Discussion

4.1. Role of object features in categorization and behavior development

The differences in performance in the four experimental conditions (reported in Fig. 3) indicate that the type of features that defined the objects' families significantly affected the ability to learn to master the task.

More specifically, the fact that for artificial agents the place-heavy–shake-light condition (in which the discriminating feature –the weight– affected the agent/object interaction) led to significantly better performance with respect to the place-circles–shake-squares and place-squares–shake-circles conditions (in which the discriminating feature –the shape– did not influence the agent/object interaction) indicates the importance of embodiment in object categorization (see explanation below). A similar trend was observed in humans, although in this case performance in the place-heavy–shake-light condition differed significantly only with respect to performance in the place-circles–shake-square condition (see section 4.3. for a discussion on the reason that may explain this difference). These results confirm the hypothesis that discriminative features affecting the agent/environment interactions, such as weight, facilitate the acquisition of the required categorization abilities with respect to alternative features that are equally informative but that do not affect the outcome of the agent actions. This facilitation effect overcomes the fact that weight is an action-dependent property, namely a feature that cannot be immediately detected by the agent but that can only be inferred by integrating sensory-motor information over time while the agent interacts with the object (Brouwer, Georgiou, Glover, & Castiello, 2006; Jenmalm, Schmitz, Forssberg, & Ehrsson, 2006; Scorolli, Borghi, & Glenberg, 2009). The effect was not observed in an additional set of experiments performed with artificial agents in which blinking was used as discriminative feature instead of weight_(results not shown). This supports an embodied account as it suggests that the effect is due not only to properties that can be detected once objects are manipulated, such as both blinking and weight, but only to properties, such as weight, that co-determine the effects of the agent/environment interaction.

The facilitation effect of weight with respect to shape can be explained by considering the action-oriented nature of the categories postulated in these experiments (the category of an object defines the way in which it should be manipulated). In this situation a feature like weight tends to produce spontaneous differentiation in the way in which objects are manipulated. The differentiation arises from agent/object interaction irrespectively of whether the agent treats objects varying with respect to that feature differently. Indeed, the fact that interaction with light versus heavy objects tends to spontaneously produce more differentiated behaviors than interaction with objects varying in shape can be observed in artificial agents already from the very first phases of the training process (see Fig. 4). In addition, as reported in section 3.3. some participants of the place-heavy–shake-light condition explicitly declared to have exploited the different reactions of light/heavy objects to accomplish the task. In other contexts purely perceptual features might be preferred with respect to features affecting the agent/environment interaction, such as weight.

The explanation of the apparently surprising result that the unstructured condition (place-light-circles-and-heavy-squares–shake-light-squares-and-heavy-circles) leads to significantly better performance in artificial agents with respect to the place-circles–shake-squares and place-squares–shake-circles conditions can be explained by the fact that artificial agents tend to converge on a relatively simple strategy that consists in shaking circles horizontally in the bottom-left area and shaking squares vertically in the center-right area. This strategy allows the agents to exploit, also in the unstructured experimental condition, the effect that weight has in the agent/object interaction, e.g. the fact that heavy objects tend to oscillate less than light ones. Indeed, artificial agents oscillated the mouse pointer horizontally in the bottom-left part of the arena in the same way, regardless of the object weight. The selected amplitude, frequency, and location of the oscillation, combined with the effect of the weight, produced an appropriate horizontal oscillatory movement of light circles and a smaller oscillatory movement inside the appropriate target area of heavy circles. The need to shake and place heavy and light squares in the unstructured condition, on the other hand, forced the agents to (often later) develop an ability to differentiate their actions for objects belonging to these two categories. This could explain why the unstructured condition outperformed the place-circle–shake-square and place-squares–shake-circles conditions only during the final phase of the training process (see Fig. 4).

It is well known that different views of embodiment exist, from more moderate to more radical ones (for discussion on these issues see Borghi & Cimatti, 2010; Chemero, 2009; Chatterjee, 2010; Goldman & De Vignemont, 2009). We define *strong* an embodied approach which assigns a central role to the body for cognition, in contrast with more moderate versions of embodied cognition, which assign less centrality to the body and to bodily states. Proponents of “mild” embodied approaches underline that cognition is “typically grounded in multiple ways, including simulations, situated action and, on occasion, in bodily states” (Barsalou, 2008, p. 619; Pezzulo, Barsalou, Cangelosi, Fischer, McRae, & Spivey, 2011). Our data on weight are more in line with a strong embodied account, as they reveal that properties affecting the agent/environment interaction are more salient than properties that only affect perceptual states (for similar behavioral results obtained with a sorting task, see Iachini, Borghi & Senese, 2008). Indeed, in our setting object weight influences the way objects can be dragged in space: heavy objects are slower than light ones in following the mouse pointer. Due to these characteristics, we believe that the importance of weight emphasizes the central role of action and of the body in constraining cognitive processes (Goldman & De Vignemont, 2009; Glenberg & Gallese, 2012).

