



Action and hierarchical levels of categories: A connectionist perspective

Action editor: Vasant Honavar

Anna M. Borghi^{a,b,*}, Domenico Parisi^b, Andrea di Ferdinando^{b,c}

^a *Dipartimento di Psicologia, University of Bologna, Viale Bertoni Pichat 5, 40127 Bologna, Italy*

^b *Institute of Cognitive Sciences and Technologies, National Research Council, Rome, Italy*

^c *Dipartimento di Psicologia, University of Padua, Padua, Italy*

Received 1 July 2003; received in revised form 24 June 2004; accepted 25 June 2004

Available online 28 August 2004

Abstract

Recent views of categorization suggest that categories are action-based rather than arbitrary symbols. Three connectionist simulations explore the hierarchical organization of categories in the framework of an action-based theory of categorization. In the simulations an organism with a visual system and a two-segment arm has to reach different points in space depending on the object seen and on context. The context indicates whether to put the object in a superordinate or in a basic category. The results show that: (a) superordinate categories are easier to learn than basic ones; (b) the more similar the actions to perform with basic and superordinate categories, the easier to learn the task; (c) violation of category boundaries leads to less good performance.

© 2004 Elsevier B.V. All rights reserved.

Keywords: Categorization; Action; Hierarchical organization; Neural networks

1. Introduction

Traditional studies of categorization have often been guided by the implicit assumption that con-

cepts, i.e., the mental representations of categories, are made of arbitrary symbols. This view implies a translation process from sensory-motor experiences into symbols which are arbitrarily related to their referents. In alternative, recent views conceive of concepts as forms of re-enactment of sensory-motor experiences (Barsalou, 1999; Barsalou, Simmons, Barbey, & Wilson, 2003; Glenberg,

* Corresponding author. Tel.: +39-051-2091838; fax: +39-051-243086.

E-mail address: annamaria.borghi@unibo.it (A.M. Borghi).

1997; Smith, 1995; Smith & Katz, 1996; Thelen, Schöner, Scheier, & Smith, 2001). In this perspective, concepts are action-based and the information they contain helps to prepare for situated actions (Barsalou, 2002; Wilson, 2002).

In this paper, we describe some connectionist simulations which explore various aspects of this view of concepts or categories as action-based. The general assumption is that neural networks that respond to sensory input with movements learn to form categories on the basis of the output, i.e., the action, with which they respond to the input rather than on the basis of the perceptual characteristics of the input.

More specifically, our aim in this paper is to explore with neural network simulations how an action-based theory of categories, can explain the formation of a hierarchical structure of categories. With the term hierarchical structure we refer to the different levels of generality at which knowledge can be organized. For example, knowledge can be organized around more general categories (animals, furniture, etc.) or more specific ones (cats, dogs, tables, etc.). The first are called superordinate, the second basic (level) categories. These two kinds of categories are linked by a class inclusion relations: basic categories are included in superordinate ones. This hierarchical organization has the advantage of being economical from a cognitive point of view, even if recently the transitivity and advantages of cognitive economy have been questioned (Hampton, 1982; Sloman, 1998).

Within the general framework of an action-based theory of concepts or categories, we aim to replicate and explain some empirical results concerning hierarchical levels of categories.

(1) *Precedence of global or superordinate categories over specific or basic categories.* In the literature on categorization much convergent evidence shows that basic categories (e.g., table, dog) are generally preferred for categorization and are acquired by children earlier than superordinate ones (Markman, 1989; Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976). However, as highlighted by Mandler (1992a, 1992b, 1998), this evidence is based on linguistic tasks. Mandler and Bauer (1988), Mandler and McDonough (1998) and Mandler, Bauer, and McDonough (1991) have

shown that with non-linguistic tasks infants first form global categories, which then help them in the acquisition of language and of more specific categories. Simulation 1 aims to test whether superordinate categories are easier to learn, and thus are learned earlier, than basic level categories. Consider that, depending on the context, it may be more adaptive to use either global (e.g., animal) or more specific categories (e.g., dog). For example, when our ancestors were in danger it may have been useful to simply discriminate between preys and predators, while when they had to feed animals it may have been useful to distinguish between dogs, cows, chickens, etc. Accordingly, context is introduced as a cue for using either superordinate or basic action-based categories.

