

Steps Towards Artificial Consciousness: A Robot's Knowledge of Its Own Body

Domenico Parisi **Marco Mirolli**

Institute of Cognitive Sciences and Technologies, National Research Council

44 Via S. Martino della Battaglia, 00185 Rome, Italy

domenico.parusi@istc.cnr.it marco.mirolli@istc.cnr.it

Abstract

Consciousness is a very difficult phenomenon to define and analyze. What we should do is try to operationalize it by reproducing in an artificial system (robot) single consciousness-related phenomena such that it is easy to say whether or not the artificial system exhibits the phenomenon. One such phenomenon is knowledge of one's own body as something which is different from other physical objects. We discuss robots that can reach with their hand specific unseen parts of their body, the role of tactile input and sensory multimodality in knowledge of one's own body, and how this knowledge can be demonstrated by robots that show "surprise" when their predictions are falsified. We also discuss how to reproduce the distinction between private and public knowledge and how more explicit and "conscious" knowledge of one's body can be demonstrated by robots that can use sentences containing the words "I" and "you".

Operationalizing Consciousness

In his seminal paper "Computing machinery and intelligence" Alan Turing proposed an operational definition of the term 'intelligence' that would provide a possibly shared ground for framing any meaningful future debate about the possibility of Machine Intelligence (Turing 1950). Turing's operational definition resulted in the well-know 'Turing test' according to which a machine can be said to be intelligent if it can engage in a computer-mediated conversation with a human without being recognized as a machine, rather than another human being. The importance of the Turing's proposal lies not so much in its success in framing the debate on machine intelligence on undisputable grounds. In fact, Turing proposal has been criticized on several grounds. And notwithstanding the fact that an annual contest (the Loebner prize) is organized with a 100.000 \$ prize for the first machine able to pass the Turing test, nowadays the most important approaches to reproducing intelligence in machines are no more focusing on the symbolic, conversational ability proposed by Turing, but rather on effective (adaptive) interactions between agents and their environment, which are today considered to be the hallmark of (any kind of) intelligence. Rather, the importance of Turing's paper lies in the fact that it set the stage for the development of Artificial Intelligence: if intelligence can be defined in behavioral terms, then it might also be possible to reproduce

intelligence by constructing machines demonstrating the same kinds of behaviors of natural intelligent beings. Furthermore, by providing an operational definition, Turing addressed a fundamental and still valid requirement advanced by the epistemologists and psychologists of his time (neopositivists and behaviorists, respectively): scientific concepts must be related to observable phenomena such as behavior. Otherwise, we are not dealing with science but with metaphysics. And by pointing to a way of speaking scientifically about mind-related concepts Turing set the stage for the birth not only of Artificial Intelligence but also of modern Cognitive Science.

We think the time has come for trying to apply Turing's operational approach not only to Intelligence but also to Consciousness (Holland, 2003; Chella and Manzotti, 2007). By providing operational definitions of consciousness and trying to build artifacts which, according to these definitions, can be reasonably said to be conscious, we can try and frame the debate on consciousness on less metaphysical and more scientific grounds than it is typically the case in traditional philosophical debates. If we take this approach, our question changes from "What does it mean to be conscious?" in "What are the conditions under which a machine can be reasonably claimed to be 'conscious'?"

As soon as one takes this perspective, it becomes quite clear that 'consciousness' is a rather tricky concepts, being both vague and ambiguous. In fact, under the same word we refer to a number of quite different phenomena: for example, to be conscious can mean to be attentive, to be aware of something, to be awake rather than sleeping, to be self-reflecting, to be able to report about something, to have a sense of being a self, and so on. Rather than trying to clarify this mess, philosophers have typically looked for a unique, core meaning of consciousness. And this has been typically identified with 'having a phenomenal experience', 'having *qualia*', 'feeling how it is like to be someone'. This has been dubbed as the 'Hard Problem' of consciousness (Chalmers 1996), and clearly separated from all the other problems, which have been regarded as 'Soft' and consequently largely ignored. But this has just pushed the debate on consciousness towards metaphysics, since the Hard problem is considered hard just because phenomenal experience can't be, by definition, operationalized. Concepts such as zombies, qualia, and the

like which have been dominating the debate on consciousness are all based on the same idea: that phenomenal consciousness can't be related to observable behavior. Not surprisingly, being based not on scientific grounds but on intuitions, the debate has not reached a consensus: different authors (and readers) have different intuitions, with Hard-problem enthusiasts stressing the mysterious nature of consciousness and questioning any possible solution, and critics denying the very existence of the problem (Dennett, 1991; Blackmore, 2003).