4.2. Integration of categorization and action skills

As reported above, the analysis of human and artificial agents revealed that they discovered and exploited overgeneralized behaviors, namely manipulations applied to objects belonging to more than one category, which allowed them to achieve relatively good and sometimes optimal performance. This result is in line with the original idea by Gibson (1979) according to which we do not need to categorize objects in order to respond to their affordances. In other words agents might produce the required differentiated manipulations for objects belonging to different categories without differentiating the way in which they are treated but rather by identifying a single manipulation that, in interaction with objects with different properties, produces the required differentiated manipulations (for another example of such implicit categorization behavior see Nolfi & Marocco, 2002).

The role of overgeneralized strategies and the analysis of the process through which lower and higher level categories are acquired in our experiments also provide interesting evidence that can shed light on the way in which categories are formed. According to one influential account (e.g. Mandler, 2004) children first form global categories, which are more general than those referred to by words. Later, they start differentiating them and forming more specific ones. Other authors instead report evidence showing that infants start learning basic level categories then form more general, superordinate ones (see for example the seminal work by Rosch, Mervis, Carolyn, Gray, Johnson, & Boyes-Braem, 1976). The behavioral analysis conducted on data from the training of artificial agents provided evidence of both processes in the same experimental scenario. We observed both cases in which agents first display an ability to discriminate between higher-level categories, i.e. object to be placed versus objects to be shaken, and cases in which agents first exhibit the acquisition of lower-level categories, i.e. object to be placed on the right versus object to be placed on the bottom-left, and then the acquisition of higher level categories. In more general terms these results confirm the tight interaction between action and categorization.

4.3. Role of cognitive biases

Artificial agents did not have previous experience and therefore initially treated all objects in the same way. The behavior exhibited before training might differ significantly only for objects varying with respect to features that affect the agent/environment interaction (weight). As expected, the conditions that differ with respect to the relation between the shape of the object and the associated desired manipulation (place-circle–shape-squares and place–squares-shake-circles) were equally difficult to master for artificial agents and produced similar results (see Fig. 3, right).

Human participants, on the other hand, were influenced by previous experience, with particular reference to objects varying in shape (as discussed in the introduction). The analysis of the way in which circle and square objects were manipulated during the first phase of the experiment indicated that the behavior was much more varied in the case of human participants than in the case of artificial agents. Moreover, it showed that humans tended to manipulate circles and squares in specific ways, for example i) by moving circles more, and more quickly, than squares, ii) by producing curvilinear trajectories with circles and rectilinear trajectories with squares, or iii) by moving squared objects along walls and placing them in the corners of the arena. In line with our predictions, this demonstrates that, due to their previous experience, participants associated specific object shapes with specific movements types and performed better in placing squares and shaking circles than doing the opposite.

The presence of these biases also explains why in humans, and not in artificial agents, the place-squares–shake-circles condition led to better performance than the place-circles–shape-square condition. Thus, biases can facilitate the acquisition of action/categorization skills that present similarities to already mastered skills.

5. Conclusion

The results allow us to draw some conclusions on the development of categories in humans and artificial agents.

First, our work allowed us to investigate the role played by different properties during learning of novel categories. It is possible to contrast three different accounts, which may alternatively explain the results we obtained concerning the role played by the different properties: sense-based, weak embodied, and strong embodied accounts. According to the first, both artificial and human agents would first form categories on the basis of properties that can be sensed directly; thus, they would categorize primarily on the basis on directly perceivable perceptual features, such as shape and color, rather than on the basis of weight and blinking, which in the present setting could only be perceived by manipulating the object (for weight this is true also in real life). This was clearly not our case. According to a weak embodied account agents should have been particularly sensitive to properties that are perceivable only through manipulation, such as weight and blinking. This was only partially our case, since we found that weight was the most important property for categories formation, but blinking was not. According to a strong embodied account, agents should have formed categories preferentially on the basis of properties that affect agent/environment interaction. Our data support the latter case. Indeed, agents categorize primarily on the basis of weight, the only property that co-determines the effect of the agent actions.

Second, our results give many insights as to the strategies developed by human and artificial agents during learning and suggest how action and category development may interact. The analysis of the course of learning indicates that, in general terms, action and category development are achieved by: i) developing smart sub-optimal overgeneralized behaviors, particularly undifferentiated manipulations applied to two or more groups of objects, which allow the agent to achieve relatively good performance, ii) exploiting the effects of agent/environment interaction, such as identifying non-differentiated manipulations that produce appropriate differentiated behaviors in interaction with objects with different features, and iii) differentiating through learning initially non-differentiated behaviors. Thus, our analysis and results provide evidence of an action oriented categorization process. This process is bidirectional, therefore it is in keeping with two accounts which appear to be only apparently in contrast: i) the theory that postulates that categorization proceeds with the formation of global categories that are then progressively differentiated and refined (Mandler, Bauer & McDonough, 1991), and ii) the theory that postulates that categorization proceeds the other way, from particular to general (Rosch et al., 1976). Overall our experiment supports enactivist theories (e.g., Chemero, 2009; van Elk, Slors, & Bekkering, 2010; Varela, Thompson & Rosch, 1991; Maturana & Varela, 1992; Bateson, 1987) that postulate that cognitive processes, including categorization, are dynamical and emerge from the agent/environment interaction.