(2) *Influence of action similarity on the formation of superordinate categories.* If action influences categorization, superordinate categories should be more easily learned the more the action to perform with them is similar to the action to perform with their basic level members. This hypothesis is tested by comparing Simulations 1 and 2.

(3) *Influence of category boundaries on the formation of superordinate level categories.* Studies on categorization distinguish between Common Taxonomic (CT) categories, as for example animals and vehicles, and Goal Derived (GD) categories, as for example birthday presents (Barsalou, 1983, 1985, 1991). CT categories, which are more stable in memory, include perceptually similar objects that belong to the same domain, while GD categories group together objects which are not necessarily perceptually similar nor of the same domain (e.g., birthday presents may include pets, flowers, books, and cars). The novelty in this study is that we build superordinate CT and GD categories which are totally action-based, i.e., the influence of perceptual similarity among the category members is ruled out. This means that superordinate categories are formed either assembling two action-based basic categories (CT superordinate), each containing two exemplars, or assembling four exemplars of four different basic categories (GD superordinate).

The question of whether more general categories are more easily learned and distinguished from basic categories when they are formed by respect-

ing category boundaries or violating them is tested by comparing the first two simulations with the third one. We predict that when category boundaries are violated it is more difficult to learn the task.

2. Neural networks and action

Neural networks are systems for transforming activation patterns into other activation patterns. An activation pattern in the input units is transformed into a different activation pattern in the first layer of internal units and then this activation pattern is transformed into another activation pattern in the second layer of internal units until the output activation pattern is generated. What are the principles that govern these transformations? Neural networks follow two principles. The first principle is: “Make activation patterns corresponding to stimuli that should be responded to with the same output activation pattern progressively more similar to each other”. The second principle is: “Make activation patterns corresponding to stimuli that should be responded to with different output activation patterns more different from each other”. The first principle can be called the principle of categorization. The second principle can be called the principle of discrimination.

We can measure similarity among activation patterns quantitatively if we assume that activation patterns are points in the abstract hyperspace which corresponds to each layer of units. The hyperspace has as many dimensions as there are units in the layer and each dimension of the hyperspace measures the activation level of one unit. Activation patterns are represented as points located, for each dimension, in the position that corresponds to the activation level of the unit represented by the dimension. Therefore, in the hyperspace corresponding to a given layer of units are represented all possible activation patterns that may be observed in that layer of units. We can measure similarity between activation patterns in each particular layer of units as the distance between the points that represent the activation patterns in that layer. The two principles of

categorization and discrimination say that in the successive layers of internal units of a neural network the distance between points that represent activation patterns which must be responded to with the same output tends to decrease (categorization), while the distance between points that represent activation patterns which must be responded to with different outputs tends to increase (discrimination).

If we designate with the term “cloud” the set of points in some particular layer of units that represents activation patterns corresponding to stimuli which must be responded to with the same output, the internal organization of a neural network can be described by saying that as activation propagates from input to output the neural network tends to make its individual “clouds” progressively smaller (categorization) and the distance between distinct “clouds” progressively larger (discrimination).

This interpretation of the internal organization of a neural network makes it clear that the network’s internal representations (i.e., the activation patterns observed in the successive layers of internal units) are dictated by the network’s output. The two principles of categorization and discrimination are output-based principles. Inputs are internally represented in terms of the output with which the network must respond to the input. The properties of the input, those which distinguish one input from another input, have some control on the internal representations only in the first layers of internal units. But even at this early stage the internal representations follow the principles of categorization and discrimination which are output-based principles. “Clouds” are defined in terms of a neural network’s motor output, not in terms of the network’s sensory input. “Clouds” are formed in such a way that the average distance of the points included in the “cloud” from the center of the “cloud” tends to be smaller than the distance between the centers of distinct “clouds”. If the neural network is a sensory-motor network, i.e., a network which maps sensory inputs into motor outputs, the sensory inputs tend to be internally represented in terms of the similarity relations of the motor outputs rather than in terms of the similarity relations of the sensory

inputs. If “clouds” are categories, categories are action-based rather than stimulus-based.

