

Language as a Cognitive Tool

Marco Mirolli · Domenico Parisi

Published online: 2 December 2009
© Springer Science+Business Media B.V. 2009

Abstract The standard view of classical cognitive science stated that cognition consists in the manipulation of language-like structures according to formal rules. Since cognition is ‘linguistic’ in itself, according to this view language is just a complex communication system and does not influence cognitive processes in any substantial way. This view has been criticized from several perspectives and a new framework (Embodied Cognition) has emerged that considers cognitive processes as non-symbolic and heavily dependent on the dynamical interactions between the cognitive system and its environment. But notwithstanding the successes of the embodied cognitive science in explaining low-level cognitive behaviors, it is still not clear whether and how it can scale up for explaining high-level cognition. In this paper we argue that this can be done by considering the role of language as a cognitive tool: i.e. how language transforms basic cognitive functions in the high-level functions that are characteristic of human cognition. In order to do that, we review some computational models that substantiate this view with respect to categorization and memory. Since these models are based on a very rudimentary form of non-syntactic ‘language’ we argue that the use of language as a cognitive tool might have been an early discovery in hominid evolution, and might have played a substantial role in the evolution of language itself.

Keywords Language · Cognitive tool · Computational model · Language evolution · Embodied cognition · High-level cognition

M. Mirolli (✉) · D. Parisi
Institute of Cognitive Sciences and Technologies, CNR, 44 Via San Martino della Battaglia,
00185 Rome, Italy
e-mail: marco.mirolli@istc.cnr.it

D. Parisi
e-mail: domenico.parisi@istc.cnr.it

Language and Cognition: The ‘Received’ View and its Critics

What is the role of language in human cognition? This is one of the most important questions we have to address if we want to understand the human mind. The standard view of classical cognitive science can be summarized with two statements: (a) cognition is, generally speaking, ‘linguistic’ in itself, in that it is the manipulation of language-like structures (propositions) according to formal rules; (b) the function of natural language is just to express these language-like structures; therefore, natural language does not affect cognition in any substantial way.

The view of cognition as symbol manipulation is at the very heart of classical cognitive science, constituting the common assumption of at least three of the sub-disciplines that gave birth to cognitive science: artificial intelligence (the symbol system hypothesis; Newell and Simon 1976), cognitive psychology (the language of thought hypothesis; Fodor 1975), and cognitive-science-related philosophy of mind (i.e., computationalism; Putnam 1963).

If one considers cognition as fundamentally linguistic, then there is no reason for viewing language as anything more than a very complex and powerful communication system. And, in fact, this view of language has been seldom if ever questioned inside traditional cognitive science.

The basic assumptions of classical cognitive science, however, have been questioned over the years from several perspectives. For example, a number of philosophical arguments have been put forward against the view of cognition as symbol manipulation (see, for example, Dreyfus 1972; Dennett 1978; Searle 1980; Churchland 1981). But in the absence of concrete alternative proposals advocates of the view of cognition as symbol manipulation could still claim that their hypothesis was “the only game in town” (Fodor 1975).

In the last 20 years a number of such alternatives have been proposed. The first one was connectionism: in their famous 1986 book, Rumelhart, McClelland, and the PDP group (Rumelhart et al. 1986) provided a concrete and detailed account of cognition which was completely alternative to the symbol manipulation paradigm. According to this alternative view cognition is not the manipulation of symbols according to formal rules, but rather the parallel and distributed processing of sub-symbolic information, that is, the transformation of purely quantitative values (the pattern of activation of groups of units) using other quantitative values (the connection weights linking groups of units) in networks of neuron-like units.

Other fundamental attacks to the classical view of cognition as symbol manipulation came in the early 1990s from behaviour-based robotics and artificial life (Brooks 1990; Parisi et al. 1990). The ‘Artificial Life route to Artificial Intelligence’ (Steels and Brooks 1994) pointed to the fact that 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 important determinants of the way a given task is solved. They are situated because the constraints provided by the environment can act also as opportunities for the task’s solution. And they are partially distributed because they do not happen only inside an organism’s head; rather, they crucially depend on the organism’s

environment which, especially in the human case, includes artefacts and other agents. (For a view of connectionism as part of artificial life, in which neural networks control the behaviour of embodied and situated agents see Parisi 2001).

