

Towards a Vygotskian Cognitive Robotics: The Role of Language as a Cognitive Tool¹

Marco Mirolli*, Domenico Parisi

Institute of Cognitive Sciences and Technologies, CNR

Via San Martino della Battaglia 44, 00185 Rome, Italy

Abstract

Cognitive Robotics can be defined as the study of cognitive phenomena by their modeling in physical artifacts such as robots. This is a very lively and fascinating field which has already given fundamental contributions to our understanding of natural cognition. Nonetheless, robotics has to date addressed mainly very basic, low-level cognitive phenomena like sensory-motor coordination, perception, and navigation, and it is not clear how the current approach might scale up to explain high-level human cognition. In this paper we argue that a promising way to do that is to merge current ideas and methods of 'embodied cognition' with the Russian tradition of theoretical psychology which views language not only as a communication system but also as a cognitive tool, that is by developing a Vygotskian Cognitive Robotics. We substantiate this idea by discussing several domains in which language can improve basic cognitive abilities and permit the development of high-level cognition: learning, categorization, abstraction, memory, voluntary control, and mental life.

1. Introduction

We can construct robots for at least two different reasons: as useful artifacts or as scientific tools. The use of robots as scientific tools is an instance of a new approach to science. Traditionally, science has been trying to understand reality through systematic observation and experiment. This can be called the 'analytic approach' to science. But since the advent of the computer in the late 1940's a new kind of approach to science has appeared, one which tries to understand reality by modeling it in computers or robots. This is what can be called the 'synthetic approach' to science. The rationale for the synthetic approach is that, once you have built a system that reproduces some phenomenon of reality, you have a candidate explanation of that phenomenon in that it is possible – even though by no means certain – that the principles you have used to build your artificial system are the same principles that underlie the real phenomenon and explain it.

The use of the synthetic approach is today very common in many scientific fields. But the synthetic approach is particularly important in Psychology since it is at the core of one of the disciplines that contributed to the birth of modern Cognitive Science, namely Artificial Intelligence (Bechtel et al., 1998). But while classical cognitive science and artificial intelligence were based on the mind-as-a-computer metaphor, in the last twenty years robots have been replacing the computer as icons of the human mind. In fact, in the last twenty years a number of overlapping fields of investigation have emerged whose aim is to understand behavior by trying to reproduce it in artifacts, like robots, which emphasize the physical and biological basis of organisms: artificial life, evolutionary robotics, adaptive behavior, ecological neural networks, embodied and situated cognition etc (Parisi, Ceconi, and Nolfi, 1990; Varela et al., 1991; Clark, 1997; Pfeifer and Scheier, 1999; Nolfi and Floreano, 2000).

This kind of research has already provided fundamental contributions to our understanding of behavior and mind, in particular by pointing to the importance for cognitive processes of the physical, dynamical interactions between the organism and its environment. But current robotic research is still quite far from addressing, let alone explaining, the high-level cognitive capacities which are characteristic of human beings. In this paper we argue that a promising way for bridging

¹ Published as: Mirolli Marco, Parisi Domenico (2011). Towards a Vygotskian Cognitive Robotics: The Role of Language as a Cognitive Tool, *New Ideas in Psychology*, 29: 298-311. doi: 10.1016/j.newideapsych.2009.07.001

* Corresponding author. e-mail: marco.mirolli@istc.cnr.it

the gap between today's research in robotics and embodied cognition and a full understanding of human cognition is by developing a Vygotskian Cognitive Robotics, that is, by endorsing and developing, within robotic research, the idea of language as a cognitive tool which originated in the Russian tradition of Theoretical Psychology.

In the next section we clarify what we mean by 'Cognitive Robotics' and we briefly highlight the significance of current research in Cognitive Robotics by framing its role with respect to general trends in contemporary Psychology and Cognitive Sciences. In section 3 we briefly present the general Vygotskian idea of language as a cognitive tool. In section 4 we discuss a number of ways in which language can transform (improve) cognitive processes like learning, categorization, abstraction, memory, voluntary control, and mental life. Finally, section 5 draws some general conclusions.

2. Cognitive Robotics as a new approach to Cognitive Science

First of all, it might be useful to make a terminological clarification. The term 'Cognitive Robotics' can have at least two different meanings: (1) the use of robots as research tools for studying cognition; (2) the use of robots for studying 'truly cognitive' phenomena, in contrast with 'mere' sensory-motor interactions with the environment. According to (1), all kind of work which has the scientific goal of understanding behaviour by reproducing it in physical artefacts (robots) can be considered as cognitive robotics. In contrast, the second reading of the term 'cognitive robotics' is more restrictive. For example, Clark and Grush (1999, p. 12) consider as 'truly cognitive' only those phenomena "that involve off-line reasoning, vicarious environmental exploration, and the like". Cognitive robotics in this more restrictive interpretation currently includes only a very few lines of research, and in fact Clark and Grush's aim was to push roboticists 'towards' a cognitive robotics, which in 1999 was, in their view, almost an empty field.

Here we do not adhere to Clark and Grush's very restrictive definition of 'cognition'. Rather, as many other researchers, both within (Beer, 2003; Harvey et al., 2005) and outside robotics (Thelen and Smith, 1994; Kelso, 1995), we prefer to be more liberal and consider as genuinely 'cognitive' all kinds of ways of solving adaptively valuable tasks, which are usually, if not always, based on direct, on-line, interactions with the environment. The reason is that by using the same term to refer to both low-level and high-level behavioural capacities we are more prepared to discover both similarities and differences among behavioural capacities and to examine how high-level capacities emerge both evolutionarily and developmentally from low-level ones. However, we recognize that clearly not all problems, and all different ways of solving problems, are equally complex and 'cognitive'. We think that a continuum exists between low-level cognitive functions and high-level ones, and it is certainly true that robotic approaches to cognition have been focusing mostly, if not exclusively, on low-level behaviors and capacities, such as perception, sensory-motor coordination, and navigation.

This has both a practical and a theoretical explanation. The practical explanation is just common sense: first study what is simple, and then move to the more complex. Since our understanding of even low level cognitive phenomena is still rather poor, it is better to focus on low level phenomena and then move on to high level ones. But this simple, common-sense rule has also a theoretical counterpart. Since the cognitive complexity continuum also has a temporal significance, with lower-level competences preceding (both phylogenetically and ontogenetically) higher-level ones, the latter are most likely to be built upon the former and to depend on them. Hence, it is important to first understand the more basic forms of cognition and ground the study of more complex forms on these more basic forms. This implies adopting a "genetic" approach according to which in order to understand the nature of X it is crucial to reconstruct how X has become what it is. This is an approach clearly related to Piaget's genetic epistemology (Piaget, 1968) and contrary to the approach of "good old-fashioned artificial intelligence" (GOFAI), which tried to understand cognition by directly focusing on the most high-level human cognitive functions (logical theorist, chess playing, complex problem solving and the like) and ignoring all the rest.