In previous work on categorization and action we found that the internal representations of an organism’s neural network appear to be organized in terms of macro-actions, that is sequences of movements (micro-actions) that allow the organism to correctly respond to perceived objects (Di Ferdinando & Parisi, 2004). In other simulations it was shown that the organism’s neural network was able to flexibly organize itself in order to adapt to different tasks. The network’s internal representations of objects reflect the current task and not the perceptual similarity between the objects. In absence of task information, however, perceptual similarity is the best predictor of categorization (Borghi, Di Ferdinando, & Parisi, 2002; Di Ferdinando & Parisi, 2004; Di Ferdinando, Borghi, & Parisi, 2002). This paper goes one step further: it aims to apply an action-based theory of categorization to a specific problem in the literature of categorization, the study of hierarchical levels.

3. The model

Imagine an organism which lives in a bidimensional environment containing eight different objects. The organism has a simulated visual system with which it sees only one object at a time and a two-segment arm which it can move to reach some specific position in space which depends on the particular object seen and on the particular context (task). In each trial the initial position of the arm is random. The organism can know the arm’s position at any given time based on proprioceptive input from the arm’s two segments. The organism’s behavior is controlled by a nervous system, which is simulated using an artificial neural network.

The visual system is a $3 \times 3 = 9$ cell matrix (nine input units) and each object is very simple: it appears as a particular filled cell in the matrix of nine cells (one of the nine input units has an activation level of one while the remaining eight cells have zero activation). There is a total of eight objects (in the central cell no object can appear). Fig. 1 shows the organism who is currently seeing an




Fig. 1. The organism and the retina with one of the eight objects.

object. Notice that all objects are identical to each other, because we wanted to rule out the influence of perceptual similarity. They are only distinguished on the basis of the spatial location in the retina in which they appear.

The proprioceptive input from the arm which tells the organisms the position of their arm at any given time specifies the angle of the forearm with the shoulder and the angle of the arm with the forearm (two input units).

The network architecture includes a third set of two input units which encode two different contexts: context A, encoded as 10, and context B, encoded as 01.

The network has two output units encoding changes in the arm’s two angles and therefore determining the arm’s movements.

The nine visual input units and the two context units are connected with the two output units through an intermediate layer of four hidden units whereas the two proprioceptive units project directly to the two output units. The network architecture is described in Fig. 2.

The organisms must respond to this input by moving the arm in such a way that the arm’s endpoint (the hand) reaches one particular location in space (i.e., presses some particular button) which depends on the particular object which is presented




Fig. 2. The neural network architecture.

to the organism in any given epoch and on the context. Thus, they have to perform a different kind of action with the same object depending on the context. This scenario addresses a problem we encounter in real life: depending on the kind of context, it may be advantageous to categorize objects at the superordinate or at the basic level. For example, it may be useful to distinguish predators and preys when we have to decide whether to run away or not, while it may be useful to distinguish dogs and cows when we have to go hunting or to drink some milk.

In our simulations, the eight objects are grouped into four basic categories, two objects for each basic category, and into two superordinate categories, four objects per each superordinate category. In context A the organisms must use the basic categories. They have to press one of four buttons depending on the object's basic level category, whereas in context B they have to press one of two buttons depending on the object's superordinate category (see Fig. 3). Notice that while basic categories are the same in all the three simulations, the way superordinate categories are formed varies in the three simulations.