Finally, another challenge to the symbolic approach to cognition came from dynamical systems theory. Proponents of the dynamical hypothesis argue that cognition should not be accounted for in computational terms, but rather using differential equations and dynamical systems concepts such as equilibrium points, cyclic behaviour, attractors, and bifurcations. More specifically, cognition must be understood by interpreting a cognitive system as a point moving in an abstract multi-dimensional space, and by identifying the trajectories that the system follows in that space and the laws that govern these trajectories (Smith and Thelen 1993; Port and van Gelder 1995; van Gelder 1998; Beer 2000).

The concepts and tools of connectionism, robotics, and dynamical systems theory opened up several very active areas of research, especially of the synthetic kind. The overall result is that contemporary cognitive science is substantially rethinking its view of cognition. In particular, the fundamental assumption of classical cognitive science that cognition is the manipulation of symbols according to formal rules is being replaced by a view according to which the mechanisms that explain behaviour are non-symbolic or sub-symbolic, and cognition consists in the adaptation of an agent to its environment. Furthermore, this adaptation critically depends on the dynamic interactions between the agent and the environment, which can also include artefacts and other agents (Bechtel et al. 1998; Clark 2001).

But apart from 'classical' connectionism, which addresses all levels of cognition but without taking into account 'embodiment' and 'situatedness', the new cognitive science has been so far concerned mostly, if not exclusively, with low-level behaviors and capacities, such as perception, learning, sensory-motor coordination, and navigation. The question remains open whether the same broad framework can scale up to explain the higher forms of cognition which characterize human beings (such as problem solving, reasoning, and planning), or if in order to explain characteristic human cognition we must go back to the symbol manipulation paradigm. From the point of view of the new cognitive science the most promising way of addressing this question, we argue, is to consider language not only as a communication system but also as a cognitive tool.

Language as a Cognitive Tool

The view of language as something that transforms all human cognitive processes dates back as early as the 1930s, with the work of Russian scholar Lev Vygotsky (Vygotsky 1962; Vygotsky 1978). According to Vygotsky, the most important moment in child development is that in which the child begins to use language not only as a social communication system but also as a tool for controlling her own actions and cognitive processes. When the child is challenged by a particularly difficult task she is often given help by an adult or a more skilled peer, and this help typically takes a linguistic form. 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

succeed in the problem. This is called ‘private speech’, which, according to Vygotsky, plays a fundamental role in the development of all human psychological processes.

The linguistic social aid coming from adults takes several forms. Social language helps a child to learn how to categorize experiences, to focus her attention on important aspects of the environment, to remember useful information, to inhibit non-useful behavior, to divide challenging problems into easier sub-problems and hence to construct a plan for solving complex tasks, and so on. When the child is talking to herself she is just making to herself what others used to do to her, that is, providing all sorts of cognitive aid 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, thus becoming inner speech. Hence, most, if not all, of adult human cognitive processes are linguistically mediated, in that they depend on the use of language for oneself.

Recently, the idea of language as a cognitive tool has been given increasing attention within the cognitive-science-oriented philosophy of mind (Carruthers and Boucher 1998). For example, Daniel Dennett (Dennett 1991, 1993, 1995) has argued that the human mind, including its most striking and hard to explain property, namely consciousness, depends mostly not on innate cognitive abilities, but on the way human plastic brains are substantially ‘re-programmed’ by cultural input coming, principally, through language: “Conscious human minds are more-or-less serial virtual machines implemented—inefficiently—on the parallel hardware that evolution has provided for us” (Dennett 1991, p. 278).

Andy Clark (Clark 1997, 1998, 2006) has further developed these Dennettian ideas by providing several arguments about how animal-like, embodied, situated, and sub-symbolic cognitive processes can be augmented by the learning and use of linguistic signs. According to Clark, language is not only a communication system, but also a kind of “external artifact whose current adaptive value is partially constituted by its role in re-shaping the kinds of computational space that our biological brains must negotiate in order to solve certain types of problems, or to carry out certain complex problems.” (Clark 1998, p. 163).