A third contribution of this work concerns the possibility of comparing the results obtained from human and artificial agents to disentangle the role played by previous experience/knowledge from that played by learning context.

Finally we would like to stress the advantages of studies in which human and artificial beings are compared. The comparison between human and artificial agents allowed us to formulate interesting hypotheses that can be tested experimentally. As an example, let us consider the limited impact of shape observed in our experiments. We interpreted this result in the framework of a strong embodied account. Indeed, we explained the relevance of weight arguing that it directly impacts interaction with objects, while shape does not. In other words, it is possible that the reason why shape played a minor role in our experiments is that it did not affect the agent/environment interaction. This opens the interesting possibility that the importance of shape for categorization (e.g., Landau et al., 1988; Jones & Smith, 1993; Smith, 2005) might be due to its role in modulating agent/environment interaction rather than to its perceptual nature or to the role of action in shape perception. In real life, when humans manipulate objects, they modify the perceived shape, for example, as a result of object squeezing or modification of the object orientation. The investigation of this hypothesis is a challenge for future research.

Acknowledgments. This research was supported by the European Commission. The first and the fourth authors were funded by ITALK (EU FP7-ICT-214668), while the second and the third authors were funded by ROSSI (EU FP7-STREP-216125).

References

Ansuini, C., Giosa, L., Turella, L., Altoè, G., & Castiello, U. (2008). An object for an action, the same object for other actions: effects on hand shaping. *Experimental Brain Research, 185* (1), 111--9.

Ansuini, C., Tognin, V., Turella, L., & Castiello, U. (2007). Distractor objects affect fingers' angular distances but not fingers' shaping during grasping. *Experimental Brain Research,178* (2), 194--205.

Ashby, E. G., & Maddox, W. D. (2005). Human category learning. *Annual Review of Psychology, 56*, 149--178.

Barsalou, L. W. (2002). Being there conceptually: Simulating categories in preparation for situated action. In N.L. Stein, P. J. Bauer, & M. Rabinowitz (Eds.), *Representation, memory, and development: Essays in honor of Jean Mandler* (pp. 1--19). Mahwah, NJ, US: Lawrence Erlbaum Associates Publishers.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology, 59* (1), 617--645.

Bateson, G. (1987). Men are grass: Metaphor and the world of mental process. In W. Thompson (Ed.), *Gaia: A way of knowing* (pp. 37--47). Hudson, NY, US: Lindisfarne Press.

Borghi, A. M. (2005). Object concept and action. In D. Pecher, & R. A. Zwaan (Eds.), *Grounding Cognition: The role of perception and action in memory, language, and thinking* (pp. 8--34). Cambridge, UK: Cambridge University Press.

Borghi, A. M., & Cimatti, F. (2010). Embodied cognition and beyond: Acting and sensing the body. *Neuropsychologia, 48* (3), 763--773.

Borghi, A. M., Flumini, A., Natraj, N., & Wheaton, L. A. (2012). One hand, two objects: Emergence of affordance in contexts. *Brain and Cognition*, *80* (1), 64--73.

Brouwer, A. M., Georgiou, G., Glover, S., & Castiello, U. (2006). Adjusting reach to lift movements to sudden visible changes in target's weight. *Experimental Brain Research, 173* (4), 629--636.

Chatterjee, A. (2010). Disembodying cognition. *Language and Cognition, 2* (1), 79--116.

Chemero, A. (2009). *Radical embodied cognitive science*. Cambridge: MIT Press.

Costantini, M., Ambrosini, E., Scorolli, C., & Borghi, A. M. (2011). When objects are close to me: affordances in the peripersonal space. *Psychonomic Bulletin & Review, 18* (2), 302--308.

Crajé, C., Lukos, J. R., Ansuini, C., Gordon, A. M., & Santello, M. (2011). The effects of task and content on digit placement on a bottle. *Experimental Brain Research, 212* (1), 119--124.

Ellis, R., Swabey, D., Bridgeman, J., May, B., Tucker, M., & Hyne, A. (2013). Bodies and other visual objects: the dialectics of reaching toward objects. *Psychological Research, 77* (1), 31--39.

Ellis, R., & Tucker, M. (2000). Micro-affordance: The potentiation of components of action by seen objects. *British Journal of Psychology, 91* (4), 451--471.

Gallese, V., & Lakoff, G. (2005). The brain's concepts: The role of the sensorimotor system in conceptual knowledge. *Cognitive Neuropsychology, 21* (3), 455--479.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Gigliotta O., & Nolfi S. (2008). On the Coupling Between Agent Internal and Agent/Environmental Dynamics: Development of Spatial Representations in Evolving Autonomous Robots, *Adaptive Behavior, 16* (2--3), 148--165.

Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences, 20* (1), 1--55.

Glenberg, A. M., & Gallese, V. (2012). Action-based language: A theory of language acquisition, comprehension, and production. *Cortex, 48* (7), 905--922.

Goldman, A. I., & De Vignemont, F. (2009). Is social cognition embodied? *Trends in Cognitive Sciences, 13* (4), 154--159.

Iachini, T., Borghi, A. M., & Senese, P. (2008). Categorization and sensorimotor interaction with objects. *Brain and Cognition, 67* (1)*,* 31--43.

