

# Why it is important to build robots capable of doing science

Răzvan V. Florian

Center for Cognitive and Neural Studies (Coneural)

Str. Saturn 24, 3400 Cluj-Napoca, Romania

florian@coneural.org

## Abstract

Science, like any other cognitive activity, is grounded in the sensorimotor interaction of our bodies with the environment. Human embodiment thus constrains the class of scientific concepts and theories which are accessible to us. The paper explores the possibility of doing science with artificial cognitive agents, in the framework of an interactionist-constructivist cognitive model of science. Intelligent robots, by virtue of having different sensorimotor capabilities, may overcome the fundamental limitations of human science and provide important technological innovations. Mathematics and nanophysics are prime candidates for being studied by artificial scientists.

## 1. Introduction

Science is one of the highest achievements of human cognition, and its technological applications are certainly extremely important to human civilization. This paper argues that, in spite of the impressive results achieved so far, human science has fundamental limitations, given by our embodiment. In order to ensure the unconstrained advancement of science, intelligent robots capable of doing science should be built.

First, we briefly present an interactionist - constructivist cognitive model of science. It sets a framework for discussing the capabilities of cognitive agents needed for doing science and the relationship between science and embodiment. We introduce next artificially made science and argue that it may overcome the limitations of human science. Mathematics seems to be the field for which it is easiest to build artificial scientists, but their most important contribution will be in domains distant from human sensorimotor experience, such as nanotechnology. We also compare the current approach to artificial science with previous work in the field of automated science and machine discovery.

## 2. Science – a cognitive point of view

Science is a systematic process by which we come to understand the structure of the surrounding environment, to generate predictions about evolutions and explanations of causality in this environment. Social factors influence the evolution of science, by establishing socially accepted truth, setting research agendas and through the communication of scientific results. But before science becomes a social issue, scientific concepts and theories must be generated, understood and verified by the scientists themselves. Here the focus will be on the cognitive properties of a single individual needed for science.

### 2.1 *Interactionist-constructivist cognition*

The cognitive model of cognition presented here has as epistemological framework the interactionist-constructivist model of cognition of Indurkha (1992), originally developed for explaining the understanding of similarity creating metaphors. According to this theory, reality does not have a mind-independent ontology: the concepts of the cognitive agent impose an ontology to the world. The concepts are internal to the agent and generated by its sensorimotor interaction with the environment. However, the structure of the world with respect to this ontology cannot be arbitrary: reality has a structure external to the cognitive agent. It is autonomous and manifested in the structure of the sensorimotor data set, directly accessible to the agent. We may assume that reality exists prior to conceptualization, but its ontology does not.

While some of the empirical findings that motivated these considerations come from cognitive anthropology and psychology (Indurkha, 1992, pp. 94–111), this theory is also largely compatible with the results of modern science. It has been established that it is pointless to look for the final theory of the world, the final truth, which would be the equivalent of discovering its “real” ontology. A scientific theory can only be refuted, not justified, and may be considered valid only until its refuta-

tion. This follows not only from theoretical considerations (Popper, 1959), but was also shown by modern physics (Feynman, 1992, chap. 7). Moreover, a theory should only approximate relevant aspects of reality, balancing predictive power with complexity, otherwise its details would be unmanageable for our finite cognitive capabilities (Tolman, 1932). A scientific theory is however useful for its coherence (viability, functional fitness) with a limited domain of the world (see also Peschl, 2001). Quantum mechanics and the theory of relativity are incompatible theories, with different ontologies; however, they are both accepted as valid theories for their explanatory domain. But scientific theories cannot be arbitrary: nature resists their predictions. The result of an experiment is given solely by the structure of reality.

As we will show later, an interactionist - constructivist model of cognition also agrees with the latest results in neuroscience and cognitive science.