All organisms live the same amount of time (number of input/output cycles of their neural network). Each individual organism is a member of a generation of 100 organisms. At the end of life the 20 individuals which perform better in the task are selected for reproduction. Five copies of the connection weights of their neural network are gener-

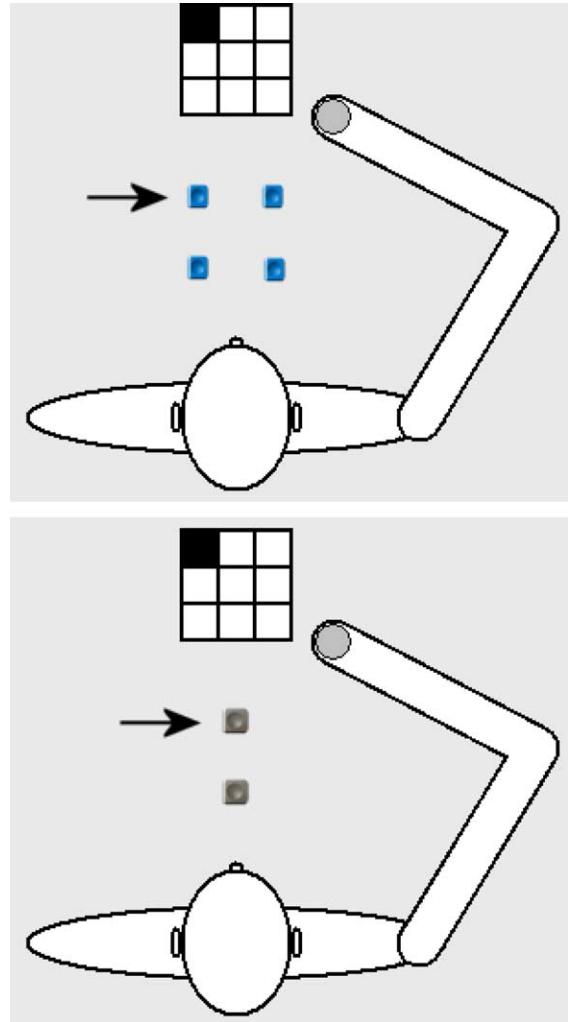


Fig. 3. The spatial locations that must be reached by the arm's endpoint to form four subordinate categories comprising two objects each and the two superordinate categories comprising four objects each.

ated and assigned to five new individuals (offspring). The $20 \times 5 = 100$ offspring constitute the next generation. Reproduction is nonsexual but an average of 10% of the inherited connection weights are randomly mutated so that offspring behave similarly but not identically to their parents.

In order to obtain more reliable results we run 10 replications of each simulation starting with different randomly generated connection weights.

4. Simulations

4.1. Simulation 1

In this simulation, we address the first hypothesis: we test whether superordinate concepts are easier to learn, and thus are learned earlier in terms of number generations needed to achieve an appropriate performance level, than basic categories.

4.1.1. Task

In this simulation the action the organism has to perform in response to all four members of a superordinate category is similar to the actions to be performed in response to the 2 + 2 members of its two basic sub-categories and dissimilar to the actions performed with the 2 + 2 members of the two basic sub-categories of the other superordinate category. Action similarity is measured in terms of spatial distance: the button (space location) that the organism has to reach with its arm in response to the four members of one superordinate category is spatially close to the buttons that must be reached in response to the 2 + 2 members of 2 basic sub-categories of that superordinate category (see Fig. 4(a)).

4.1.2. Results

We trained the organisms for 3000 generations. As Fig. 5 shows (condition “Close”), the organisms were able to learn the task.

In order to verify whether superordinate level categories are easier to learn and thus formed before, in terms of number of generations, than basic categories, we performed a test at successive stages during the training (every 500 generations). In this test the organisms were presented with all possible trials starting from five different random positions of the arm for each trial and we computed the percentage of targets reached for context A (basic) and for context B (superordinate) (see Fig. 6).

On these results we performed two sets of Anovas. In the first set we compared the performance of the best individuals while in the second set we compared the performance of the whole population (average).

In all cases we found a clear advantage of superordinate over basic level categories ($p < 0.01$).

4.1.3. Discussion

Given that all objects are perceptually identical, our results show that organisms form action-based categories independently from the perceptual similarity of the category members.