Jenmalm, P., Schmitz, C., Forssberg, H., & Ehrsson, H. H. (2006). Lighter or heavier than predicted: neural correlates of corrective mechanisms during erroneously programmed lifts. *The Journal of Neuroscience, 26* (35), 9015--9021.

Jones, S., & Smith, L. B. (1993). The place of perception in children's concepts. *Cognitive Development, 8* (2), 113--139.

Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3* (3), 299--321.

Landau, B., Smith, L.B., & Jones, S. (1998). Object perception and object naming in early development. *Trends in Cognitive Sciences, 2* (1), 19--24.

Lederman, S. J., & Wing, A. M. (2003). Perceptual judgment, grasp selection and object symmetry. *Experimental Brain Research, 152* (2), 156--165.

Mandler, J. M. (2004). Thought before language. *Trends in Cognitive Science, 8* (11), 508--513.

Mandler, J. M., Bauer, P. J., & McDonough, L. (1991). Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology, 23* (2), 263--298.

Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin, 129* (4), 592--613.

Maturana, H., & Varela, F. (1992). *The Tree of Knowledge: The biological roots of human understanding*. Boston & London: Shamabala.

Mizelle, J. C., & Wheaton, L. A. (2010). Why is that hammer in my coffee? A multimodal imaging investigation of contextually based tool understanding. *Frontiers in Human Neuroscience, 4*, 233.

Natraj, N., Poole, V., Mizelle, J. C., Flumini, A., Borghi, A. M., & Wheaton, L. (2013). Context and Hand Posture Modulate the Neural Dynamics of Tool-Object Perception. *Neuropsychologia, 51* (3), 506--519.

Nöe, A. (2004). *Action in Perception*. Cambridge, MA: MIT Press.

Nolfi, S., & Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books.

Nolfi, S., & Marocco, D. (2002). Active perception: A sensorimotor account of object categorization. In B. Hallam, D. Floreano, J. Hallam, G. Hayes, & J-A. Meyer (Eds.), *From Animals to Animats 7*, proceedings of the VII International Conference on Simulation of Adaptive Behavior (pp. 266--271). Cambridge, MA: MIT Press.

Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. A., McRae, K., & Spivey, M. (2011). The mechanics of embodiment: A dialogue on embodiment and computational modeling. *Frontiers in Psychology, 2* (5), 1--21.

Pfeifer, R., & Scheier, C. (1999). *Understanding Intelligence*. Cambridge, MA: MIT Press.

Poirier, P., Hardy-Vallee, B., & De Pasquale, J-F. (2008). Embodied categorization. In H. Cohen, & C. Lefebvre (Eds.), *Handbook of Categorization in Cognitive Sciences*. Amsterdam: Elsevier.

Rosch, E., Mervis, Carolyn, B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8* (3), 382--439.

Samuelson, L. K., & Smith, L. B. (1999). Early noun vocabularies: Do ontology, category structure, and syntax correspond? *Cognition, 73* (1), 1--33.

Sartori, L., Straulino, E., & Castiello, U. (2011). How Objects Are Grasped: The Interplay between Affordances and End-Goals. *PLoS ONE, 6* (9), e25203.

Scorolli, C., Borghi, A. M., & Glenberg, A. M. (2009). Language-induced motor activity in bimanual object lifting. *Experimental Brain Research, 193,* 43--53.

Smith, L. B. (2005). Action Alters Shape Categories. *Cognitive Science, 29* (4), 665--679.

Tipper, S. P., Paul, M., & Hayes, A., (2006). Vision-for-action: The effects of object property discrimination and action state on affordance compatibility effects. *Psychonomic Bulletin & Review, 13* (3), 493--498.

Tuci, E., Massera, G., & Nolfi, S. (2010). Active categorical perception of object shapes in a simulated anthropomorphic robotic arm. *Transaction on Evolutionary Computation Journal, (14)* 6, 885--899.

Tucker, M., & Ellis, R. (2001). The potentiation of grasp types during visual object categorization. *Visual Cognition, 8* (6), 769--800.

van Elk, M., Slors, M., & Bekkering, H. (2010). Embodied language comprehension requires an enactivist paradigm of cognition. *Frontiers in Psychology, 1* (234).

Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. Cambridge MA: MIT Press.

Vassie, K., & Morlino, G. (2012). Natural and Artificial Systems: Compare, Model or Engineer? In T. Ziemke, C. Balkenius, & J. Hallam (Eds.), *From Animals to Animats 12*, LNAI 7426 (pp. 1--11). Berlin / Heidelberg: Springer-Verlag.

Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review, 9* (4), 625--636.

Yoon, E.Y., Humphreys, W. W., & Riddoch, M. J. (2010). The paired-object affordance effect. *Journal of Experimental Psychology: Human Perception and Performance, 36* (4), 812--824.

Appendix A: Scoring

The score S was computed on the basis of the following equations:

$$S = \begin{cases} 0 & S' < .02 \\ S' & .02 \leq S' \leq .9 \\ 1 & otherwise \end{cases}$$

$$S' = \frac{1}{1+e^{6-12\sqrt{s}}} \qquad (1)$$

where s depended on the required manipulation.