In fact, the shift of focus in contemporary research on cognition from high-level to low level cognitive capacities has been paralleled, or, more correctly, guided by a theoretical shift from the symbol-manipulation paradigm of classical cognitive science to the sub-symbolic, embodied, situated, and distributed approaches to cognition which characterize current research on cognition (Bechtel et al., 1998; Clark, 2001). And there is no doubt that robotics has been playing a major role in the generation and development of these new ideas.

After the connectionist revolution of the mid-eighties (Rumelhart and McClelland, 1986), which still involved disembodied and non-situated neural networks, in the early 1990s behavior-based robotics (Brooks, 1991) and the use of neural networks in an Artificial Life perspective (Parisi et al., 1990) set the stage for the development of a new theoretical framework. The 'Artificial Life route to Artificial Intelligence' (Steels and Brooks, 1994) pointed to the fact that natural cognitive processes are always 'embodied', 'situated' and (partially) 'distributed' in an organism's environment. They are embodied in that the body and its physical properties are critical for the way a given task is solved. They are situated because the constraints provided by the environment can also act as opportunities for the task's solution. And they are partially distributed because they do not happen only inside an organism's head, but they substantially depend on the organism's environment, which, especially in the human case, includes also artefacts and other agents. These ideas have converged with several different lines of previous work in theoretical psychology (i.e., Gibson, 1979; Bickhard, 1980; Norman, 1980), and have been inspiring a lot of important work in contemporary cognitive science (i.e., Hutchins, 1995; Barsalou, 1999; Borghi et al., 2004).

Notwithstanding the importance, from both the practical and the theoretical point of view, of the shift from high-level to low-level cognitive capacities in cognitive science research, still, the ultimate goal of psychology, and of the cognitive sciences more generally, is to understand human cognition, and a skeptic could still claim that it is not clear how far the new ideas of embodied, situated, and distributed cognition can go in explaining the more complex phenomena of the cognitive continuum. What is the relationship between sensory-motor coordination, active perception, low-level adaptive behavior, on one side, and high-level, human-specific cognitive capacities of abstract reasoning, complex decision making, logico-mathematical capacities, reflexive thinking, on the other? How can we bridge the gap currently existing between Embodied Cognition principles, ideas, and models, and human cognitive capacities?

3. Robotics meets Vygotsky: language as a cognitive tool

In line with other recent proposals (i.e., Clark, 1998, 2006; Zlatev, 2001; Lindblom and Ziemke, 2003; Clowes and Morse, 2005), here we argue that the most promising way of addressing high-level, specifically human cognitive capacities using the new embodied, situated, distributed, interactive, view of cognition, is to develop a Vygotskian cognitive robotics, i.e., to put together the cognitive robotics approach with the ideas developed in the first half of the XX century by the Russian psychologist Lev Vygotsky, who stressed the importance of language for human cognition. (For other, slightly different, attempts to outlining a Vygotskian robotics see also Dautenhahn and Billard, 1999; Kulakov and Stojanov, 2002).

Vygotsky (1962, 1978) claimed that the most important event in human development is when two previously unrelated lines of development, that of practical abilities and that of language, converge. According to Vygotsky, language or, more properly, linguistically mediated social interactions, cause a radical transformation of elementary cognitive abilities into the high-level, specifically human, psychological functions. This happens thanks to a process of internalization of the social interactions and relationships which the child entertains with adults and more skilful peers in the course of development. It is through this process of internalization that the child develops an ability to accomplish increasingly complex cognitive tasks. The basic idea is the following. When the child is challenged by a task which she cannot solve but through the help of an adult or a more skilled peer, she asks for help, which typically takes the form of linguistic help. Later on, when the child is facing the same or a similar task all alone, she can rehearse the social linguistic aid which

helped her to solve the problem. This is what is called 'private speech'. The linguistic social aid coming from adults takes several different forms. Social language helps a child in learning how to categorize experience, to focus her attention on important aspects of the environment, to remember useful information, to inhibit irrelevant, spontaneous behavior, to sub-divide challenging problems into easier sub-problems, to construct a plan for solving complex tasks, and so on and so forth. When the child is talking to herself she is just doing to herself what others do to her, that is, providing all sorts of cognitive aids through linguistic utterances. Once the child has mastered this linguistic self-aid, private speech tends to disappear, but only apparently. In fact, it is just abbreviated and internalized, becoming inner speech. Hence, most, if not all, high-level human cognitive processes are linguistically mediated, in that they depend on the use of language for oneself.

These ideas have been developed in the Russian psychological tradition (i.e. Solokov, 1975; Luria, 1979) and elsewhere (i.e. Cole, 1996; for a review, see Diaz and Berk, 1992). But the Vygotskian view of language as a cognitive tool has been largely ignored in mainstream cognitive science, principally because Vygotskian works began to be translated in English only in the 1960s, when developmental psychology was dominated by the more individually oriented theory of Jean Piaget. Of course, Piaget recognized the importance of sociality in human behaviour but, at least in explaining cognitive development in children, he put more emphasis than Vygotsky on internal organizational processes and on the child's interactions with the non-social environment. More recently, the Vygotskian idea of language as a cognitive tool and its role in enhancing cognitive functions has been raising increasing interest both in cognitive-science oriented philosophy of mind (Dennett, 1991; Carruthers, 2002; Clark, 1998, 2006), and in several areas of psychology (see, for, examples, Gentner, 2003; Spelke, 2003; Tomasello, 2003).

Of course, language did not arise suddenly in its modern form during human evolution, nor is it magically acquired during child development. On the contrary, language presupposes a number of evolved, specifically human, cognitive capacities that made it possible for it to take its modern form during evolution and make it possible for it to be easily acquired during development. In particular, human language has probably gone through various stages of evolution during which it has co-evolved with human non-linguistic cognition, exploiting new and more complex forms of cognition: for example, an increased ability to predict the consequences of one's own action (Parisi and Mirolli, 2006), or the ability to self-generate sensory-motor inputs inside the brain, without relying on sensory-motor input from the external environment (Parisi, 2007). This co-evolution of linguistic and non-linguistic abilities in hominids has most probably led to significant differences between human and non-human non-linguistic cognition, in contrast with what Vygotsky appears to have assumed. However, as soon as language began to be used (or begins to be used in child development), it had a fundamental influence on all sorts of human cognitive functions, in such a way that it may be impossible to develop a human-like cognitive robotics without endowing robots with the capacity of using language for themselves as humans do.

4. Towards a (more) Cognitive Robotics

In this section we discuss several ways in which language can improve basic cognitive processes thus permitting the development of the high-level cognitive functions which are characteristic of humans. We do this by both reviewing (some of) the existing empirical and computational literature on the topic and developing new ideas and hypotheses. Hence, some of our ideas have some empirical and computational support, while others are of a more speculative nature. In any case, we consider them as promising starting points for future robotic research which aims at understanding high-level cognition from an embodied and situated perspective.