We think that a more fruitful way to address the problem(s) of consciousness is to try and clarify the concept itself by disentangling the several meanings that this words has. In order to do that, we take the synthetic stance, and try to look for operational definitions of the several 'Soft' meanings of consciousness, so that we might try and reproduce in an artificial system specific phenomena that can be considered as aspects or components of consciousness and whose implementation is at least *thinkable*, if not already feasible (Parisi, 2007). Some of these aspects are related to the possession of a 'sense of self' which appears to be a crucial component of consciousness or self-consciousness. In this paper we try and clarify two such aspects by relating them to an agent's knowledge of its own body: (1) the knowledge of a distinction between the self and its (external) environment, and (2) the knowledge of possessing both *private* and *public* experiences. In very general terms, knowledge in real organisms, and in artifacts that aim at reproducing real organisms, is what makes it possible for the organism/artifact to behave in adaptive/useful ways. So our questions become: "What behaviors should be exhibited by an artifact that (1) can be said to possess knowledge of its own body as distinct from other objects, and (2) can recognize the private character of at least some of the sensory inputs originating in its own body? In what follows we will propose some reflections on how to provide answers to these questions.

The Distinction Self vs. Environment

In line with recent advances in Artificial Intelligence and Cognitive Science, we presupposes a physically embodied AI which is different from traditional AI (Pfeifer and Scheier, 2001). Traditional AI is not physically embodied since its emphasis is on intelligence as pure manipulation of symbols. Physically embodied AI is robotics, i.e., the construction of artifacts whose intelligence depends critically on their possessing a physical body which interacts autonomously with a physical environment. A robot's behavior can be controlled by a simplified model of the nervous system, i.e., by a neural network. The neural network has input units that receive information from both the external environment and the robot's own body, output units that cause movements of (parts of) the robot's body or changes within the robot's body, and a more or less complex architecture of internal units that map input into output. So how can a robot know (and demonstrate of

knowing) that some of its inputs come from its own body while others come from the external environment?

Reaching Specific Parts of one's own Body

Consider a typical robot with a visual field and a single multi-segment arm. What appears in the robot's visual field is encoded in the neural network's input units while the output units encode changes in the position of the different segments of the robot's arm. When an object appears in the robot's visual field, the robot responds by moving its arm so that the arm's endpoint (the hand) reaches the object. Now imagine that some part of the robot's body enters the robot's visual field. The robot moves its arm so that its hand reaches that part of the body. However, there are parts of the robot's body that cannot enter the robot's visual field, for example the back of the body. Furthermore, since the robot by moving its eyes can see some parts of the external environment and fail to see other parts, in some circumstances the robot can fail to see even parts of its own body that in principle could enter the robot's visual field.

Real organisms can reach with their hand (or other parts of their body such as their snout or muzzle) specific parts of their own body even if these parts are not seen. This is possible if their body generates another type of sensory input which is different from visual input. For example, some specific parts of the body can produce pain stimuli or itches. If these stimuli are localized in specific parts of the body, the organism can respond by moving its arm so that the hand reaches the specific part of the body that generates the pain or itch stimuli and can either massage or scratch that part.

An important problem to be solved is how to encode in the input units of the robot's neural network the information originating in the robot's body, say, a localized pain or itch. Whatever the solution adopted for this problem, it is important to recognize that spatial knowledge is action-based, not stimulus-based. I know where things are in space, or what is the distance between two things in space, because I know how to reach things with my hand or my eyes or, in the case of the distance between two things, because I know how long I have to move my arm or my eyes to go from one thing to the other. For example, I know that an object is on my left and not on my right because, given the same starting position of my arm (or eyes), in order to reach the object I have to move my arm (or eyes) in certain ways which are systematically different from how I should move my arm (or eyes) if the object were on my right. Whatever the manner in which we encode information about the body in the robot's sensory units, it should therefore be possible to train a robot to reach with its hand particular unseen portions of its body where pain or itch is felt in the same way as robots are routinely trained to reach with their hand visually perceived objects.