More interestingly, the comparison between learning superordinate and basic categories shows that superordinate categories are acquired more easily (earlier) than basic level categories. The explanation seems to be quite straightforward: forming more specific categories by associating the same set of objects with a larger number of distinct actions poses a greater burden on the same pool of computational resources (the network's




Fig. 4. The spatial locations that must be reached by the arm's endpoint in Simulation 1 (panel a), Simulation 2 (panel b), and Simulation 3 (panel c).

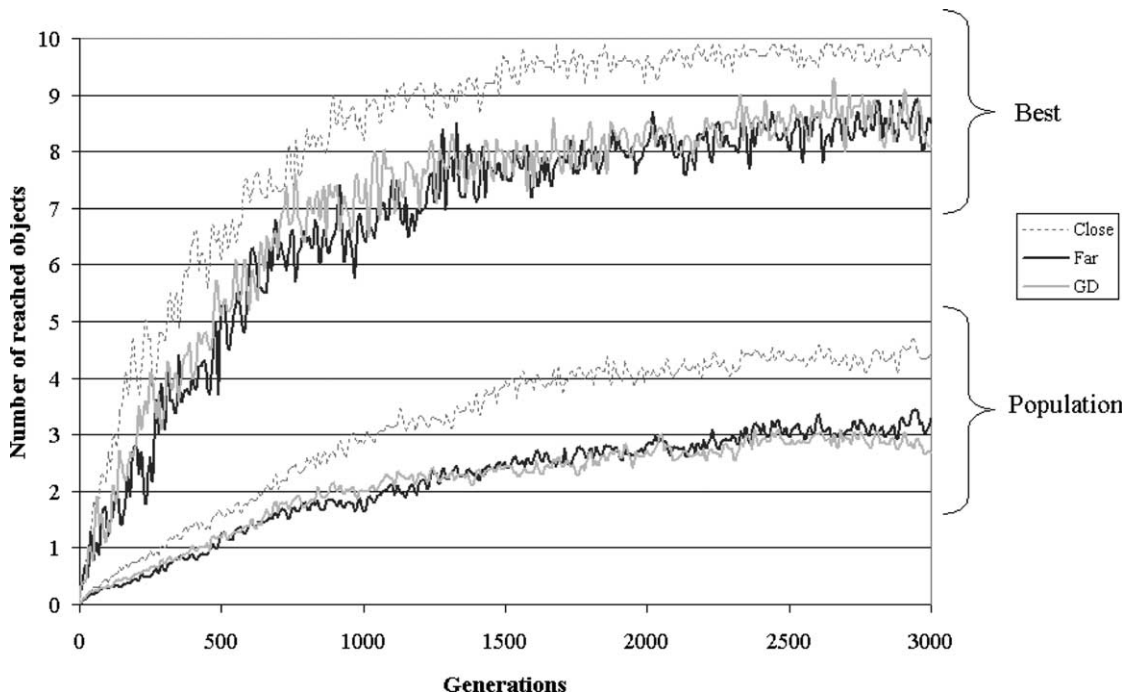


Fig. 5. Comparison between Simulation 1 (close), 2 (far), and 3 (GD). Number of objects reached by the best organism and the whole population. Average of 10 replications.




Fig. 6. Number of objects reached with superordinate and basic categories by the best individual and the whole population. Simulation 1.

connection weights) than forming fewer more general categories by associating the objects with fewer different actions.

This result has interesting implications for psychological studies of categorization. In fact, it supports the view according to which, before learning

language, infants first form more global categories and only later they form more specific categories (Mandler et al., 1991). Furthermore, it helps to explain why this happens: simply because it is simpler to associate different objects with a smaller number of different actions than with a greater number of actions.

4.2. Simulation 2

The aim of Simulation 2 is to test whether the performance is less good than that of Simulation 1 when the action in response to the members of a superordinate category is dissimilar to the actions performed in response to the members of the two basic level sub-categories of that superordinate category. In Simulation 1 the action to be performed at the superordinate level is to reach a button which is spatially located between the two buttons to be reached in response to the members of the category's to subcategories. In contrast, in Simulations 2 the button to be reached at the superordinate level is spatially located on the opposite side with respect to the two buttons to be reached when responding at the basic category level. In real life, an example of action similarity between the actions to perform with superordinate and basic categories is the following: if I have to wash an indument, I typically put it into the washing machine (I simply categorize it at the superordinate level); however, I would better wash a delicate skirt (basic level) by hand, and I would better bring an elegant jacket (basic level) to the laundry.