4.1 Learning

The importance of linguistic labels on individual learning is surely the most studied effect of language on cognition, both through empirical investigations and through computational modeling.

From an empirical point of view, in the recent literature on language acquisition there exists an important line of research which has repeatedly and consistently demonstrated that language can facilitate category learning. Several studies with subjects of different ages (from 9 month old children to adults) have in fact demonstrated that providing linguistic input to somebody who is learning to categorize objects can substantially ease and speed up the learning process (see, for example, Waxman and Markow, 1995; Nazzi and Gopnik, 2001; Yoshida and Smith, 2005). These findings strongly suggest that labels have the function of 'inviting' category formation, by guiding our attention onto 'meaningful' aspects of our environment and by providing important cues about how to categorize them. Nonetheless, the field is only at its very beginnings and much more work has still to be done in order to provide a full account of the mechanisms that underlie the role played by language in category learning.

The effects of language on category learning has also been the subject of several computational models. For example, Schyns (1991) and Lupyan (2005) have shown with neural network simulations how linguistic label can simplify category learning. In fact, providing neural networks with labels accompanying perceptually presented objects during learning has been shown to speed up category learning (Schyns, 1991) or to improve internal representations of objects, specifically of those categories of objects which are more difficult to learn (Lupyan, 2005). In a language-game model, Steels and Belpaeme (2005) have shown how the self-organized emergence of a linguistic system in a collection of agents can co-evolve with the process of categorization of perceptual experiences because of a structural coupling between the 'conceptual' and the 'linguistic' systems. In other words, while a populational linguistic system self-organizes, the agents' conceptual system adapts itself in order to maximize communicative success. Finally, using artificial life simulations, Cangelosi and colleagues (Cangelosi and Harnad, 2000; Cangelosi et al., 2000) have shown how organisms with language can learn to categorize their experience in adaptive ways not only through genetic evolution or individual learning by trial and error, but also through social learning, with what they call 'symbolic theft'. In the symbolic theft condition learning happens thanks to (a) a pre-existing ability to categorize some of the stimuli, and (b) exposure to others' language, which incorporates information on how to categorize new experiences. The results of those simulations have shown that symbolic theft can give an adaptive advantage with respect to standard phylogenetic or individual learning in that it is both significantly faster and less dangerous (you don't risk to suffer the cost of errors).

But why should category learning be facilitated by language? What are the mechanisms which may underlie this effect? We argue that the facilitatory effect of labels on category learning derives from the following two mechanisms: (a) linguistic inputs constitute additional stimuli that focus the learner's attention to the specific aspects of perception that are relevant for categorization, and (b) language itself can sometimes represent the principal (or even the only) ground on which the learner can develop the discriminative capacities that constitute categorization. Let's explain these two points in order.

Hearing the same linguistic stimulus, let's say the word 'red', when perceiving red cars, red apples and red flowers facilitates -- or may even induce -- the acknowledgement that all those different stimuli have something in common, namely the red color. In neural network terms, this means that the occurrence of the same activation pattern in the acoustic input units -- namely the pattern that corresponds to the sound 'red' -- increases the similarity of the internal representations of all red stimuli, and this in turn can help -- or induce -- the network itself to learn that all those stimuli belong, to some extent, to the same category, namely that of red things. A similar point has already been made in the empirical literature on the topic (see, for example, Waxman, 2004).

But this is not necessarily the whole story. Though fundamentally correct, this account is partial, in that it assumes a substantially 'passive' view of perception and cognition in general. The point is that perception depends in a fundamental way on action (similar action-based views of perception and cognition in general have been developed, among others, by Gibson, 1979; Bickhard, 2001; Di Ferdinando and Parisi, 2004; Gallese and Lakoff, 2005). To categorize means to produce a given

behavior A when perceiving a certain class of stimuli and another behavior B when perceiving another class of stimuli. It is the need to respond appropriately (and discriminately) that makes an organism perceive the first class of stimuli as different from the second one. So, when dealing with a categorization process, we must always ask the following question: which is the differential behavior exhibited by an agent that makes us say that the agent is categorizing some experiences as belonging to the same category and other experiences as belonging to another category? We argue that, for many human categories the answer might just be: the production of different words! In other words, in many cases the human brain learns to represent some patterns of inputs (for example those produced by red cars, red apples, and red flowers) as similar to each other and different to other patterns (those produced by yellow cars, yellow apples, and yellow flowers) principally because it is learning to produce, through its phono-articulatory output units, the same (or different) action, which consists in the production of the same (or different) word: 'red' (or 'yellow'). This is not to deny that there is some internal 'appreciation' of different colours, nor that there is no genetic tendency to discriminate colours: we all know, from direct experience, that we have different appreciation of redness and yellowness, and we know from scientific investigations that there is some genetically based capacity to discriminate colours in certain ways. The point is that part of the specific way a given human being categorizes experiences (in this case, related to colour) is in fact due to the way s/he learns to name them. In other words, according to this view, the amazing amount of categories that humans can have is in great part due to the role played by language in providing a behavioural ground for categorization.

4.2 Categorization

The effect of language on learning is surely very important, and deserves further attention, but it still cannot be considered as a demonstration of the Vygotskian thesis of language as a cognitive tool. It is rather a demonstration of the power of social learning: in the special case, which is clearly fundamental for human beings, in which social learning is mediated by language. But language does not only facilitate category learning. It can also improve categorization once categories have already been learned. We have recently demonstrated this with a neural network model (Mirolli and Parisi, 2005a, 2006).

We modeled the 'child's brain' as a modular neural net, composed of two sub-networks (which we call 'sensory-motor' and 'linguistic', respectively) with reciprocal interconnections between the two layers of hidden units. In a first phase (from birth to 1 year) the sensory-motor net learns to respond to perceived objects by producing motor actions which are appropriate to objects' categories, while the linguistic net learns to imitate heard sounds (words). In a subsequent phase (from 1 year on), the connection weights which link the two sub-networks are trained so that the whole network learns to associate the internal representations of perceived objects (i.e., the patterns of activation of the sensory-motor hidden units) with the internal representations of instances of the names of the appropriate categories (i.e. the patterns of activation of the linguistic hidden units). After learning the network could both correctly name perceived objects (i.e., produce an instance of the correct word through its linguistic output units) and comprehend perceived words (i.e., produce the appropriate action through its sensory-motor output units). We then analyzed the effects of the learned mapping between objects and their names on categorization.

In a neural network, the quality of categorization can be quantitatively measured in the following way. We can consider the particular activation pattern observed in the network's hidden units at any given time as one specific point in an abstract hyperspace with as many dimensions as the number of hidden units, where the coordinate of the point for each dimension is the activation level of the corresponding unit. Categories are 'clouds' of points in this abstract hyperspace, that is, sets of points elicited by sensory inputs that must be responded to with the same motor output. Different categories are different clouds of points. Good categories are clouds of points that are (a) small (objects that must be responded to with the same action are represented as similar), and (b) distant from each other (objects that must be responded to with different actions are represented as

different). Property (a) can be measured as the average distance between all the points belonging to a category and the category centroid, while property (b) can be measured as the distance between the categories' centroids.