Tactile Input

Another aspect of knowledge of one's own body arises from the existence of another kind of sensory input originating in one's body: tactile input. When one part of one's body comes into physical contact with an object, a tactile input is generated in that part of the body and sent to the central nervous system. It has been shown (Schlesinger and Parisi, 2001) that adding a tactile input to the neural network of a robot that has to reach a visually perceived object with its hand favors the acquisition of the ability to reach for the object in that the robot's neural network is directly informed by its touch sensors when its hand comes into physical contact with the object. The sense of touch can play other important roles in the emergence of self-consciousness if it becomes an haptic sense, that is, if by moving its hand on the surface of an object an organism can recognize the object's physical shape. (For a robot that distinguishes spherical vs. cube-shaped objects by exploring the surface of the objects with its hand, see Nolli and Marocco, 2002; for a very clear treatment of the importance of haptic perception for consciousness, see Morasso, 2007.)

The sense of touch is at the origin of an important component of an organism's knowledge of its own body as something different from other physical objects present in the environment. As we have said, when one part of an organism's body comes into physical contact with an object, a tactile sensory input is generated in that part of the organism's body. This already distinguishes the organism's own body from other objects, since a tactile input is generated only when one part of the organism's body comes into visually perceived physical contact with an object but not when the organism sees two physical objects which enter in physical contact with one another. In addition, when one part of the organism's body comes into physical contact with another part of the organism's body, a tactile input is generated in both parts of the organism's body, which of course is not true when one part of the organism's body comes into physical contact with an object. Hence, the sense of touch appears to contribute importantly to an organism's knowledge of its own body as different from other objects.

Sensory multimodality

Sensory multimodality is another source of knowledge that differentiates one's body from other physical objects existing in the environment. An important property of the environment, both external and internal, with which an organism's nervous system interacts is that one and the same entity existing in the environment can be a source of a variety of sensory inputs belonging to different sensory modalities. However, from this point of view too the external environment and the organism's body are different. Objects existing in the environment can be a source of visual input and in some cases of acoustic input and smell input, while tactile and taste inputs can be generated only when some portion of the organism's body

comes into physical contact with the object. In contrast, in addition to these inputs one's body generates proprioceptive input, vestibular input, and a variety of other inputs collectively called somatosensory (pain, itches, sense of fatigue, etc.) What is even more important is that there is systematic co-variation among the different sensory inputs and that an organism's nervous system can recognize and exploit this co-variation. When my hand comes into physical contact with an object I can both see that my hand and the object are in physical contact and get tactile input from my hand. When I move my arm I can both see my arm and hand moving and feel the corresponding changes in proprioceptive input from my arm.

Neural networks can incorporate multimodal co-variation in their knowledge of the environment if, for example, different sensory modalities have each its own separate set of sensory input units that map into separate sets of hidden units but the different sets of hidden units have interconnections that link each set of hidden units with the other sets. We can consider the activation patterns in the different sets of hidden units as composing a single, composite pattern. By using Hebbian learning to train the interconnections, we can obtain a complete overall pattern even when one of the sensory inputs is missing. For example, when an organism, or a robot, feels that its hand is touching an object, the organism can "imagine" to see that its hand and the object are in physical contact even if there is no actual visual input from the hand and the object.

The reconstruction of missing parts of sensory input gives different results when the source of the existing or missing sensory inputs is the organism's body and when only the external environment is such a source. This therefore appears to be another basis for the distinction between one's body and the rest of the environment. Furthermore, different parts of the organism's body, e.g., arms, legs, head, tend to be associated with different types of patterns of co-variation between multiple sensory input, and this has been replicated with robots that demonstrate with their behavior that they can distinguish between different parts of their body (cf. Hafner and Kaplan, 2005). This might be at the origin of the fact that organisms do not only distinguish between their own body and other objects in the environment but possess an articulated knowledge of their body as made up of different parts in specific spatial relations with one another.

Controlling one's own Body

So far, we have been discussing how a robot might learn the distinction between inputs which come from the environment and those which come from its own body by just relying on sensory information. But one fundamental distinction between an agent's body and the rest of its environment lies in the fact that the agent's body is *directly under the agent's control*, while objects in the environment are not. As discussed above, an agent's control system is composed not only of sensory units but also of output units, whose patterns of activation determine