For object to be placed, s was computed as $s = 1 - \frac{d}{a}$, where d was the distance from the object and the target area (positioned in the bottom-left corner or on the right-center of the arena, depending on the specific manipulation required) at the end of the trial, when the agent skipped to the next trial or when the time provided for the trial expired, and a was the side length of the square arena in which the objects were manipulated.

In the case of objects to be shaken s, whose value was initially set to *0*, was computed through an exponential moving average updated each time that the direction of the movement changed along the relevant direction (horizontally or vertically, depending on the required manipulation):

$$s = \alpha s + (1 - \alpha) s_i$$

where $\alpha = 0.7$ was the time constant of the moving average and s_i was the score gained at the i-th change of the (horizontal/vertical) direction. This, in turn, was computed as:

$$s_i = \sqrt{1 - \frac{|\delta_i - a/2|}{a/2}}$$

where δ_i was the distance between the horizontal/vertical projections of the *i*-th and *(i-1)*-th points (the *0*-th point was the position in which the object was at the beginning of the trial) in which the movement along the relevant direction changed.

Notice that the non-linear transformation in (1) provided the agent with a bad/neutral/good but still continuous feedback (in the range [0.02, 0.98]) that pushed the agent more strongly toward close-to-optimal manipulation. In order to be more easily read by humans the score *S* was linearly mapped from [0, 1] to [0, 100].

Appendix B: Instructions and questionnaire for human participants

Preliminary instructions:

"You are in an alien world where there are a number of objects with which you have to appropriately interact using the mouse.

The game is divided into two phases.

In the first phase (TRAINING) you are asked to learn to interact properly with the objects, which will be presented one at a time. You can interact with each object, for a maximum time of 30 seconds, how many times you want. At the end of each period of interaction you will be given a score (between 0 and 100). This score, during training, allows you to understand how well you interacted with each object and DOES NOT PENALIZE YOU. If you are not interested in an item you can safely skip to the next. Once trained, you can get the maximum score for each object after interaction with it for a few seconds. The value of the score is also represented in the vertical bar on the right side of the window.

The training phase lasts 25 minutes and then you will pass to the TESTING phase, where you have five minutes to interact with AS MANY OBJECTS as possible (the same experienced during the training phase) trying to achieve, for each, the MAXIMUM SCORE (100). The total points gained at this stage will indicate your ability to survive in this alien world.

During the game 'THINK ALOUD' saying what you see, think, do, and how do you feel.

Press 'SPACE' to start the experiment.

Click the LEFT mouse button to jump to the next object. Remember that moving to the next object does not penalize you.

Click the RIGHT button to interact again with the last object (only possible during training). Remember that interacting again with the last object does not penalize you."

Instructions at the end of the training phase:

"Training phase finished.

Now you have 5 minutes in which, for every object with which you'll interact, you'll receive a score.

In this phase, unlike in the training phase, you cannot interact again with the last object.

The total score received in this phase will indicate your chance of surviving in this alien world.

Press 'Space' to start."

Questionnaire:

1. *Please describe with your own words the objects that you experienced in the alien world.*
2. *Did you notice specific object properties? Can you remember/describe these properties? Do you think that all properties were relevant to the task, or some characteristics were more important than others?*
3. *Can you describe your strategies of interaction? Do you think they were successful?*
4. *How did you feel during the experiment? Did you ever feel bored/tired?*
5. *Did you think-aloud? If not, do you think it could have helped you to do so?*