Our model demonstrated that learning the mapping between pre-linguistically learned concepts and linguistic labels changes the internal representations of objects. In fact, after having learned the mapping, the internal representations of the non-linguistic, sensory-motor network tend to be influenced not only by the non-linguistic sensory inputs but also by the linguistic network. This has the effect that the non-linguistic categories tend to become better categories, in the sense that the clouds of internal representations of objects belonging to different categories become smaller and more distant from each other. And since an organism's categories influence the organisms behaviour by making it easier for the organism to select the appropriate action in response to sensory inputs, an organism endowed with language will have a more effective behaviour. In our model we demonstrated this facilitatory effect of language on categorization not only in the case in which language is social, that is, when perceived words are produced by other individuals (Mirolli and Parisi, 2005a), but also when language is self-produced, that is, when the agent (the network) talks to itself. And this happens both when talking to oneself is external, as in private speech, and when it is internal, as in inner speech (Mirolli and Parisi, 2006). These results represent, to the best of our knowledge, the first example of a computational model which directly supports the Vygotskian idea that social language can improve basic individual cognitive abilities through a process of internalization of linguistically-mediated social aid.

4.3 Abstraction

The effects of linguistic labels on internal categorization have further important consequences with respect to the process of abstraction. Categorization *requires* abstraction. In order to respond in the same way to different stimuli which belong to the same category you need to abstract from their differences. And, vice versa, in order to respond in different ways to similar stimuli which belong to different categories you have to abstract from their similarities. So, reducing the size of a category's cloud is improving the first kind of abstraction (ignoring differences between intra-category stimuli), while increasing the distance between clouds is improving the second kind of abstraction (ignoring the similarities between inter-category stimuli).

The model just discussed had one fundamental simplification: each object had one and only one specific action associated with it. This is a clear limitation, since for real organisms typically the same object can evoke several different responses, depending on the context. For example, an apple can be eaten, thrown, given to another person, and so on. In other words, depending on the context, the same object can be categorized in different ways – in the apple example, as food, weapon, or present. This means that the internal representation of an object must be multi-functional, in the sense that it must allow the organism to consider the same object as belonging to different categories depending on the circumstances. Linguistic labels can help organisms to abstract away from the possible ways in which an object can be categorized which are not relevant from the current situation, and focus only on the categorization which is relevant.

Furthermore, labels can also induce hierarchical categorization. Representations can be hierarchically organized in the sense that two sets of sensory inputs can be responded to by two different actions, and therefore they constitute two distinct clouds of points, but there is a third action with which the organism responds to both sets of sensory inputs. Therefore there is a third cloud of points that includes both the first and second clouds of points. Language can favour the creation of hierarchies of clouds of points just because it provides hierarchies of labels: there are two linguistic signals, e.g., "dog" and "cat", that correspond to two distinct clouds of points, and there is third linguistic signal, "pet", that evokes the point located centrally in the larger cloud of points including the "dog" cloud and the "cat" cloud .

The idea that labels can play a critical role in the process of abstraction has received empirical support in the contest of relational thinking, with studies on both humans and chimps (for a review,

see Gentner, 2003). In two series of experiments with children, Gentner and her colleagues (Rattermann and Gentner, 1998; Loewenstein and Gentner, 2005) clearly demonstrated that the use of relational language helped children to solve analogical (relational) mapping tasks across a wide range of ages and task difficulties. Indeed, Gentner et al. provided evidence for the Vygotskian hypothesis that linguistic aid undergoes a process of internalization: while younger children need to be provided with relational language even for solving simple tasks, older children do not, but they need linguistic help if confronted with more difficult tasks. These studies clearly point to the importance of language for reasoning on abstract (relational) properties of the world. In fact, it seems that acquiring and using a name for describing a relational pattern helps the child to abstract that pattern from the concrete context in which it has been experienced and thus it increases the probability that the same abstract pattern is recognized the next time it is encountered. In other words, labeling an abstract (in this case, relational, but the point can be generalized to any kind of abstract labels) property changes the perceptual apparatus of the child, in that it can render the property in some sense directly perceivable.

The most impressive evidence for this view comes from empirical work on chimpanzees (Thompson et al., 1997, see also Oden et al., 2001). Chimps (as several other animals) can learn quite easily to succeed in a match-to-sample task: that is, to choose, among two different objects, the object which matches a sample object (given an A as the sample, the chimp has to choose an A against a B). Normal, i.e., non-linguistically-enculturated, chimps were not able to learn and solve a relational-matching task: that is, to choose, among two pairs of objects, the pair whose objects are in the same relation as the ones of the sample pair (given an AA as sample, the chimp has to choose a BB pair against a CD one). Most strikingly, chimps which had been previously trained to use two different symbols for the two relations 'sameness' and 'difference' were able to solve the relational-matching task. Note that in order to solve the relational-matching task chimps must apply the same/different distinction at the relational level, that is, at the level of the relation between objects. So, in order to solve this second-order problem (judging relations between relations) all it is needed is to reduce it to a first order one (judging relations between objects), which we know chimps are able to solve. And this is exactly what language training seems to do: practicing with relational symbols seem to change chimps' perception, so that linguistically trained chimps are able to 'see' the relation holding between two objects so that they have just to decide whether two perceived relations are same or different (beside the original paper of Thompson and colleagues, see Clark, 2006 for a similar analysis of the same experiments).

These studies seem to suggest that training in linguistic tasks changes the way an agent perceives its world. The interactions with the world of a linguistically trained animal is mediated by linguistic forms, which render some of the aspects of experience more salient than others. The process is recursive: once you have learned to see the world in certain ways you can also discover new, more abstract patterns. For example, recognizing that sometimes you are looking at the colour while in other occasions you are looking at the form of objects, thus allowing the development of more and more abstract concepts, like the concepts of 'colour' and 'form'. In other words, a non-linguistically trained animal which is able to react appropriately to stimuli of different colours (red vs. yellow) has surely different representations of those stimuli. But it probably doesn't have the concept of 'redness' vs. 'yellowness', and it surely doesn't have the even more abstract concepts of 'colour' vs. 'form', not to speak about the still more abstract concept of 'property' (of which 'colour' and 'form' are two instances). Our abilities of constructing more and more abstract categories depends on our ability to label discovered categories and hence reason on those categories themselves. The reason is that the language user can apply his/her cognitive processes not only to simple perceptual experiences, but also to the 'concepts' themselves, which are made perceivable through the linguistic labels.