the movements of the agent's effectors. Each time the agent's output units send commands to the effectors, this results in movements of the parts of the agent's body which are controlled by those units, which in turn will result in some change in the agent's sensory state related to the part of the body which has been moved: in particular, each output command will result in *systematic* changes in the proprioceptive input, and, in some circumstances, also in the visual and somatosensory inputs. What is important here is that these changes are systematic, in the sense that there is a high correlation between changes in the agent's output units and changes in the relevant sensory units. For example, each time the motor neurons controlling the contraction of my arm muscles fire this always result in corresponding changes in the proprioceptive (and sometimes visual) input from my arm. This is not true for the inputs which come from the environment: the relationship between an agent's output commands and the environmental input is much more indirect and less systematic. For example, the changes in visual input which result from changes in the activation of the output units which control the movements of my eyes are highly unpredictable, depending on the current details of the environment, which are different from time to time. Hence, another fundamental source of information which can be used by a robot for distinguishing its own body with respect to the rest of the environment comes from considering which part of the input is under the agent's direct control and which is not.

The Distinction Private vs. Public

Why are some sensory inputs private and other sensory inputs public? How can a robot be aware that some sensory inputs are private and other sensory inputs public?

The answer to the first question appears to be rather simple. For purely physical (physiological) reasons, certain events that take place within an organism's body (or brain) cause sensory effects only in the particular organism's nervous system, and cannot cause similar effects in the nervous systems of other organisms. These sensory effects are therefore necessarily private and lead to the emergence of the organism's private world. On the other hand, events in the external environment can cause sensory effects in the nervous systems of all the organisms that happen to be sufficiently close to the physical origin of the physical (or chemical) cause-effect chain leading to the sensory effect. These sensory effects are public and cause the emergence of the organism's public world, a world shared with other organisms (Parisi, 2004).

When we turn to our second question, "How can a robot be aware that some sensory inputs are private and other inputs public", the first thing to note is that this a question that only arises for social organisms. Organisms that tend to live all alone have no basis for distinguishing a private world from a public world. More correctly, for these organisms the distinction makes no sense at all. The concept of a public world acquires meaning only in a social

contest, in which the same experience is accessible to more than one individual. Consequently, the very first requirement for a robot to acquire an awareness that its world includes a private part to which it only has access and a public part which it shares with other robots (or with humans) is that it must be a highly social robot, with frequent sensory access to other robots (or humans) and with frequent interactions with them. These social interactions can be both cooperative and competitive, and both kinds of interactions might have an important role in the genesis of the private vs. public distinction. It is in fact probably thanks to the constant engagement in cooperative tasks which require *joint attention* that a sense of *public experience* might arise. But it is probably thanks to the constant engagement in competitive interactions in which the usefulness of dissimulating one's own internal state (like being hungry or tired) might be discovered that the awareness of possessing *private knowledge* may arise.

Assessing Robot's Self-Knowledge

'Surprise' Behaviors Resulting from Falsified Predictions

How can self-knowledge be assessed in robots? As we have said, an organism's knowledge can be defined in very general terms as anything in an organism's structure that makes it possible for the organism to behave appropriately. But both real organisms and robots can possess various ways of knowing that their body is something special with respect to the rest of the environment. This knowledge can have several levels and grades of explicitness, from being entirely implicit to being entirely explicit when organisms such as humans express it through language. (The distinction between a more implicit and a more explicit knowledge of one's own body is sometimes captured in the literature by distinguishing between "body schema" and "body image". Cf. De Preester and Knockaert, 2005.)

We have already indicated some implicit ways in which a robot can distinguish between its own body and the rest of the world. For example, a robot can reach with its hand some particular unseen part of its body when that part of the body generates some sensory input (pain or itch) that only one's body can generate. (For a more dynamical knowledge of one's own body in a robot, cf. Tani, 1998.)

A somewhat more explicit (and interesting) form of knowledge of the distinction between one's body and the external environment can be demonstrated by robots that possess prediction abilities, that is, robots whose neural network is able to internally generate sensory input that matches the sensory input that will actually arrive to the neural network at some future time. With robots that possess prediction abilities their knowledge can be more explicitly and specifically demonstrated by exposing the robots to conditions in which their predictions fail. Prediction abilities are of two types. The first type is predicting an event that follows another event independently of the actions of the organism. An example

is predicting the weather. The second type of predictions are predictions of events that are caused by the actions of the organism. My knowledge of the room in which I am now can be demonstrated by my ability to predict that if I turn to the right I will see the window, whereas if I turn to the left I will see the door. Another example is predicting that the glass will fall down and break if I open my hand that holds the glass.