Apart of the interesting philosophical ideas of Dennett and Clark, the Vygotskian view of language as a cognitive tool has recently been raising increasing interest also in empirical cognitive science (see, for example, Gentner and Goldin-Meadow 2003). Indeed, a growing body of empirical evidence demonstrates the importance of language for a number of cognitive functions including learning (Waxman and Markov 1995; Nazzi and Gopnik 2001), memory (Gruber and Goschke 2004), analogy making (Gentner 2003), cross-modal information exchange (Spelke 2003), problem solving (Diaz and Berk 1992), abstract reasoning (Thompson et al. 1997), and logico-mathematical abilities (Dehaene et al. 1999).

In our work, we explore and articulate the hypothesis of language as a cognitive tool by the aid of artificial life simulations which use neural networks as models of the nervous system and genetic algorithms as models of evolution by natural selection. Computer simulations can provide fundamental tools in the development of new ideas and in the formulation of theories in that (a) they force the theory to be stated clearly and in full details, (b) they uncontroversially generate the

consequences of the assumptions of the theory as the simulation results, and (c) they suggest new ideas and directions of research. In what follows we describe some recent computational models of the use of language as a cognitive aid and of its role in human evolution.

Language and Categorization

Basically, organisms respond to sensory inputs by generating motor outputs. The motor output which is generated in response to some particular sensory input tends to have consequences that increase the individual's survival/reproductive chances. Evolution and learning are processes, respectively at the population and individual level, that result in acquiring the capacity to respond to sensory inputs with the appropriate motor outputs. We model organisms using neural networks, and evolution and learning as changes in the networks' connection weights that allow the organism to respond appropriately to sensory input.

If we look at sensory-motor mappings, we see that it is not the case that each different sensory input requires a different motor output. *Different* sensory inputs may require the *same* motor output, and different sensory inputs that require the same motor output are said to form 'categories'. (Motor outputs can be 'the same' at some more abstract level than the level of specific physical movements. An organism can respond to an object with the same action of 'reaching' for the object although in different occasions the specific physical movements of the organism's arm can be different, for example as a function of the arm's starting position.) What are categories in terms of a neural network model of behaviour? To answer this question we have to consider how a simple sensory-motor neural network is structured and functions.

In a neural network some particular sensory input is encoded as some particular activation pattern in the network's input units. This activation pattern elicits another particular activation pattern at the level of the hidden units, which in turn elicits a particular activation pattern in the output units. The activation pattern appearing in the output units determines the particular movement with which the organism responds to the sensory input. Neural networks learn to respond appropriately to sensory input by modifying their connection weights (either by genetic evolution or through individual learning) so that *different sensory inputs that must be responded to with the same motor output will elicit similar activation patterns in the hidden units*, and *similar sensory inputs that must be responded to with different motor outputs will elicit different activation patterns in the hidden units*. (For an artificial life model of this action-based view of categories, see Di Ferdinando and Parisi 2004.)

We can consider the activation pattern observed in the network's hidden units at any given time as one 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 (activation patterns that must be responded to with the same motor output are made more similar by the connection weights linking the input units to the hidden units) and (b) distant from each other (activation patterns that must be responded to with different motor outputs are made more different by these weights). The reason is that effectiveness of the organism's behaviour depends on the quality of these categories. With good categories the organism will be less likely to respond in different ways to sensory inputs that require the same response, or in the same way to sensory inputs that require different responses.