4.4 Memory

Another example of the importance of language for cognition concerns memory. For example,

Clowes and Morse (2005) presented artificial life simulations in which agents had to respond to external commands in order to perform a simple task. By comparing the results of three different experimental conditions, they showed that agents who could repeat to themselves heard commands performed significantly better than agents who could not. Similarly, with a simple artificial life model we have investigated the effect of using signals as individual memory aids on the evolution of a simple communication system (Mirolli and Parisi, 2005b). In particular, our simulations show that using signals not only for communicating information to other individuals, but also for remembering the information received by other organisms, the evolution of signals is favoured, and this has a clear positive impact on the organisms' fitness as well.

These are just the very first steps in the study of the importance of language for memory processes. For example, in contrast with what happens in our simplified models, in real brains the non-linguistic and the linguistic networks have a significantly different size. Specifically, the non-linguistic network appears to be much larger than the linguistic one, in terms of number of units and connections of which the two networks are composed (Mirolli et al., 2007). This simple fact provides a clear advantage for the specifically human linguistic memory system with respect to the older non-linguistic memory system which we share with other animals. In other words, it is generally easier to remember words than the actual sensory-motor experiences to which words are associated. This, in turn, has the consequence that an organism possessing language can work more easily with linguistic (sound) information and translate this information into the associated non-linguistic information when necessary. Furthermore, possessing a linguistic memory system in addition to the older sensory-motor one has a second, fundamental advantage: delegating the memory function to the linguistic system leaves the sensory-motor system free to process other information useful for acting in the environment while linguistically remembering previous information.

The importance of language for human memory has been also demonstrated by empirical research. Both psychological and neuroscientific evidences in fact demonstrate that humans have at least two distinct working memory systems: the first is a multi-modal system which we share with non-human primates; the second is linguistic, and hence species-specific (see, for example, Baddeley, 1992; Petrides et al., 1993; Becker and Morris, 1999). Recent neuro-imaging studies have also been discovering the neural basis of these memory systems, suggesting that the linguistic memory system is subserved by mostly left-hemispheric areas which underlie normal (audible) speech (see, for example, Gruber, 2002; for a detailed review see Gruber and Goschke, 2004). Furthermore, in many circumstances verbal memory seems to be more efficient and flexible than the older multi-modal system, and therefore it seems to be the predominant rehearsal mechanism, although it may function in cooperation with the older multi-modal system. For example, a number of studies using different experimental paradigms have consistently shown that articulatory suppression significantly increases the difficulty of a task by making it more difficult to retrieve the task's goal (see, for example, Baddeley et al., 2001; Emerson and Miyake, 2003; Miyake et al., 2004). This strongly suggests that normally it is inner speech that is used for retrieving and activating relevant information for solving a given task.

The advantages of possessing a linguistic memory system for short-term processes may extend to long-term memory. Instead of memorizing full experiences a human being can label them and memorize their verbal description. Thanks to the abstracting power of language, what is to be remembered now are just the most relevant features of a given experience. And thanks to the small size of the linguistic network, those relevant features are coded in a very efficient way, so that they can be easily memorized and recalled. It may well be that it is thanks to the possibility of this linguistic coding, memorizing, and recalling that human beings are able to recall events which happened in their very distant past. This hypothesis could also explain why our memories never go too far in our past. Typically, the first memories never date back before about the third or fourth year of life. This is just the age at which private speech begins to appear in the child (Berk, 1994). It is possible that we can't recall anything about our first three or four years of life because in those years we haven't learnt to label our own experiences and to memorize them in verbal form yet.

4.5 Voluntary control

Finally, another fundamental consequence of the internalization of linguistic stimulation for the purpose of guiding one's actions is the development of voluntary control. As discussed above, an animal's ability to voluntarily control its own actions is very limited. Non-human animals are fundamentally stimulus-driven: they are both very limited in their capacity to include non-perceived objects in their problem-solving, and easily distracted by non-relevant stimuli, in that they can hardly inhibit their instinctual responses to highly motivating stimuli. Of course, under the same conditions human beings experience the same kind of difficulties, but we are able to overcome them in ways which are not accessible to other animals. We can control our behavior, we can focus on our tasks, and we can inhibit our instinctual responses to even the most motivating stimuli. And we develop these abilities through the internalization of suggestions and commands which we receive from other individuals -- most notably, our parents -- during our infancy.

The idea is once more the same. The behavior of the child is constantly controlled, through linguistic stimulation, by other people: during all our infancy, we are continually instructed about all kinds of do's (wash your hand, clean your teeth, do your home-work, etc.) and don't's (don't get dirty, don't eat sweets, don't watch too much TV, etc.). Once an individual has experienced the positive effects of being guided by linguistic stimuli produced by other individuals, she learns, by imitation, to linguistically stimulate herself in the same way to produce the same effects. That is, she starts to talk to herself as a means for controlling her own behaviour. It is in this way, we argue, that we learn to do what we know is important for us but we don't like too much to do and, conversely, to prevent ourselves from doing what we are motivated to do but we know we shouldn't. This is what, later on, we call voluntary control, or Will.

The fundamental role that language (in the form of inner speech) plays in the voluntary control of action is demonstrated by empirical work on the effects of articulatory suppression on the task-switching cited above. In fact, these studies consistently show that talking-to-oneself is essential when we engage in goal-directed behaviour. But besides providing an efficient means for remembering task information (the goal you are pursuing), language can help the development of voluntary control also as a powerful means of abstraction. Recent studies on language-trained chimpanzees seem to provide this kind of evidence (Boysen et al., 1996). Boysen and her colleagues presented chimps with two bowls containing different numbers of candies. The animal is given the bowl which it doesn't point to. Hence, to get the most rewarding result the chimp has to point to the bowl containing the smaller number of candies. Surprisingly, chimps never learned to do this. But if the same chimps, which had previously been taught symbols referring to numerals, were presented with numerical symbols instead of candies, they quickly learned to point to the smaller numerical symbol in order to get the larger quantity of candies. This result seems to demonstrate that the use of (numerical) symbols can enable chimps to control their otherwise overwhelming food-related responses. And this is possible just because of the abstractness of the symbols. When presented with symbols, in fact, the chimps could just focus on the information which is relevant for solving the task, while when presented with real food, the richness of the sensory stimulus prevented them from inhibiting their responses in order to reason on the strategy to adopt.

4.6 Mental life

Finally, language might play a major role in the most striking and peculiar characteristic of human mind, that is, human mental life (as argued also by Dennett, 1991). Human beings have a very rich mental life which includes visual and motor images, rememberings, dreams, hallucinations, and so on. Mental life can be considered as the self-generation of one's own input (Parisi, 2007). For example, mental images are self-generated input (typically visual input, but we can imagine any kind of sensory-input, including proprioceptive ones) which generally we have not recently received but is actively produced by the nervous system itself; rememberings are self-generated input which we know we have not received recently but we also know that we have received in the

past; dreams are self-generated inputs that occur when one is asleep; hallucinations are self-generated input that we erroneously believe is coming from the external environment; and so on. The first and most obvious reason why language plays a fundamental role in human mental life is simply because a great part of our self-generated input is linguistic: in other words, a large part of our mental life is constituted by internally talking to oneself, that is, by inner speech. But the role of language in mental life might go well beyond this. In what follows we will try to explain why.