## *2.2 Neurobiological foundations for knowledge and science*

It is generally accepted that the biological basis of cognition is the neuronal activity. Part of the interaction of the environment with the organism is captured by sensors and is translated in neuronal patterns of activation. Activation of motor neurons results in actions that are externalized in the environment and may change it, or change its perception by the agent (as in a movement). Synapses transmit neuronal activation to other neurons, and in general it is sustained even in the absence of significant sensorial input.

Quasi-hebbian learning, associating concomitant activations, is widely considered as an important learning mechanism in biological neural networks and it was demonstrated experimentally in some areas of the mammalian brain (Kandel, Schwartz, & Jessel, 2000, chap. 63; Rolls & Treves, 1998, chap. 1; Fuster, 1995, chap. 3). In general, learning corresponds to changes in the chemical composition of the neuronal environment, and modifications of synaptic strengths and synaptic connectivity. Long term modifications in the synaptic connectivity yields a neural constructivism (Quartz and Sejnowski, 1997) that mirrors and is the physical basis of the conceptual one. These changes are caused by the neuronal activity, which is largely environmentally derived: learning and cognitive development in biological agents is thus inextricably related to their sensorimotor capabilities, and thus to their embodiment (see also Tolman, 1932, Gibson, 1979, MacDorman et al., 2000).

The representational properties of the brain, which are mostly a result of learning, thus also largely depend on the sensorimotor capabilities of the sup-

porting body. Moreover, it is believed that imagery and short term memory share many common neural mechanisms with perception or motor action (Kosslyn and Thompson, 2000, Fuster, 1995, Jeannerod, 1994, Jeannerod, 1999). Many results point out that the neural correlates of a certain concept, activated, for example, by a word, are activations of the neural networks that were also active during the experiences of the person with the significant of that word (Damasio, 1990, Pulvermuller, 1999, Martin et al., 2000). The representational properties of a symbol cannot thus go further than the perceptual, motor or emotional states that are associated by learning with the phonological or visual form of the symbol. It was argued that even the understanding of abstract concepts is ultimately grounded like this (Barsalou, 1999). Such grounded representations are internal to the cognitive agent, and do not have the shortcomings of the classical symbolic representations (Harnad, 1990, Bickhard, 1993, Pfeifer and Scheier, 1999, Ziemke, 2001). Further associations between symbols are also possible, but ultimately their representational content will be grounded in the sensorimotor states associated with them.

These facts seem to confirm an interactionist-constructivist view of cognition: representations depend on the interaction of the cognitive agent with the external environment and are constructed according to his individual history of interactions. Previous experiences induce long term changes in the synaptic connectivity, and thus each perception is influenced by the past. Each agent thus has a different ontology of the world, the differences being attenuated only by genetically induced similarities between bodies, commonality of the environments and by communication.

Motor capability is of the utmost importance for the possibility of representation for the cognitive agent himself (in contrast to the representation for the user and designer, as in classical symbol-based artificial intelligence). Bickhard (1993) argues on theoretical grounds that genuine representational content can be generated only by an embodied, action capable, goal directed agent. This content is in fact the potentialities for action activated by the current perceptual input and by the internal state of an agent (which correspond to the "behavioral possibilities" of Tolman, 1932, and to the "affordances" of Gibson, 1979). In robotics, the importance of sensorimotor coordination was demonstrated for solving the problems that appear in information processing approaches (Pfeifer and Scheier, 1999, Steels and Brooks, 1995).

In the primate brain, the main locus of integration between the perceptual and the motor pathways is

the prefrontal cortex, which is considered essential for certain types of memory, but also for planning, initiation of action, and creativity (Fuster, 1995). On one hand, associations between perceptual and sensorimotor networks capture part of the environment structure, together with the structure that may be detected from a static sensorial input or from the temporal changes of sensorial input. On the other hand, mutual activations between these networks may allow planning of actions by mental simulation and the prediction of results. For example, it was experimentally shown that visuomotor anticipation—the prediction of the visual consequences of a future motor action—is likely to be also the mechanism that drives mental rotation (Wexler et al., 1998). It is thus possible for the cognitive agent to internally simulate and predict the evolution of the environment in reaction to a given action or to a chain of actions. Causality may also be represented, alternatively, by direct associations between cause and result, without temporal continuity.