What are the consequences of the possession of language for an organism's categories? We can model language as a second sensory-motor network which is added to the basic sensory-motor network that we have already described and which underlies the organism's non-linguistic behaviour. We will call the two networks the 'sensory-motor network' and the 'linguistic network', respectively. Like the sensory-motor network, the linguistic network has a layer of sensory input units connected to a layer of hidden units connected to a layer of motor output units. The sensory units of the linguistic network encode linguistic (heard) sounds and the motor output units encode phono-articulatory movements that produce linguistic sounds. During the first year of life of the child, the linguistic and the sensory-motor network are not functionally (or perhaps even anatomically) connected and they are used separately. The child uses the sensory-motor network to learn to map non-linguistic sensory inputs from objects and persons into the appropriate motor actions (e.g., reaching for, grasping, and manipulating objects, following another person's gaze, turning towards another person, etc.) and uses the linguistic network to learn to generate phono-articulatory movements that result in sounds corresponding to heard sounds (that is, imitating the linguistic sounds of the particular language spoken in her environment).

At around 1 year of age proper language learning begins. The two networks become functionally connected and the child begins to learn the appropriate synaptic weights of the two-way connections linking the hidden units of the sensory-motor network to the hidden units of the linguistic network. What are the appropriate synaptic weights for these connections? These are weights such that a particular sound which is heard by the child, i.e., which is encoded in the sensory units of the linguistic network, will tend to elicit an activation pattern in the hidden units of the sensory-motor network which is similar to the activation pattern elicited by some perceived object or action, and thus in a non-linguistic action which is appropriate to the heard sound. This is language understanding. And, conversely, a particular perceived object or action, which is encoded in the sensory units of the sensory-motor network, will tend to elicit an activation pattern in the hidden units of the linguistic network that result in some appropriate phono-articulatory movements. This is language production.

What are the consequences of this reciprocal functional linking of the sensory-motor network and the linguistic network, i.e., of possessing a language, for the organism's categories? The answer is that categorization is enhanced by language (Mirolli and Parisi 2005b). When the child hears and understands the language spoken by others, the child's categories tend to become better categories, i.e.,

smaller and more distant clouds of points in the child's neural network. If the child perceives an object and at the same time she hears the linguistic sound that designates the object in the particular language spoken in her environment, the activation pattern that results in the hidden units of the child's sensory-motor network depends on both the sensory input from the object and the sensory input from the linguistic sound. The consequence is that this activation pattern is more similar to the activation patterns elicited in other occasions by other objects belonging to the same category (that must be responded to with the same action) and more dissimilar to the activation patterns elicited by objects belonging to other categories, compared with the activation pattern observed in an organism without language.

But this is not the whole story. An important characteristic of human language, which distinguishes it from the communication systems of other animals, is that human language is used not only for communicating with others but also for communicating with oneself. Indeed, the use of language for oneself starts as soon as language is acquired, and represents a significant proportion of the child's linguistic production. Empirical studies demonstrate that 3–10 year old children use language for themselves 20–60% of the time (Berk 1994).

As discussed above, the use of language for talking to oneself can be related to the 'language as a cognitive tool' hypothesis: private speech happens as the child discovers that she can exploit the advantages provided by language by talking to herself. Later on, the child can internalize this linguistic self-aid, by just 'thinking' linguistic labels without producing them out aloud. Can this interpretation of private and inner speech be applied to the advantages produced by language on categorization? In order to answer this question we need to model both ways in which humans can talk to themselves: externally, as private speech, or internally, as inner speech.

The simulation of private speech is quite straightforward. The network encounters an object and it responds to the object by producing the sound that designates the object using its linguistic sub-network. Then, the network hears the sound it has just produced and responds, using its sensory-motor sub-network, to the internal representation of the self-produced sound. Inner speech can instead be simulated as follows. When the network perceives an object, it does not produce any sound. Nonetheless, the sight of the object does induce the internal representation of the name of the object in the linguistic hidden units. In inner speech, it is this internal representation of the label associated to the perceived object that influences the non-linguistic response of the network.

As it turns out, the advantage for the network's categories provided by social language, when the network hears linguistic signals produced by somebody else, can be observed even if the network is all alone and talks to itself. In fact, both self-produced and internally-thought linguistic signals improve sensory-motor internal representations of perceived objects more or less to the same extent as social linguistic input. That is, compared to the representations of the pre-linguistic network, internal representations of objects belonging to the same category are more similar (close) to each other, and those of objects belonging to different categories are more different (distant) to each other (see Mirolli and Parisi 2006).