Learning to self-generate one's own input appears to be associated with learning to predict the consequences of one's own actions. By learning to self-generate one's own input one can learn how to control one's own behaviour. Afterwards, this capability can become internal by means of the association of the context to the sensory-motor trace of the effects of one's own actions. In other words, after I have learned to induce to myself a sensory-motor experience by manipulating the external environment, I can learn to do the same but just internally by predicting the sensory-motor experience I would have if I would do the appropriate actions. This is how we learn to internally self-generate our own input. Indeed, it may well be that we learn to do this first with words, and then with all other modalities, because language constitutes the easiest domain for learning to predict the consequences of one's actions.

The linguistic system is different from the other sensory-motor systems because in the linguistic system the mapping between input and output is more direct, complete, and stable: when you produce a sound, you can predict very reliably the acoustic input that you will receive as a result. This is not true for other kind of sensory-motor systems. In fact, it is true only for very simple movements (you can reliably predict the proprioceptive input given some motor command) and in fact there are many circuits devoted to this kind of prediction (Kawato, 1999; Wolpert and Flanagan, 2001) whose principal function is to provide self-generated proprioceptive feed-back in cases in which the real feed-back would be too slow for the behaviour to occur properly (Clark and Grush, 1999). But the input which is self-generated by internal models is just proprioceptive, so it is useful just in the sensorymotor coordination of the organism itself, not for other more complex kinds of activity.

The situation for the linguistic system is different, for several reasons. First, the linguistic input which we receive as a consequence of our phono-articulatory actions is almost completely determined by the phono-articulatory action itself. Second, these consequences of our actions are always present (we cannot but hear what we say), and, thanks to the importance of language for our cognition, we always pay attention to linguistic input. This makes learning to predict our own linguistic stimuli given our linguistic action very straightforward. Furthermore, linguistic tokens have meanings, in the sense that they are associated with other sensory-motor experiences which are relevant to the organism and which they tend to restore. This simple fact renders the self-generation of linguistic stimuli particularly important in that self-generating a word will cause the self-generation of a non-linguistic sensory-motor experience which is associated to the word. In other words, both the association between the phono-articulatory movements and the resulting acoustic sounds and the association between a given sound (word) and its meaning (that is the sensory-motor experience associated with that word) are very systematic, reliable, and almost immediate for a developing child. This makes it easy for the child to learn to self-generate his or her own (non-linguistic) input by simply producing the words which are associated to the particular sensory-motor experience he or she wants to have.

On the other hand, how can a non-linguistic animal self-induce some experience? On the one side, it can self-generate proprioceptive inputs by its internal models; on the other side, it can control the input it receives by (overtly) directing its attention towards the desired objects. But how could it learn to direct its attention towards something which is not in the immediate surroundings? And hence, how could it learn to have mental images or rememberings, that is, to self-generate the experience of something which is not present here and now?

There is no real (external) action you can do in order to let you perceive the visual image of the Coliseum if you are not in front of it or very close to it. So there is no way in which you can

internalize this ability by just thinking about that same action and self-generating (predicting) the consequences of it. But if you have learnt to associate the visual stimulus of the Coliseum and the word 'Coliseum' so that hearing that word will tend to re-activate the internal experience of seeing the Coliseum, then you just need to produce the word by yourself and listen to what you have produced and you can re-experience the Coliseum. Furthermore, since it is very easy to predict the acoustic effects of your phono-articulatory movements, the whole process can be easily internalized: next time you want to re-experience the visual image of the Coliseum, just think about producing its name, this will activate in you the experience of the word which in turn will activate the desired visual image.

5. Conclusion

A new framework is emerging in the contemporary science of behavior which considers cognition as "environmentally embedded, corporeally embodied, and neurally embrained." (van Gelder, 1999, pag. 244). Within this framework, a fundamental role is played by Cognitive Robotics, that is the use of robots as research tools for studying cognition. But up to now robotics research has been able to address only low level cognitive phenomena, and it is currently not clear how can these basic models be extended to explain higher level cognition. In this paper we argued that a new, very promising line of research for understanding increasingly high-level cognitive phenomena from the standpoint of contemporary research in embodied and situated cognition is to merge the basic ideas and synthetic methods of contemporary robotic research with the theoretical tradition in Russian psychology which considers language not only as a complex communication system but also a powerful cognitive tool. We have elaborated this idea by discussing a number of ways in which language can improve cognition. In particular, by reviewing both the empirical and computational literature and developing original ideas and hypotheses we have discussed the possible role that language might play in such cognitive phenomena as learning, categorization, abstraction, memory, voluntary control, and mental life.

An important aspect of our discussion that should be emphasized is that most of the advantages provided by talking to oneself do not seem to require a complex, syntactic language, but just the symbolic capacity to associate meanings, i.e., internal representations of significant experiences, with linguistic labels. This is important for at least two reasons. First of all, it suggests that the discovery of individual, cognitive uses of language could have happened quite early in language evolution, in particular, before the transition from an holistic proto-language to the full-blown compositional language of modern humans. This means that the evolution of language itself could have been favoured and partially driven by this individual, cognitive function of language (Mirolli and Parisi, 2005b, 2006; Parisi and Mirolli, 2007). Steels (2003) has provided some preliminary computational support to this idea: discussing one of his computational experiments in which agents had to evolve a compositional language, Steels reports that focussing one's attention to the self-produced linguistic utterances proved to be necessary for a population of agents playing a language game to develop a linguistic system with case grammar. But still more important for our purposes is that, since most of the advantages of talking to oneself seem to require only the symbolic property of language, i.e., the mapping between linguistic forms and meanings, studying the effects of language on cognition through robotic models is possible today, without the need for waiting for good, working robotic models of complex (grammatical) language acquisition. On the contrary, for the reasons just discussed, it is also possible that the development of self-talking robots might help us develop good models of grammatical language acquisition.

Of course, we are aware that most of our ideas are currently just speculations, even if sometimes supported by empirical findings and/or simple computational models. In fact, these ideas represent for us just a possible starting point for future Cognitive Robotics research. We hope we have demonstrated that studying the effects of language on cognition is one of the most important and promising way of trying to understand human cognition from the standpoint of embodied and situated cognition. Indeed, we are convinced that a fundamental step forward in our understanding

of the human mind will be made as theoretical (and empirical) psychology and robotics will unite in studying the role of language as a cognitive tool, that is, as current robotics will become a Vygotskian Cognitive Robotics.

Acknowledgements

The research presented in this paper has been supported by the ECAGENTS project funded by the Future and Emerging Technologies program (IST-FET) of the European Community under EU R&D contract IST-2003-1940.