### 2.3 *A cognitive model of science*

We are now close to introducing science in this picture. The technological breakthroughs made possible by science in the last few centuries are relatively recent on evolutionary timescales. The cognitive mechanisms that allow us to do science could not be favored by evolution for science itself. Those mechanisms must thus be the same as those involved in more mundane, though evolutionary adaptive, cognitive processes (Nersessian, 1992, Peschl, 1999, Peschl, 2001). Some of these mechanisms are:

- Causality detection and mental simulation. Repeated associations between actions and perceptions of the results, or between evolution in time of perceptions, may yield permanent associations under the form of causality relations. Further, once these causalities are learned, one may mentally simulate chains of actions and their causally related results to construct plans of actions and predict results of these plans. This is an adaptive mechanism, which may be found also in animals, and which serves in science as a basis for causality detection and the formulation of predictions. Computational modeling of simulation and prediction based on sensorimotor experience is an active field of research (Jirenhed et al., 2001, Stojanov, 2001, MacDorman et al., 2000, Clark, 2001).
- Coherence detection. When executing a mentally simulated plan, the result may or may not be the same as what was planned. It is adaptive for the agent to evaluate the coherence of the predicted and actual result and to consequently enforce the causal associations used in the plan, or alterna-

tively to loose them, or to explain the result with an alternative causal chain.

- Projective reasoning. This comprises the type of reasoning used for the understanding and generation of similes, analogies, metaphors and models. It may be considered a “projection” of the structure of a source domain on a target domain (Indurkha, 1992). While a definite, biologically plausible model of this phenomenon remains to be developed, we may speculate that it arises from simultaneous activation of the neural networks activated with the source and target domain. The common features of the two domains are revealed, if there are any. If the target domain has little sensorimotor structure associated with it, the structure of the source domain will dominate the conceptualization of the target. However, an asymmetry between the source and the target domain always exists during a metaphor or an analogy, and a mechanism that accounts for it has to be determined. Emergent capability for projective reasoning should be an important test for every model of human knowledge representation.

Reasoning by analogy is also an adaptive mechanism, as it may suggest ways to deal with new situations, based on experience acquired with previous situations. While nonhuman primates can successfully reason analogically after training only, young children do it naturally (Holyoak and Thagard, 1996) and have no problems understanding metaphors.

- Abstraction. Abstraction implies extraction of common sensorimotor structure and may be related to categorization, which is easily obtained with neural networks. Another type of abstraction is related to the schematicity of concepts (Barsalou, 1999), which is implicitly realized if the concepts are grounded in distributed neural networks.
- Symbolic association. This capability may have evolved for communication, which is also an adaptive mechanism.
- Subitizing. It was shown that newborns and some animals are capable to precisely and instantly discriminate (without counting) between small quantities, up to numbers of the order of 4 or 5 (Wynn, 1993, Dehaene, 1997). The numerosity of greater quantities can be also imprecisely estimated by animals, the variability being proportional with the magnitude (Gallistel and Gelman, 2000). This is also an adaptive mechanism, which may be used to bootstrap the understanding of numbers. Subitiz-

ing was simulated with connectionist models (Dehaene and Changeux, 1993).

Out of these mechanisms, the most important is probably the projective reasoning, as it allows the formation of theories and the reasoning about domains of reality where there is no direct sensorimotor access. Through projective reasoning, the structure of a (source) domain, learned by sensorimotor interaction with it, is projected to a different (target) domain. If the projection is coherent with the target domain, it may allow predictions, through simulations that follow the sensorimotor associations in the source domain. Moreover, the source domain can induce a different structure in the target domain, thus inducing a creative conceptualization of this domain (Indurkha, 1992).