The ability to predict can be simulated in robots by adding a set of predicting units to the neural network that controls the robot's behavior. The predicting units receive connections from both the sensory input units and the motor output units but a prediction is generated before the activation pattern in the motor output units with which the neural will respond to the current sensory input is actually translated into a physical movement of some specific part of the robot's body. The robot can be said to be able to predict if the activation pattern that appears in the predicting units matches the next sensory input, i.e., the sensory input that will arrive to the neural network's sensory input units after the movement specified by the motor output units has been physically executed (Jordan and Rumelhart, 1992).

In organisms that are able to predict their knowledge of the environment may be expressed in their predictions. If we create special conditions in which an organism's predictions are falsified, that is, if we allow an action on the part of the robot to be followed by a sensory consequence which is different from the predicted sensory consequence, and the organism exhibits signs that demonstrate its awareness that its predictions have been falsified, we have a well-defined method for accessing an organism's knowledge. (This idea underlies a very popular experimental paradigm, the 'habituation paradigm', often used in research with very young children. Cf., for example, Wynn, 1992.) What we should do is to have a predicting robot react in specific ways when its predictions are falsified, for example by showing "surprise". In this way, we may use this experimental paradigm to assess a robot's knowledge with respect to both the distinction self vs. environment and the distinction private vs. public experience.

For example, a robot which moves its arm in order to reach an object, will predict that when the hand will make visually perceived physical contact with the object, there will be a tactile sensory input from the hand. If we manipulate the robot so that this tactile input fails to appear, the failed prediction will generate some "surprise" behavior in the robot. On the contrary, if we cause a robot to have a tactile sensation in its hand when one object in the environment comes into physical contact with another object, equally the robot should show "surprise". These signs of "surprise" are one way of demonstrating that the robot's knowledge of its environment includes two very different parts: its own body and the rest.

The same approach can be applied with respect to the private vs. public distinction. Given some sensory input that originates in the external environment, a robot can

predict that other robots will react in specific ways because they also have access to this sensory input. Given sensory input that originates from within its own body, the robot will predict that the other robots will not respond because they do not have access to this sensory input. As we have discussed above, if we manipulate the conditions so that these predictions turn out to be wrong, we expect the robot to show "surprise" and this will demonstrate that the robot has incorporated in its knowledge/prediction ability a distinction between a public world and a private world.

Use of I/You Sentences

Of course, organisms can also possess a more explicit and "conscious" knowledge of their own body and, more generally, of themselves, compared to the knowledge demonstrated by the robots discussed so far. This is typically true of humans because humans have language and, using language, they can explicitly communicate their "self-knowledge" to both others and themselves. Given our operational approach, we will consider a very specific ability of robots that would demonstrate this more explicit and "conscious" knowledge of their body and, more generally, of themselves. We assume that our robots already possess language, that is, they are able to both produce sequences of sounds associating a specific meaning to the sequence and understand sequences of sounds when the sounds are produced by other robots or by humans. We are proposing that to demonstrate a "sense of self" robots should not only be able to use language in general but they should be able to use one particular type of sentences, that is, sentences that contain the words "I" and "you" (I/You-sentences). We might construct robots that are able to produce/understand all sorts of sentences except I/You-sentences. For example, a robot might be asked "Where was John born?" and the robot would respond appropriately "John was born in Chicago". However, when asked "Where were you born?", this kind of robot would be unable to respond "I was born in Tokyo" (assuming that the robot was born in Tokyo). The robot's ability to understand the question about John and to respond appropriately shows that the robot can use proper names, that is, words that refer to individual entities (people, towns, etc.). But although the words "I" and "you" refer in specific contexts to specific individual entities, they are clearly different from proper names for at least two reasons: they require that the robot has identified itself as one specific entity in the world different from all other entities, and that it understands when a sentence containing the word "you" is addressed to itself or to another robot/human. We claim that a robot that can use I/You sentences would possess one (higher) form of "sense of self".

To answer a question may require to retrieve some knowledge already existing in one's memory but in some

cases it can presuppose an ability to acquire new knowledge. For example to answer the question “What time is it?” normally one has to consult a watch. Robots can be asked questions about themselves requiring knowledge that they do not already possess, and in particular questions concerning their body as a physical entity (“What is your weight?”), motivational and emotional states (“Are you hungry?”, “Are you preoccupied?”), mental states (“What you expect?”, “What you think?”). To answer these questions presupposes that the robot has identified some particular physical object as its own body, and can articulate its world as containing not only physical objects and events that can be accessed by other robots/humans (public world) but also motivational, emotional, and mental states that can only be accessed by itself and for which the particular robot is the only, final, authority (private world). Therefore, it can be said that a robot that can answer these questions is a robot which has a “sense of self”.