References

- Baddeley, A. (1992), 'Working memory', *Science* 255, 556--559.
- Baddeley, A.; Chincotta, D. & Adlam, A. (2001), 'Working memory and the control of action: Evidence from task switching', *Journal of Experimental Psychology: General* 130, 641--657.
- Barsalou, L.W. (1999), 'Perceptual symbol systems', *Behavioral and Brain Sciences* 22, 577--609.
- Bechtel, W.; Abrahamsen, A. & Graham, G. (1998), 'The Life of Cognitive Science', in William Bechtel & George Graham, ed., *A companion to cognitive science*, Blackwell, Oxford, MA.
- Becker, J.T. & Morris, R.G. (1999), 'Working memory(s)', *Brain and Cognition* 41, 1--8.
- Beer, R.D. (2003), 'The dynamics of active categorical perception in an evolved model agent', *Adaptive Behavior* 11(4), 209--243.
- Berk, L.E. (1994), 'Why children talk to themselves', *Scientific American*, 78--83.
- Bickhard, M.H. (1980), *Cognition, Convention, and Communication*, Praeger Publishers, New York.
- Bickhard, M.H. (2001), 'Why Children Don't Have to Solve the Frame Problems: Cognitive Representations are not Encodings', *Developmental Review* 21, 224--262.
- Borghì, A.M.; Glenberg, A.M. & Kaschak, M.P. (2004), 'Putting words in perspective', *Memory & Cognition* 32(6), 863--873.
- Boysen, S.T.; Bernston, G.; Hannan, M. & Cacioppo, J. (1996), 'Quantity-based inference and symbolic representation in chimpanzees (Pan troglodytes)', *Journal of Experimental Psychology: Animal Behavior Processes* 22, 76--86.
- Brooks, R.A. (1991), 'Intelligence Without Representation', *Artificial Intelligence Journal* 47, 139--159.
- Cangelosi, A.; Greco, A. & Harnad, S. (2000), 'From robotic toil to symbolic theft: Grounding transfer from entry-level to higher-level categories', *Connection Science* 12(2), 143--162.
- Cangelosi, A. & Harnad, S. (2000), 'The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories', *Evolution of Communication* 4, 117--142.
- Carruthers, P. (2002), 'The cognitive functions of language', *Behavioral and Brain Sciences* 25, 657--726.
- Clark, A. (1997), *Being There: putting brain, body and world together again*, Oxford University Press, Oxford.
- Clark, A. (1998), 'Magic words: How language augments human computation', in Peter Carruthers & Jill Boucher, ed., *Language and thought: Interdisciplinary themes*, Cambridge University Press, Cambridge, pp. 162--183.
- Clark, A. (2001), *Mindware: an introduction to the philosophy of cognitive science*, Oxford University Press, Oxford.
- Clark, A. (2006), 'Language, embodiment, and the cognitive niche', *Trends in Cognitive Sciences* 10(8), 370--374.
- Clark, A. & Grush, R. (1999), 'Towards a Cognitive Robotics', *Adaptive Behavior* 7(1), 5--16.
- Clowes, R. & Morse, A.F. (2005), Scaffolding Cognition with Words, in L. Berthouze; F. Kaplan; H. Kozima; H. Yano; J. Konczak; G. Metta; J. Nadel; G. Sandini; G. Stojanov & C. Balkenius, ed., 'Proceedings

- Fifth International Workshop on Epigenetic Robotics: Modeling Cognitive Development in Robotic Systems', Lund University Cognitive Studies, Lund, pp. 101—105.
- Cole, M. (1996), *Cultural psychology: A once and future discipline*, Harvard University Press, Cambridge, MA.
- Dautenhahn, K. & Billard, A. (1999), Studying robot social cognition within a developmental psychology framework, in Schweitzer et al., ed., 'Proceedings of the Third European Workshop on Advanced Mobile Robots', IEEE Computer Society, Los Alamitos, CA, pp. 187--194.
- Dennett, D.C. (1991), *Consciousness Explained*, Little Brown & Co., New York, NY.
- Diaz, R. & Berk, L.E., ed. (1992), *Private speech: From social interaction to self regulation*, Erlbaum, New Jersey, NJ.
- 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, thinking and knowing', Kluwer, Dordrecht, pp. 115--141.
- Emerson, M.J. & Miyake, A. (2003), 'The role of inner speech in task switching: A dual-task investigation', *Journal of Memory and Language* 48, 148--168.
- Gallese, V. & Lakoff, G. (2005), 'The Brain's Concepts: The Role of the Sensory-Motor System in Reason and Language', *Cognitive Neuropsychology* 22, 455--479.
- Gentner, D. (2003), 'Why we are so smart', in Dendre Gentner & Susan Goldin-Meadow, ed., *Language in mind*, MIT Press, Cambridge, MA, pp. 195--235.
- Gibson, J.J. (1979), *The ecological approach to visual perception*, Houghton Mifflin, Boston.
- Gruber, O. (2002), The co-evolution of language and working memory capacity in the human brain, in M. Stamenov & V. Gallese, ed., 'Mirror neurons and the evolution of brain and language', Benjamins, Amsterdam, pp. 77--86.
- Gruber, O. & Goschke, T. (2004), 'Executive control emerging from dynamic interactions between brain systems mediating language, working memory and attentional processes', *Acta Psychologica* 115, 105--121.
- Harvey, I.; Di Paolo, E.a.; Quinn, M. & Wood, R. (2005), 'Evolutionary Robotics: A new scientific tool for studying cognition', *Artificial Life* 11(1--2), 79--98.
- Hutchins, E. (1995), *Cognition in the Wild*, MIT Press, Cambridge, MA.
- Kawato, M. (1999), 'Internal models for motor control and trajectory planning', *Current Opinion in Neurobiology* 9, 718--727.
- Kelso, J. (1995), *Dynamic Patterns*, MIT Press, Cambridge, MA.
- Kulakov, A. & Stojanov, G. (2002), Structures, inner values, hierarchies and stages: essentials for developmental robot architectures, in C.G. Prince; Y. Demiris; Y. Marom; H. Kozima & C. Balkenius, ed., 'Proceedings of the Second International Workshop on Epigenetic Robotics', Lund University Cognitive Studies, Lund, pp. 63--69.
- Lindblom, J. & Ziemke, T. (2003), 'Social Situatedness of Natural and Artificial Intelligence: Vygotsky and Beyond', *Adaptive Behavior* 11(2), 79--96.
- Loewenstein, J. & Gentner, D. (2005), 'Relational language and the development of relational mapping', *Cognitive Psychology* 50, 315--353.
- Lupyan, G. (2005), Carving nature at its joints and carving joints into nature: How labels augment category representations, in A. Cangelosi & R. Bugmann, G. & Borisjuk, ed., 'Modelling Language, Cognition and Action: Proceedings of the 9th Neural Computation and Psychology Workshop', World Scientific, Singapore, pp. 87--96.
- Luria, A.R. (1979), *The Making of Mind: A Personal Account of Soviet Psychology*, Harvard University Press, Cambridge, MA.
- Mirolli, M.; Cecconi, F. & Parisi, D. (2007), A Neural Network Model for Explaining the Asymmetries between Linguistic Production and Linguistic Comprehension, in S. Vosniadou & D. Kayser & A. Protopapas, ed., 'Proceedings of the European Cognitive Science Conference 2007', Lawrence Erlbaum,

Hove, pp. 670--675.