For example, we may know that if we put a cup of water in the fridge, or if we put it outside in the winter, it will transform into ice. A sensorimotor conceptualization of this target domain would involve associations between the perception of cold, visual and tactile perceptions of water and ice, the motor actions involved in the experiment, other associations with previous encounters with water and ice, and so on. However, we may also think about this phenomenon in the terms of the molecular structure of water: the solidification would thus be seen as a change in the movement of the molecules. The source domain would be here spatial perception (we would imagine the molecules ranged orderly in three dimensional space) and the sensorimotor interactions with objects, which serve as a source for the conceptualization of the molecules. The sensorimotor grounding of the source domain is thus different than the one of the target domain, and through projection it imposes a new conceptualization of water and ice.

The mechanism of projective reasoning may explain also the incommensurability of different paradigms (Kuhn, 1962). If in different paradigms there are different source domains that structure the target domain, and if there are no similarities between the different source domains that could be abstracted, there is simply no way to reconcile the two views: they are simply different and activate different neural networks. Projective reasoning also introduces an extra degree of variability in the ontology of the world, which has special importance in science: the ontology not only depends on the experiential history of the cognitive agent, but also on which sensorimotor domain it is grounded on, and in which way.

Abstraction is another important cognitive mechanism used in science. For example, the concept of number is abstracted from the manipulation and construction of objects, measurement of linear dimensions and quantities, and movement on a path (Lakoff and Nunez, 2000). They are several differ-

ent sensorimotor domains, but they have a certain common structure that is abstracted in the concept of number. This structure is also associated with subitizing mechanisms, counting, and the different forms of symbolic representation of numbers (Dehaene, 2000, Pesenti et al., 2000, Dehaene, 1992, Dehaene and Cohen, 1995).

In general, the structure of a source domain has a certain plasticity; some new associations are added, and some older associations are inhibited, after the application in a target domain. This modified structure may be abstracted, or restructured with a projection from another domain, in a more and more elaborated construction that eventually leads to the creation and understanding of modern scientific concepts.

Science is thus the process of selecting source domains and adapting their structure, initially acquired from sensorimotor processes, such as the resulting structure is coherent with the structure of a certain domain of the environment where is no direct sensorimotor access, or where the sensorimotor generated structure is not rich enough. The projected structure of the source domain allows mental simulations, which yield predictions in the target domain.

### 3. Limitations of human science

Human scientific concepts thus crucially depend on human sensorimotor capabilities, given by the human body. The limited range of these sensorimotor capabilities thus fundamentally limits the class of abstract concepts, including scientific concepts, that a human can understand and use. The animal world exemplifies some biologically implemented sensorial capabilities that are beyond human experience, such as space perception through sonar-like interactions, magnetoreception and electroreception (Hughes, 1999). Presumably, if we had these perceptual pathways in addition to our current ones, theories about waves and electromagnetism would have been much simpler to generate and acquire, and much closer to our intuition.

Artifacts obtained with current technologies may extend further the domain, spectral range, precision and intensity of receptors and effectors. Measurement instruments that use those enhanced sensors and effectors are routinely used in science. Via transducers, the signals can be perceived by the human sensorimotor apparatus, and human movements can be translated to other types of actions, thus leading to novel sensorimotor couplings, which may associate an extra sensorimotor grounding with the theories. For example, it was argued that tools extend action and perception capabilities (Hirose, 2002). The sensorimotor contingencies generated by the use of tools and instruments integrate smoothly, after training, with the sensorimotor con-

tingencies of our own body (O'Regan and Noe, 2001, Stojanov and Gerbino, 1999).

However, current scientific theories constrain the design of the experimental apparatus, and the output of the measurement instruments has to be accessible to human sensors. The expansion of this productive cycle, from theories to new sensorimotor groundings and back, thus cannot fully escape the limitations of human body.

Other limitations of human science may come from limits of short-term memory and slow reasoning performance (Riegler, 1998). General limitations of human cognition were also discussed by McGinn (1994).

## 4. Artificial science

### 4.1 *Introducing artificial scientists*

Having reviewed the cognitive capabilities needed for science, it is naturally to ask if it is possible to implement them in artificial cognitive agents. It seems that there is, in principle, no impediment for such agents to develop science—artificial science.