Figures

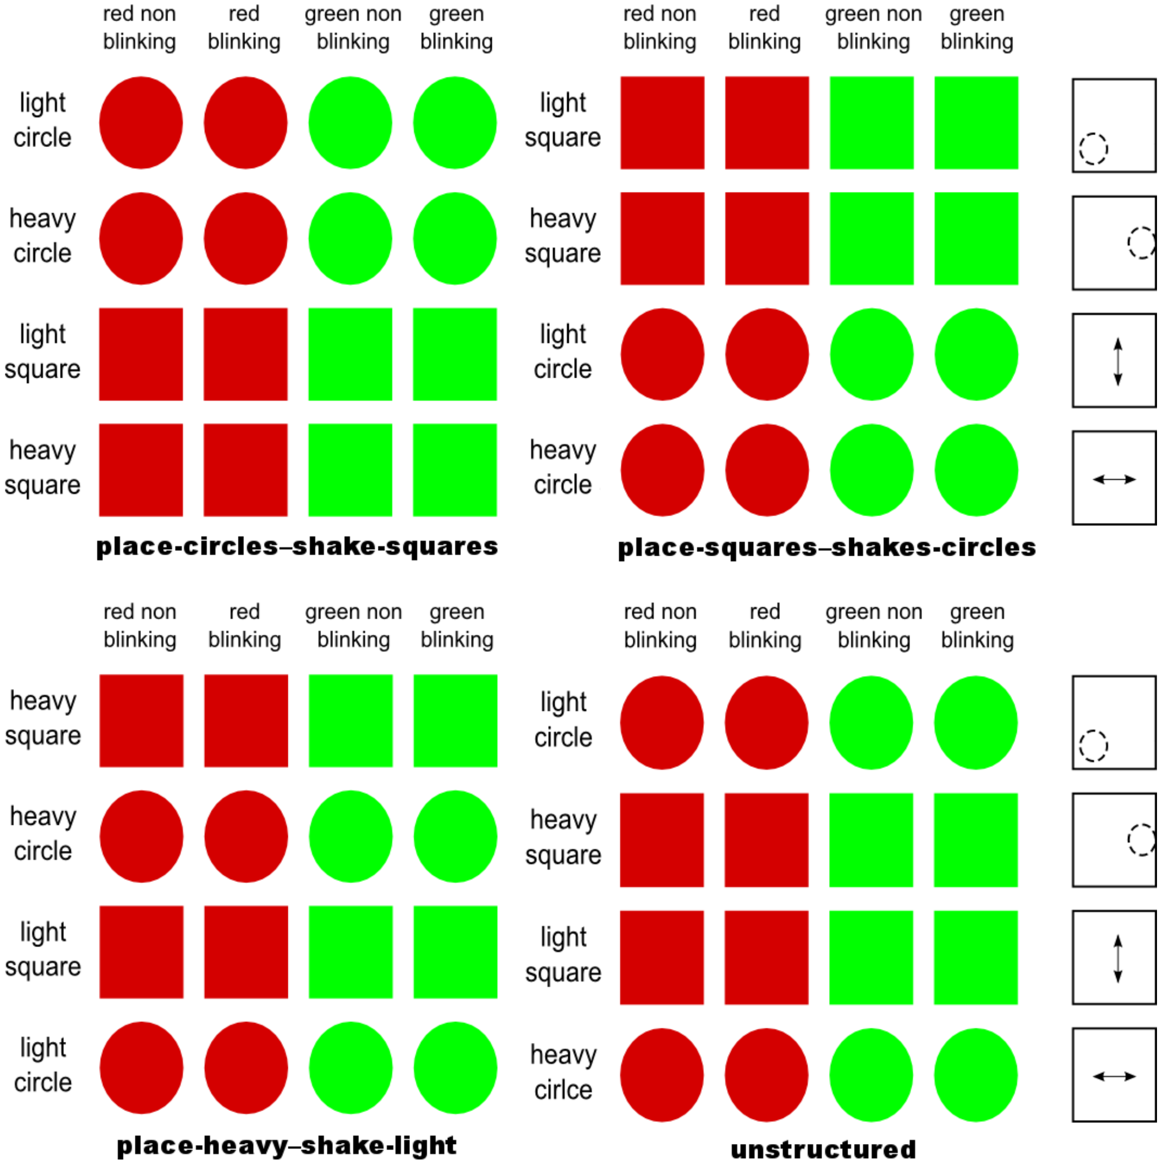

Fig. 1. Experimental conditions. In each condition the 16 objects, grouped in categories, have to be manipulated differently. All the objects in a row have to be manipulated accordingly to the manipulations depicted on the right, so objects from the first, second, third, and fourth row have to be placed left, placed right, shaken vertically, and shaken horizontally, respectively. The circle/square objects have a diameter/side of 100 pixels (corresponding to 3 cm on the screen used to perform the experiment). The dimensions of the area in which the

objects can be moved are 600x600 pixels (18x18 cm). The arrows indicate the shaking behaviors. The dashed circles indicate the target areas where objects have to be placed. The target areas (depicted with dashed circles) are centered at (150, 150) and (450, 300), respectively (coordinates relative to the bottom-left corner of the arena), and their size, defined as the diameter of the area in which the score is maximum, is 200 pixels (6 cm).

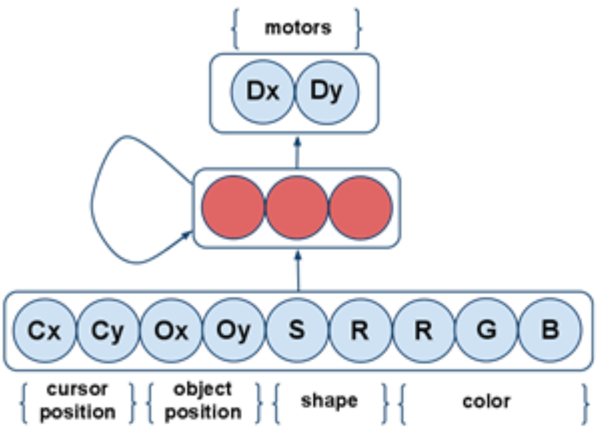

Fig. 2. Architecture of the artificial neural controller. Sensory, hidden, and motor neurons are depicted at the bottom, center, and top of the illustration, respectively. Arrows indicate that all neurons of one layer are connected to all neurons of the other. Sensory neurons were relay units, hidden neurons were dynamic units, and motor neurons were standard logistic units (see text for details).