Our hypothesis is that the organism has more difficulties in forming action-based categories if the actions to be performed with superordinate level categories are dissimilar from those to perform with basic categories (hypothesis 2). Accordingly, performance in Simulation 2 should be worse than performance in Simulation 1.

4.2.1. Task

In this simulation the action to be performed with the four members of a superordinate category is different from the two actions to perform with the 2 + 2 members of the respective basic sub-categories. In fact, when responding to a superor-

dinate category (context B) the organism has to reach with the hand the central location on the opposite side with respect to the two locations to be reached when it has to respond to the category's basic sub-categories (context A) (see Fig. 4(b)).

4.2.2. Results

As predicted, the performance of Simulation 2 is worse than that of Simulation 1 (see Fig. 5, condition "Close" vs. "Far").

In order to compare the data obtained in the two simulations, we performed two sets of three within subjects Anovas. In the first set, we compared the average number of objects reached with superordinate and basic categories by the best individuals of each of the 10 replications at generations 1000, 2000, and 3000. In the second set we compared the average number of objects reached with superordinate and basic categories by the whole population of each replication at generation 1000, 2000, and 3000. The independent variable, the type of simulation performed (Simulation 1 or 2), was manipulated between-subjects.

The results are straightforward and confirm our prediction. The performance of both the best individual and the population average in Simulation 1 is significantly ($p < 0.01$) better than in Simulation 2. This clearly indicates that there is an effect of action similarity, as predicted.

4.2.3. Simulation 3

In the preceding simulation we found an effect of the similarity of the actions to be performed at the basic level and at the superordinate level. In Simulation 3 the superordinate category is formed by four members, each of which belongs to a different basic category (Fig. 4(c)). This means that the location the organism has to reach while categorizing at the superordinate level is close (similar) to the location it has to reach with two of the superordinate category members, while it is far from the location it has to reach with the other two members. Thus, if only action similarity influences performance, the results of Simulation 3 should be intermediate between those of Simulations 1 and 2.

We predict, instead, that performance is influenced not only by action similarity but also by

respecting vs violating category boundaries. In fact, these simulations differ from both preceding simulations because categories boundaries are violated. This means that a superordinate category (e.g., animal) is not formed by grouping together the members of two already existing basic categories (e.g., dog and robin) (CT categories), but by putting together the members of the four different basic categories (e.g., the superordinate category ‘birthday presents’ is formed by ‘dog’, ‘doll’, etc.). In the literature on categorization categories which violate category boundaries are called GD or goal derived categories (Barsalou, 1983, 1985, 1991; Vallée-Tourangeau, Anthony, & Austin, 1998).

Summarizing, in Simulation 3 it should be more difficult to form action-based superordinate categories in that superordinate categories are GD categories, i.e., they assemble together one member of each of the four different basic categories (hypothesis 3).

4.2.4. Task

In this simulation, when responding to superordinate categories the organism has to reach with the hand the central button on one side of the button space, thus performing an action similar to that performed with two of the objects included in the category and dissimilar to that performed with the other two members (see Fig. 4(c)).

4.2.5. Results

Fig. 5 shows that, as predicted, organisms perform worse in Simulation 3 than in Simulation 1, whereas there is no difference between Simulations 2 and 3. This indicates that the bad performance of organisms in Simulation 3 does not depend only on action similarity between superordinate and basic categories. If this were the case, in fact, performance in Simulation 3 should be better than performance in Simulation 2, given that, as we have seen, the average distance between the buttons to be reached with superordinate and basic categories is smaller in Simulation 3 than in Simulation 2.

The Anovas and the post hoc Newman–Keuls comparing the performance of the best individual and of the whole population in the three simulations every 1000 generations indicate ($p < 0.1$) that

the performance of both the best individual and the average of the population in Simulation 3 is not significantly different from that of Simulation 2 while it differs from that of Simulation 1.