As discussed in Section 2.3, they would have to be embodied and to have both sensors and effectors. They will conceptualize on their own their environment. At first, this conceptualization will be through sensorimotor interaction. Later, these sensorimotor structures generated for some domains of the environment may be projected on other domains. As human science has shown, if the structure of the sensorimotor data set is rich enough, it may be coherent with other structures from other parts of the environment. We may ensure this emergent phenomenon in the artificial scientists by giving them access to a complex environment and to a wide range of sensors and effectors. The control apparatus of the agents, which would probably be implemented in artificial neural networks, will have to implement causality detection, internal simulation, abstraction, projective reasoning and eventually symbolic association. In biological agents, the need for coherence is imposed by their need for survival: a lack of coherence may result in injuries or death. In the artificial agents, the goal may be simply to maximize the coherence and the diversity of their predictions. Alternatively, their goal may be to conceive and later build technological applications of their science. In this way, they may generate technological innovations.

For example, with the recent advent of nanotechnology, cognitive agents could be build that would have direct perceptual access to quantum phenomena. Quantum mechanics would then be at least as easy to understand for them as classical mechanics is for humans. Moreover, it is possible that their conception of quantum mechanics would be much simpler than the conception that humans painfully acquire through an elaborate construction from an

unadapted grounding. This new way of conceiving quantum phenomena may lead to novel applications.

In general, having a different embodiment, the artificial agents will have access to different classes of concepts than the one available in human science. The science generated by artificial agents will thus escape the limits imposed by the human body.

Once enough of the structure of the environment is acquired through sensorimotor interaction, the artificial scientist could also continue the scientific process of searching coherences “offline”, without permanent interaction with the environment.

### 4.2 *Communication between humans and artificial scientists*

As in any other communication between two agents with different types of bodies, the communication between humans and the artificial scientists may prove difficult. We will not be able to understand in their entirety the conceptualizations of the artificial agents. Understanding an utterance means an internal simulation involving the sensorimotor significants of the communicated symbols, previously learned by association. As the sensorimotor groundings of the agents are different than ours, there will be an important loss of information in the communication between the two parties. The situation may be different if the environment and the sensorimotor capabilities of the agents are moderately similar to the human ones.

### 4.3 *Artificial mathematics*

Many results point out that the groundings of human mathematics are mainly the conceptualization of space and objects. A brain imagery study (Dehaene et al., 1999) has shown that arithmetic uses bilateral areas of the parietal lobes involved in visuo-spatial processing, for estimation of numerical magnitudes, besides part of the brain involved in word association, for the addition tables. Lakoff and Nunez (1999, 2000) have theoretically studied the embodiment of the mathematical concepts. Their results also point to the (active) perception of space and objects as the grounding of mathematics. Many psychological experiments have shown a strong correlation between spatial abilities and mathematical and scientific abilities (Siemankowski and MacKnight, 1971, Poole and Stanley, 1972, Bishop, 1973, Guay and MacDaniel, 1977, Mitchelmore, 1980, Pallrand and Seeber, 1984). It seems thus that an artificial agent that will acquire an accurate conceptualization of space and object perception and manipulation, through sensorimotor interaction, and which will also have symbolic capabilities, will be able to understand and generate mathematics, and eventually, later, classical mechanics.

There already exist many neuroscience results about space and object perception, as they may be collected also from direct brain recordings from animals. They may guide the realization of the artificial mathematician. The discreteness of mathematics, based on the discreteness of objects, may ease the communication between the agent and humans. The results of human mathematics may be communicated to the artificial mathematician, and at its turn it may provide us its innovative results. Finally, the result would be an artifact that would be able to master the whole human mathematics and continuously generate new mathematical results (which yields the question, maybe senseless until now, if interesting mathematics could ever be complete, or infinitely new important results may be generated). The artificial mathematician will be able to escape the limits of Gödel's theorem, as a human mathematician does, because its demonstrations will not be limited to the use of a single grounding. Through projective reasoning, it will be able to change groundings, and use results from a grounding in different ones, as a human does.