Talking to Oneself in the Evolution of Language

Why did language evolve? What was the adaptive function of language? This question is surely of the most importance, if one wants to understand the evolution of language and of man in general. Nonetheless, in the contemporary literature on language evolution there is not much debate on this topic (see, for example, Knight et al. 2000; Christiansen and Kirby 2003). One reason seems to be the common assumption that the only function of language is communication. As we have discussed above, the 'received view' holds that language is nothing but a very complex and powerful communication system. But once one has acknowledged the importance of language in the development of human cognition one can no more assume that the evolution of language has been driven only by the pressures for better communication. On the contrary, an interesting question immediately rises: when did hominids started to use language for themselves as a cognitive tool?

Generally, there is a tendency to think that language was used by humans to communicate with themselves only when language was already well developed and was sophisticated and syntactically complex; hence, quite recently compared with the first appearance of a proto-language. However, this is not necessarily the case. Even a very simple proto-language, for example, a language made up of single words (or holophrases), may be used to talk to oneself, with advantages for the individual that uses the language in this way. Based on this hypothesis, we have developed another set of simulations in which we studied the effect of talking-to-onself for the evolution of a simple communication system (Mirolli and Parisi 2005a).

In this simulation a population of artificial organisms (whose behaviour is controlled by neural networks) evolve in a simple world which contains both other organisms and poisonous and edible mushrooms. Organisms must avoid poisonous mushrooms, which decrease an individual's probability to reproduce, and eat edible ones, which increase individual fitness. Furthermore, organisms can communicate to each other the quality of encountered mushrooms by emitting signals through their linguistic output units. But in order to exploit the advantages provided by communication the population must evolve an appropriate communication system. Since each individual mushroom is different from all other mushrooms belonging to the same category, organisms must evolve the capacity to send similar signals every time they encounter an edible mushroom and another class of signals when they encounter poisonous mushrooms.

The evolution of such a communication system proves to be quite difficult, especially because in this simulation there is no direct selective pressure for producing appropriate signals: an individual's reproductive chances depend on the number and quality of mushrooms the individual eats, not on the signals it produces. Indeed, by producing good signals an individual can increase the probability of reproduction of other individuals, thus providing a direct advantage to competitors. The result is that in the standard simulation, in which signals are used only for social communication, a good and stable communication system never evolves.

In another simulation we let organisms use signals not only for social communication, but also for talking to themselves, as aids to memory. In particular,

organisms can hear their self-produced signals and use them in order to remember the information received by other organisms. The results of this second simulation are clear: if organisms can use language not only as a social communication system but also as a cognitive (memory) aid the evolution of language itself is favoured, and this has a positive impact on the organisms' fitness as well. In other words, organisms which can talk to themselves develop a better communication system and reach a higher fitness with respect to organisms which can use signals only for communication. The reason is that in order to exploit the advantages provided by using language as a memory aid organisms must produce useful signals, because otherwise they would mislead themselves. Hence, talking to oneself associates a direct individual advantage to producing useful signals, which was not the case in the previous simulation.

Using language as an aid to memory can be advantageous for at least two reasons: (a) delegating the memory function to the linguistic system can leave the sensory-motor system free to process other information useful for acting in the environment while linguistically remembering previous information, and (b) linguistic signals may occupy less space in memory than the sensory-motor information they refer to.

Using language as a cognitive tool may have had a fundamental impact not only on categorization and memory. For example, other neural network simulations have shown that language can improve the *learning* of categories (Schyns 1991; Lupyan 2005). Furthermore, the artificial life simulations of Cangelosi et al. (Cangelosi and Harnad 2000; Cangelosi et al. 2000) have demonstrated that language can also allow 'symbolic theft', that is, a way of learning useful categories not by direct sensory-motor experience with the world but through cultural transmission mediated by language. And it can be argued that talking to oneself can be useful in many additional ways. It can allow an individual to direct her attention to specific aspects of the environment, to make explicit predictions of future states of the environment, and to explicitly plan future actions (see Parisi and Mirolli 2006).