This shows that not only action similarity has an effect but that there is also an effect of the violation of category boundaries.

In order to isolate the effect of the violation of category boundaries from the effect of action similarity, we performed a control simulation in which the output of the neural network does not control the movement of the arm, but directly determines the action to perform with a certain category member. More specifically, there are four output units encoding in a localistic way the four basic categories and two additional output units encoding the two superordinate categories. The task is to activate the two output units that correspond to the basic and superordinate categories to which the object currently seen belongs.

Given that the arm has not to be moved in this simulation, the proprioceptive input is not considered. The new neural architecture is shown in Fig. 7.

To train the network to solve the task we used a back-propagation algorithm. Given that there is no movement of the arm, there is no difference between the condition “spatially close” and the condition “spatially distant”. We refer to both these




Fig. 7. Not ecological network architecture used for the control with back-propagation.




Fig. 8. Control simulation. Performance with CT ad GD categories.

conditions as CT. We compared this condition with the condition GD. In this new scenario the only difference between CT and GD categories is the violation of category boundaries. Fig. 8 shows that CT categories are learned faster than GD categories, thus confirming our hypothesis.

5. Conclusion

In this research, we show with three simulations that neural networks can form action-based categories. In fact, rather than based on perceptual similarity between category members, in our simulations categories are formed on the basis of the kind of action to be performed (see also Borghi et al., 2002; Di Ferdinando et al., 2002). In addition, we show that neural networks are able to use different action-based categories, at both the superordinate and basic levels, depending on the context.

The similarity of the action to be performed at different hierarchical levels influences categorization: in fact, superordinate action-based categories are easier to form if their members have to be responded to with actions similar to those used with the members of their basic sub-categories.

In our simulations, we not only show that an action-based theory of categorization is plausible and that it works, but also that it replicates and explains various empirical results discussed in the psychological literature on categorization.

First, we show that an action-based theory of categories can explain the formation of a hierarchical organization in categorization, and the role played by context in selecting a categorization criterion. Consider that the hierarchical structure of knowledge has traditionally been explained in symbolic terms. In our work, we show that hierarchical organization may be action-based, suggesting that we might form superordinate categories on the basis on the similarity of the actions performed with their members rather than on the basis of perceptual similarity between them. Consider, however, that in this study we do not address the point of whether action similarity overcomes perceptual similarity. By underlying the role of action we do not intend to exclude that in real life perceptual similarity may play an important role for categorization (Goldstone, 1994).

Second, our results indicate that, in absence of linguistic input, superordinate categories are acquired earlier than basic ones. This result has

interesting implications for psychological studies of categorization. In fact, it supports the view according to which, before learning language, infants first form more global categories and only later they form more specific basic level categories (Mandler et al., 1991). Consider, however, that our results are not necessarily in conflict with studies showing the importance of basic level categorization (Rosch et al., 1976). In fact, studies on categorization generally use linguistic tasks and focus on linguistic categories (Malt, Sloman, Gennari, Shi, & Wang, 1999). Once categories are designed with verbal labels, basic level categories may acquire an advantage over superordinate level categories due to the higher frequency of basic level labels in comparison with superordinate level labels. This issue is not addressed here and it is worth of further exploration.

Third, we show that neural networks have more difficulties in learning superordinate GD categories, i.e., categories which violate basic level category boundaries, than superordinate CT categories that do not violate such boundaries. In general, in the literature of categorization it is argued that GD category are less stable in memory than CT categories (Barsalou, 1991). This instability is considered to be due to the less frequent use of these categories due to the perceptual dissimilarity between the category members. The novelty in our study is that the categories we use are completely action-based – thus the influence of perceptual similarity is ruled out. Using only action-based categories our study suggests that the instability of GD categories may depend on the frequency with which an action is performed with a certain category member, i.e., on the existing consistency between the canonical actions we perform with an object and the action we are required to perform in a new context.

Acknowledgments

Part of this research has been presented at the conference EuroCogSci03 with the title A.M. Borghi, A. Di Ferdinando and D. Parisi, Violation of category boundaries: a connectionist study with action-based categories.