But the relationship between the use of words such as ‘I’ and ‘You’ and the sense of self might be even stronger than this. In fact, the use of I/You sentences might not only testify the presence of a preexisting self-awareness, it may also significantly contribute to the formation of higher levels of such a self-awareness. Language is in fact not only a complex communication system, but it is also a powerful cognitive tool, which influences and transforms all cognitive functions (Vygotsky 1978). Linguistic labels transform human thinking in many ways (Clark, 2006; Mirolli and Parisi, submitted), one of the most important one being the provision of ‘anchors’ for thought: by creating a label for a concept language makes it possible to attend and reflect on that same concept, acquiring a higher level of awareness of it. This is true for any concept, including the concept of a self. Hence, the meaningful use of sentences containing the words ‘I’ and ‘You’ by a robot may not only demonstrate the robot’s self-knowledge but it may also constitute a fundamental step towards higher levels of such a self-knowledge, and hence of robot’s consciousness.

Acknowledgements

The research presented in the 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

Acosta-Calderon, C.A. and Hu, H. 2005 Robot imitation: body schema and body percept. *ABBI*, 2, 131-148.
 Blackmore, S. 2003 *Consciousness: an Introduction*. Oxford: Oxford University Press
 Chalmers, D. 1996 *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press

Chella, A. and R. Manzotti eds. 2007 *Artificial Consciousness*. New York: Imprint-Academic.
 Clark, A. 2006 Language, embodiment, and the cognitive niche. *Trends in Cognitive Sciences*, 10, 370-374
 Dennett, D. 1991 *Consciousness Explained*. Boston: Little & Company
 De Preester, H. and Knockaert, V. eds 2005 *Body Image and Body Schema: Interdisciplinary Perspectives (Advances in Consciousness Research)*. Amsterdam: John Benjamins.
 Edelman, G.M. 2006 *Second Nature: Brain Science and Human Knowledge*. New Haven, Yale: Yale University Press.
 Edelman, G.M. and Tononi, G. 2000. *A Universe of Consciousness. How Matter Becomes Imagination*. London: Allen Lane.
 Hafner, V.V. and Kaplan, F. 2005 Interpersonal maps and the body correspondence problem. *Proceedings of the AISB Third International Conference in Imitation in Animals and Artifacts*. Hatfield, U.K., 48-53.
 Holland, O. ed. 2003 *Machine Consciousness*. New York: Imprint-Academic.
 Jordan, M.I. and Rumelhart, D.E. 1992 Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16, 307-354.
 Mirolli, M. and Parisi, D. submitted Towards a Vygotskian cognitive robotics: the role of language as a cognitive tool. *New Ideas in Psychology*.
 Morasso, P. 2007 The critical role of haptic perception: consciousness as the emergent property of the interaction between brain, body, and the environment. In A. Chella and R. Manzotti eds. *Artificial Consciousness*. Exeter: Imprint-Academic.
 Nolfi, S. and Marocco, D. 2002 Evolving robots able to integrate sensory motor information over time. In R.J. Duro, J. Santos and M. Grana eds. *Biologically Inspired Robot Behavior Engineering*. Heidelberg: Springer, 199-214.
 Parisi, D. 2004 Internal robotics. *Connection Science*, 16, 325-338
 Parisi, D. 2007 Mental robotics. In A. Chella and R. Manzotti eds. *Artificial Consciousness*. Exeter: Imprint-Academic.
 Pfeifer, R. and Scheier, C. 2001 *Understanding Intelligence*. Cambridge, MIT Press
 Schlesinger, M. and Parisi, D. 2001 Multimodal control of reaching: Simulating the role of tactile feedback. *IEEE Transactions on Evolutionary Computation*, 5, 122-128.
 Tani, J. 1998 An interpretation of the self from the dynamical systems perspective: a constructivist approach. *Journal of Consciousness Studies*, 5, 516-542.
 Turing, A. 1950 Computing machinery and intelligence. *Mind*, 236, 433-460
 Vygotsky, L.S. 1978 *Mind and society: The development of higher mental processes*. Cambridge: Harvard University Press
 Wynn, K. 1992 Addition and subtraction by human infants. *Nature*, 358, 749-750