In as much as the advantages of talking to oneself do not require a complex syntactic language, it is reasonable that the discovery of the 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. And this is just what the computational models reviewed here suggest: none of them included any kind of syntax, but just the 'symbolic' capacity to associate 'meanings' (as internal representations of significant experiences) with linguistic labels. Nonetheless, they demonstrated that addressing to oneself even simple linguistic labels can provide important individual advantages. Trying to sort out what could have been the consequences of this early use of language for oneself in the subsequent evolution of language is an interesting topic for future research.

Conclusion

A crucial, but often neglected, characteristic of human language is that language is used not only for communicating with others but also for communicating with

oneself, whereas we seem not to have evidence for this type of use of animal communication systems. Talking to oneself, in the form of both private and inner speech, has tremendous consequences for the development of the human mind. Indeed, we have argued that considering the cognitive role of language can provide the missing link for addressing the high-level cognitive capacities which characterize humans within the new, emerging framework which considers cognition as “environmentally embedded, corporeally embodied, and neurally embrained.” (van Gelder 1999, p. 244). In the present paper we have described some simple computer simulations that show that language can improve one’s categories and can be an useful aid to memory, both if it mediates social communication and if it is used to talk to oneself as private or inner speech. But we argue that the use of language for oneself does not improve only categorization and memory, but almost any human cognitive function (for a discussion of other possible roles of language in enhancing cognitive processes, see Mirolli and Parisi 2009). Therefore, much more work needs to be done in order to understand the relationships between language and cognition. And we think that computer simulations will play an important role in our understanding of this fundamental topic.

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

References

- Bechtel, W., Abrahamsen, A., & Graham, G. (1998). The life of cognitive science. In B. William & G. George (Eds.), *A companion to cognitive science*. Oxford: Blackwell.
- Beer, R. D. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3), 91–99.
- Berk, L. E. (1994). ‘Why children talk to themselves’, *Scientific American*, (pp. 78–83).
- Brooks, R. A. (1990). Elephants don’t play chess. *Robotics and Autonomous Systems*, 6, 3–15.
- 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., & Boucher, J. (Eds.). (1998). *Language and thought: Interdisciplinary themes*. Cambridge: Cambridge University Press.
- Christiansen, M. H., & Kirby, S. (Eds.). (2003). *Language Evolution: The State of the Art*. Oxford: Oxford University Press.
- Churchland, P. M. (1981). Eliminative materialism and the propositional attitudes. *Journal of Philosophy*, 78(2), 67–90.
- Clark, A. (1997). *Being there: putting brain, body and world together again*. Oxford: Oxford University Press.
- Clark, A. (1998). Magic words: How language augments human computation. In P. Carruthers & J. Boucher (Eds.), *Language and thought: Interdisciplinary themes* (pp. 162–183). Cambridge: Cambridge University Press.
- Clark, A. (2001). *Mindware: An introduction to the philosophy of cognitive science*. Oxford: Oxford University Press.
- Clark, A. (2006). Language, embodiment, and the cognitive niche. *Trends in Cognitive Sciences*, 10(8), 370–374.
- Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., & Tsivkin, S. (1999). Sources of mathematical thinking: Behavioral and brain-imaging evidence. *Science*, 284, 970–974.