Mirolli, M. & Parisi, D. (2005a), Language as an aid to categorization: A neural network model of early language acquisition, in Angelo Cangelosi; Guido Bugmann & Roman Borisyuk, ed., 'Modelling language, cognition and action: Proceedings of the 9th Neural Computation and Psychology Workshop', World Scientific, Singapore, pp. 97--106.

Mirolli M. & Parisi D. (2005b), 'How can we explain the emergence of a language that benefits the hearer but not the speaker?' *Connection Science*, 17(3-4): 307--324

Mirolli, M. & Parisi, D. (2006), 'Talking to oneself as a selective pressure for the emergence of language', in A. Cangelosi, A.D.M. Smith & K. Smith eds., *The Evolution of Language. Proceedings of the 6th International Conference on the Evolution of Language*. Singapore, World Scientific: 214--221.

Miyake, A.; Emerson, M.J.; Padilla, F. & Ahn, J. (2004), 'Inner speech as a retrieval aid for task goals: The effects of cue type and articulatory suppression in the random task cuing paradigm', *Acta Psychologica* 115, 123--142.

Nazzi, T. & Gopnik, A. (2001), 'Linguistic and cognitive abilities in infancy: When does language become a tool for categorization?', *Cognition* 80, 303--312.

Nolfi, S. & Floreano, D. (2000) *Evolutionary Robotics*, MIT Press, Cambridge, MA.

Norman, D.A. (1980), 'Twelve issues for cognitive science', *Cognitive Science* 4, 1--32.

Oden, D.; Thompson, R. & Premack, D. (2001), Can an ape reason analogically? Comprehension and production of analogical problems by Sarah, a chimpanzee (*Pan troglodytes*), in D. Gentner; K.J. Holyoak & B.N. Kokinov, ed., 'The analogical mind: Perspectives from cognitive science', MIT Press, Cambridge, MA, pp. 471--498.

Parisi, D. (2007), 'Mental Robotics', in A. Chella & R. Manzotti eds., *Artificial Consciousness*, Imprinting Academic, Thorverton, 191--210.

Parisi, D.; Cecconi, F. & Nolfi, S. (1990), 'Econets: Neural networks that learn in an environment', *Network* 1, 149--168.

Parisi, D. & Mirolli, M. (2006), 'The emergence of language: How to simulate it', in C. Lyon; C. Nehaniv & A. Cangelosi, ed., *Emergence of Communication and Language*, Springer Verlag, Berlin.

Petrides, M.E.; Alivisatos, B.; Evans, A.C. & Meyer, E. (1993), 'Dissociation of human mid-dorsolateral from posterior dorsolateral frontal cortex in memory processing', *Proceedings of the National Academy of Sciences* 90, 873--877.

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

Piaget, J. (1968), *Genetic Epistemology*, Columbia University Press, New York.

Rattermann, M.J. & Gentner, D. (1998), The effect of language on similarity: The use of relational labels improves young children's performance in a mapping task, in K.J. Holyoak; D. Gentner & B.N. Kokinov, ed., 'Advances in analogy research: Integration of theory & data from the cognitive, computational, and neural sciences', New Bulgarian University, Sophia, pp. 274--282.

Rumelhart, D.; McClelland, J. & the PDP Research Group (1986), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 1 & 2, MIT Press, Cambridge, MA.

Schyns, P.G. (1991), 'A Modular Neural Network Model of Concept Acquisition', *Cognitive Science* 15(4), 461--508.

Sokolov, A.N. (1975), *Inner speech and thought*, Plenum Press, New York.

Spelke, E. (2003), 'What makes us smart? Core knowledge and natural language', in Dendre Gentner & Susan Goldin-Meadow, ed., *Language in mind*, MIT Press, Cambridge, MA, pp. 277--311.

Steels, L. (2003), 'Language-reentrance and the Inner Voice', *Journal of Consciousness Studies* 10(4-5), 173--185.

Steels, L. & Brooks, R., ed. (1994), *The artificial life route to artificial intelligence: Building Situated Embodied Agents*, Lawrence Erlbaum Ass., New Haven.

- Steels, L. & Belpaeme, T. (2005), 'Coordinating perceptually grounded categories through language: A case study for colour', *Behavioral and Brain Sciences* 28(4), 469--529.
- Thelen, E. & Smith, L.B. (1994), *A Dynamic Systems Approach to the Development of Cognition and Action*, MIT Press, Cambridge, MA.
- Thompson, R.K.R.; Oden, D.L. & Boysen, S.T. (1997), 'Language-naive chimpanzees (Pan troglodytes) judge relations between relations in a conceptual matching-to-sample task', *Journal of Experimental Psychology: Animal Behavior Processes* 23, 31--43.
- Tomasello, M. (2003), 'The key is social cognition', in Dendre Gentner & Susan Goldin-Meadow, ed., *Language in mind*, MIT Press, Cambridge, MA, pp. 47--57.
- Varela, F.; Thompson, E. & Rosch, E. (1991), *The Embodied Mind*, MIT Press, Cambridge, MA.
- van Gelder, T.J. (1999), Dynamic Approaches to Cognition, in R. Wilson & F. Keil, ed., 'The MIT Encyclopedia of Cognitive Sciences', MIT Press, Cambridge MA, pp. 243--245.
- Vygotsky, L.S. (1962), *Thought and language*, MIT Press, Cambridge, MA.
- Vygotsky, L.S. (1978), *Mind in society*, Harvard University Press, Cambridge, MA.
- Yoshida, H. & Smith, L.B. (2005), 'Linguistic cues enhance the learning of perceptual cues', *Psychological Science* 16(2), 90—95.
- Waxman, S.R. (2004), Everything had a name, and each name gave birth to a new thought: Links between early word-learning and conceptual organization, in D. Geoffrey Hall & Sandra R. Waxman, ed., 'From many strands: Weaving a lexicon', MIT Press, Cambridge, MA, pp. 295—335.
- Waxman, S. & Markow, D. (1995), 'Words as invitations to form Categories: Evidence from 12 to 13-month-old infants', *Cognitive Psychology* 29(3), 257--302.
- Wolpert, D. & Flanagan, J. (2001), 'Motor Prediction', *Current Biology* 11(18), 729--732.
- Zlatev, J. (2001), 'The epigenesis of meaning in human beings, and possibly in robots', *Minds and Machines* 11(2), 155—195.