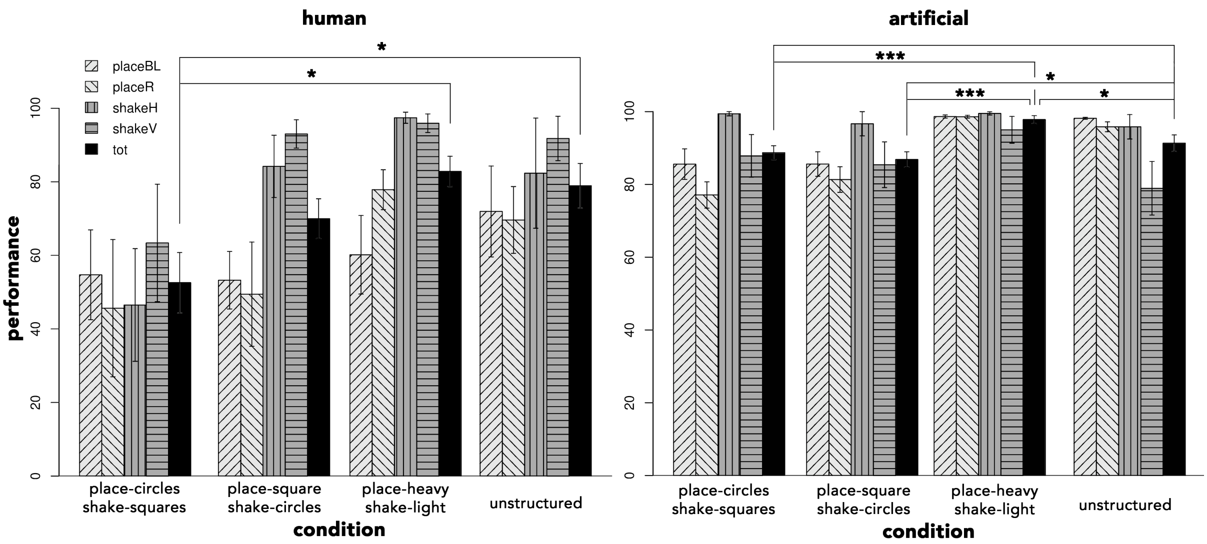

Fig. 3. Overall performance and performance by target manipulations obtained in the four experimental conditions by human participants and artificial agents. Error bars represent the standard error of the means. Significant differences in the overall performance (rightmost bars) are marked with stars ("*" means the difference is significant at .05 level and "***" means the difference is significant at .01 level). **Left**: performance of the 20 human participants in the test phase. **Right**: average performance of the best 120 artificial agents (30 replications for each experimental condition) at generation 50. Data obtained by post-evaluating the agents for 640 trials.

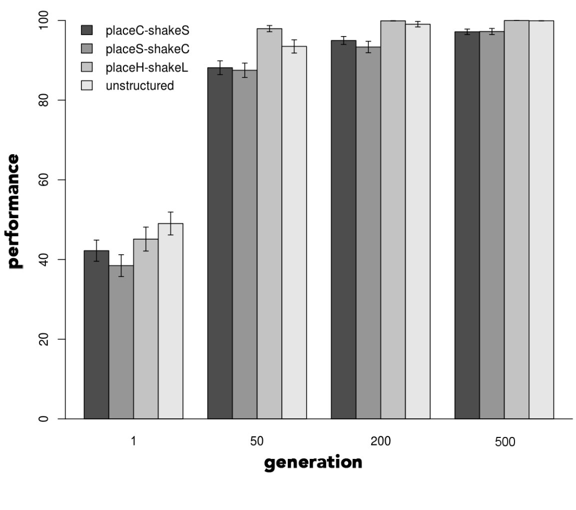

Fig. 4. Average performance of the best artificial agents of 30 replications of the experiment in the four experimental conditions and at different phases of the training process (generations 1, 50, 200, and 500). Error bars represent the standard error of the means. Data obtained by post-evaluating the best agents for 640 trials.

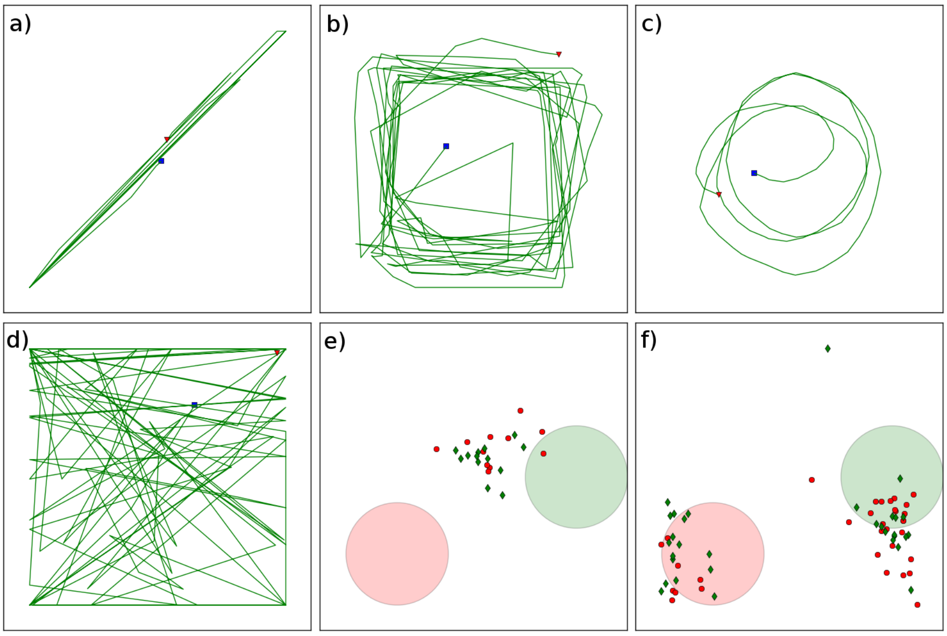

Fig. 5. Type of the overgeneralized strategies displayed by human participants. Little squares and triangles in panels a, b, c, and d represent the start and end points of four exemplary trajectories, respectively. Little circles and diamonds in panels e and f represent the end points of the trajectories performed by two participants during the test phase with objects to be placed in the bottom-left and right target areas (big circles), respectively.