- Dennett, D.C. (1978). *Brainstorms*, Bradford Books, Montgomery, Vt.
- Dennett, D. C. (1991). *Consciousness explained*, Little Brown & Co., New York, NY.
- Dennett, D. C. (1993). Learning and labeling. *Mind and Language*, 8(4), 540–547.
- Dennett, D. C. (1995). *Darwin's dangerous idea: evolution and the meanings of life*, Simon and Schuster, New York, NY.
- 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* (pp. 115–141). Dordrecht: Kluwer.
- Diaz, R., & Berk, L. E. (Eds.). (1992). *Private speech: From social interaction to self regulation*, Erlbaum, NJ: New Jersey.
- Dreyfus, H. (1972). *What computers can't do: A critique of artificial reason*. New York: Harper & Row.
- Fodor, J. (1975). *The language of thought*. Cambridge: Harvard University Press.
- Gentner, D. (2003). Why we are so smart. In D. Gentner & S. Goldin-Meadow (Eds.) *Language in mind* (pp. 195–235). Cambridge: MIT Press.
- Gentner, D., & Goldin-Meadow, S. (Eds.). (2003). *Language in mind*. Cambridge: MIT Press.
- 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.
- Knight, C., Hurford, J. R., & Studdert-Kennedy, M. (Eds.). (2000). *The evolutionary emergence of language: Social Function and the Origins of Linguistic Form*. Cambridge: Cambridge University Press.
- Lupyan, G. (2005). Carving nature at its joints and carving joints into nature: How labels augment category representations. In A. Cangelosi, G. Bugmann, & R. Borisyuk (Eds.) *Modelling language, cognition and action: Proceedings of the 9th neural computation and psychology workshop* (pp. 87–96). Singapore: World Scientific.
- Mirolli, M., & Parisi, D. (2005a). How can we explain the emergence of a language which benefits the hearer but not the speaker? *Connection Science*, 17(3–4), 325–341.
- Mirolli, M., & Parisi, D. (2005b). Language as an aid to categorization: A neural network model of early language acquisition. In A. Cangelosi, G. Bugmann, & R. Borisyuk (Eds.) *Modelling language, cognition and action: Proceedings of the 9th neural computation and psychology workshop* (pp. 97–106). Singapore: World Scientific.
- Mirolli, M., & Parisi, D. (2006). Talking to oneself as a selective pressure for the emergence of language. In A. Cangelosi, A. Smith & K. Smith (Eds.) *The evolution of language: Proceedings of the 6th international conference on the evolution of language* (pp. 214–221). World Scientific Publishing.
- Mirolli, M., & Parisi, D. (2009). Towards a Vygotskian cognitive robotics: The role of language as a cognitive tool. *New Ideas in Psychology*. doi:10.1016/j.newideapsych.2009.07.001.
- Nazzi, T., & Gopnik, A. (2001). Linguistic and cognitive abilities in infancy: When does language become a tool for categorization? *Cognition*, 80, 303–312.
- Newell, A., & Simon, H. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19, 113–126.
- Parisi, D. (2001). Neural networks and Artificial Life. In David, Baltimore, R. Dulbecco, F. Jacob, & R. Levi-Montalcini (Eds.), *Frontiers of life* (Vol. 4). San Diego: Academic Press.
- 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 (Eds.) *Emergence of communication and language*. Berlin: Springer.
- Port, R. F., & van Gelder, T. (Eds.). (1995). *Mind as motion*. Cambridge: MIT Press.
- Putnam, H. (1963). Brains and behavior. In R. Butler (Ed.) *Analytical philosophy*. Oxford: Basil Blackwell.
- Rumelhart, D. E., & McClelland, J. L., and the PDP research group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*, Cambridge: MIT Press.
- Schyns, P. G. (1991). A modular neural network model of concept acquisition. *Cognitive Science*, 15(4), 461–508.
- Searle, J. R. (1980). Minds, brains and programs. *Behavioral and Brain Sciences*, 3, 417–458.
- Smith, L. B., & Thelen, E. (Eds.). (1993). *A dynamic systems approach to development*. Cambridge: MIT Press.
- Spelke, E. (2003). What makes us smart? Core knowledge and natural language. In D. Gentner & S. Goldin-Meadow (Eds.) *Language in mind* (pp. 277–311). Cambridge: MIT Press.

- Steels, L., & Brooks, R. (Eds.). (1994). *The artificial life route to artificial intelligence: Building situated embodied agents*. New Haven: Lawrence Erlbaum Ass.
- 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.
- van Gelder, T. J. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, *21*, 1–14.
- van Gelder, T. J. (1999). Dynamic approaches to cognition. In R. Wilson & F. Keil (Eds.) *The MIT Encyclopedia of cognitive sciences* (pp. 244–246). Cambridge: MIT Press
- Vygotsky, L. S. (1962). *Thought and language*. Cambridge: MIT Press.
- Vygotsky, L. S. (1978). *Mind in society*. Cambridge: Harvard University Press.
- 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.