References

- Barsalou, L. W. (1983). Ad hoc categories. *Memory and Cognition*, 11, 211–217.
- Barsalou, L. W. (1985). Ideals, central tendency, and frequency of instantiation as determinants of graded structure in categories. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 4, 629–654.
- Barsalou, L. W. (1991). Deriving categories to achieve goals. In G. H. Bower (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 27, pp. 1–64). New York: Academic Press.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 577–609.
- 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. 1–16). Mahwah, NJ: Erlbaum.
- Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Science*, 7, 84–91.
- Borghi, A., Di Ferdinando, A., & Parisi, D. (2002). The role of perception and action in object categorization. In J. A. Bullinaria & W. Lowe (Eds.), *Connectionist models of cognition and perception* (pp. 40–50). Singapore: World Scientific.
- Di Ferdinando, A., Borghi, A. M., & Parisi, D. (2002). The role of action in object categorization. In S. Haller & G. Simmons (Eds.), *Proceedings of the 15th international affairs conference* (pp. 138–142). Pensacola, FL: AAAI Press.
- Di Ferdinando, A., & Parisi, D. (2004). Internal representations of sensory input reflect the motor output with which organisms respond to the input. In A. Carsetti (Ed.), *Seeing and thinking* (pp. 115–141). Dordrecht: Kluwer.
- Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences*, 20, 1–55.
- Goldstone, R. L. (1994). The role of similarity in categorization: Providing a groundwork. *Cognition*, 52, 125–157.
- Hampton, J. (1982). A demonstration of intransitivity in natural categories. *Cognition*, 12, 151–164.
- Malt, B. C., Sloman, S. A., Gennari, S., Shi, M., & Wang, Y. (1999). Knowing versus naming, similarity and the linguistic categorization of artifacts. *Journal of Memory and Language*, 40, 230–262.
- Mandler, J. M. (1992a). The foundations of conceptual thought in infancy. *Cognitive Development*, 7, 273–285.
- Mandler, J. M. (1992b). How to build a baby: II. Conceptual primitives. *Psychological Review*, 99, 587–604.
- Mandler, J. M. (1998). The rise and fall of semantic memory. In M. A. Conway, S. E. Gathercole, & C. Cornoldi (Eds.), *Theories of memory* (Vol. II, pp. 147–169). Hove, UK: Psychology Press.
- Mandler, J. M., & Bauer, P. J. (1988). The cradle of categorization: Is the basic level basic?. *Cognitive Development*, 3, 247–264.

- Mandler, J. M., & McDonough, L. (1998). On developing a knowledge base in infancy. *Developmental Psychology, 34*, 1274–1288.
- Mandler, J. M., Bauer, P. J., & McDonough, L. (1991). Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology, 23*, 263–298.
- Markman, E. M. (1989). *Categorization and naming in children: Problems of induction*. Cambridge, MA: MIT Press.
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, M. D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382–439.
- Sloman, S. A. (1998). Categorical inference is not a tree: The myth of inheritance hierarchies. *Cognitive Psychology, 35*(1), 1–33.
- Smith, L. B. (1995). Stability and Variability, the geometry of children's novel-word interpretations. In F. D. Abraham & A. R. Gilgen (Eds.), *Chaos theory in psychology*. London, UK: Praeger.
- Smith, L. B., & Katz, D. B. (1996). Activity-dependent processes in perceptual and cognitive development. In R. Gelman & T. K. Au (Eds.), *Perceptual and cognitive development* (pp. 413–445). San Diego, CA: Academic Press.
- Thelen, E., Schöner, G., Scheier, C., & Smith, L. B. (2001). The dynamics of embodiment: A field theory of infant perseverative reaching. *Behavioral and Brain Sciences, 24*, 1–86.
- Vallée-Tourangeau, F., Anthony, S. H., & Austin, N. G. (1998). Strategies for generating multiple instances of common and ad hoc categories. *Memory, 6*(5), 555–592.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review, 9*, 625–636.