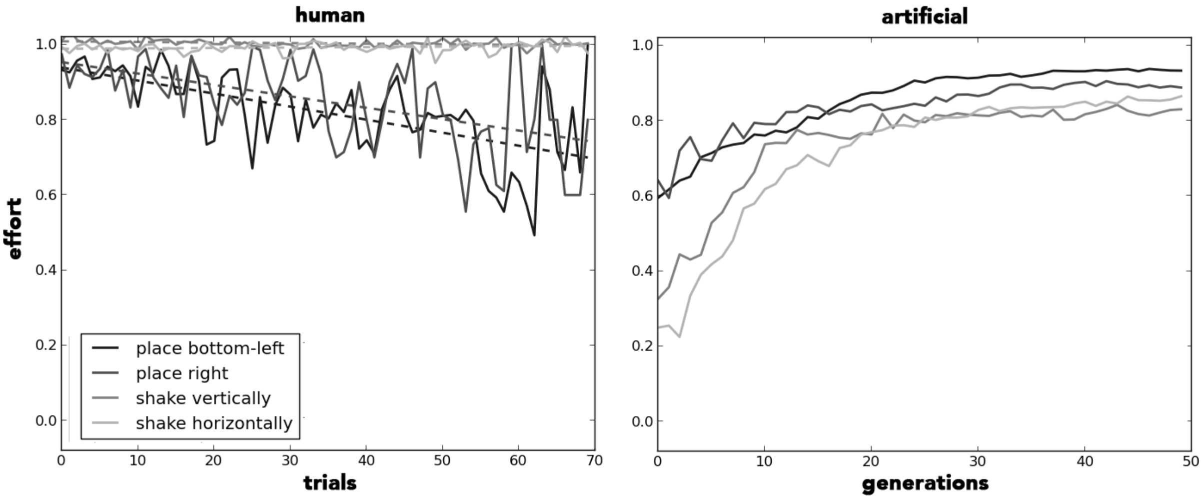

Fig. 6. *Effort* of human participants and artificial agents. The left picture reports the average *effort* of the human participants for each target action over the first 70 trials. Linear regression lines are drawn to show the trends (dashed lines). The right picture reports the average *effort* of the best artificial agents computed over the first 50 generations.

Tables

Tab. 1. Number and frequency of overgeneralized strategies.[2]

Gen. level		2									3			4			
Strategy		a	b	c	d	e	f	g	h	a or b	a+l	b+l	c+l	a+i	c+i	a+e	d+e
Human	**n°**	3	2	1	7	3	2	0	0	2	0	0	0	0	0	2	1
	%	15.00	10.00	5.00	35.00	15.00	10.00	0.00	0.00	10.00	0.00	0.00	0.00	0.00	0.00	10.00	5.00
Artificial	**n°**	10	0	9	0	1	2	8	1	0	3	1	1	58	27	0	0
	%	8.33	0.00	7.50	0.00	0.83	1.67	6.67	0.83	0.00	2.50	0.83	0.83	48.33	22.50	0.00	0.00

Tab. 2. Overall frequency of the overgeneralized strategies.[3]

Gen. level	2	3	4	**Total**
Human %	50.00	0.00	15.00	**65.00**
Artificial %	12.92	4.17	70.83	**87.92**

[2] data related to the 20 human participants and the 120 artificial agents and listed by level of generalization, defined as the number of categories of objects manipulated in the same way; the strategies from a to f are depicted in Fig. 5.; g and h stand for “shake vertically” and “shake horizontally”, respectively, while i stands for “place the objects in the correct target areas” and l stands for “place all objects in one of the two target areas”.

[3] the level of generalization is defined as the number of object categories manipulated in the same way.

Tab. 3. Features reported as relevant by human participants.[4]

Condition	1					2					3					4				
Participant	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Shape	r	r	-	r	r	p	r	r	r	p	-	p	-	-	r	r	p	r	r	p
Weight	p	p	p	p	-	p	p	-	p	r	r	r	r	r	r	-	-	r	r	p
Color	-	-	-	-	-	-	-	r	-	-	-	p	-	-	-	-	p	-	-	-
Blinking	-	-	-	-	-	r	-	-	-	-	-	-	-	-	-	-	p	r	-	r

[4] ‘r’ / ‘p’ stand for feature reported as relevant / partially relevant. Conditions 1, 2, 3, and 4 represent place-circles–shake-squares, place-squares–shakes-circles, place-heavy–shake-light, and unstructured, respectively.