#### 4.4 Comparison with previous approaches

There is an important volume of work in artificial intelligence dedicated to the computational discovery of new scientific knowledge. Some of this work was successful and concluded with publications in speciality journals of the scientific results discovered by the programs (for reviews, see Langley, 2000, Langley, 1998, Colton and Steel, 1999, Valdes-Perez, 1999, Langley et al., 1993). However, this work is only superficially similar with the approach proposed here. In general, previous approaches to machine discovery or automated science were concerned with clustering, searching of qualitative and quantitative laws, and formulation of models, always using human provided data. It is up to the developer of the automated discovery program to choose a problem formulation adapted to the input data, and also to choose a representation for the data fed to the program. Often, the input data is preprocessed and filtered by the developer, and the developer manipulates the program algorithm to modulate its performance for certain inputs (Langley, 2000).

In most cases, the program has no effectors and has no access to the environmental phenomena studied, other than the preprocessed data. An exception is a project of closed-loop scientific discovery, where experiments are planned automatically and carried out by robots (Bryant et al., 2001); that paper also introduces the term "artificial scientist". However, in all cases the ontology of the world, as seen by the program carrying out the automatic scientific discovery work, is fixed by its human developer.

In contrast, the approach presented here suggests

that the most interesting fact about doing science with robots is that they are able to come up with novel ontologies of the world, by virtue of having an embodiment different than the human one. This is why they should have both perceptual and effector access to the phenomena they study, and the capability to conceptualize on their own the environment, according to their sensorimotor affordances.

## 5. Conclusion

The paper presented a cognitive model of science and a blueprint for artificial scientists. While they are sometimes speculative, they open a research direction that may finally yield important technological applications. Results from neuroscience, artificial life and robotics accumulate at a fast rate, and may soon prove that the construction of an artificial mathematician is a question of years.

## 6. Acknowledgements

This work was supported by Arxia. I thank to Bipin Indurkha, Mihai Preda, Lucia Faiciuc and two anonymous reviewers for useful suggestions.

## References

- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, 22:577–660.
- Bickhard, M. (1993). Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5:285–333.
- Bishop, A. (1973). Use of structural apparatus and spatial ability: A possible relationship. *Research in Education*, 9:43–49.
- Bryant, C. H., Muggleton, S. H., Oliver, S. G., Kell, D. B., Reiser, P., and King, R. D. (2001). Combining inductive logic programming, active learning and robotics to discover the function of genes. *Electronic Transactions on Artificial Intelligence*, 5:1–36.
- Clark, J. J. (2001). Learning sensorimotor contingencies. *Workshop on the Genesis of Perception and the Notion of Space in Machines and Humans, Paris, France*.
- Colton, S. and Steel, G. (1999). Artificial intelligence and scientific creativity. *Journal of the Society for Artificial Intelligence and the Study of Behaviour*, 102.
- Damasio, A. (1990). Category-related recognition defects as a clue to the neural substrates of knowledge. *Trends in Neurosciences*, 13:95–98.

- Dehaene, S. (1992). Varieties of numerical abilities. *Cognition*, 44:1–42.
- Dehaene, S. (1997). *The number sense: How the mind creates mathematics*. Oxford University Press, Oxford, UK.
- Dehaene, S. (2000). Cerebral bases of number processing and calculation. In Gazzaniga, M. S., (Ed.), *The new cognitive neurosciences, 2nd ed.*, pages 987–998. MIT Press, Cambridge, MA.
- Dehaene, S. and Changeux, J.-P. (1993). Development of elementary numerical abilities: A neuronal model. *Journal of Cognitive Neuroscience*, 5 (4):390–407.
- Dehaene, S. and Cohen, L. (1995). Towards an anatomical and functional model of number processing. *Mathematical Cognition*, 1:83–120.
- Dehaene, S., Spelke, E., Pinel, P., Stanescu, R., and Tsivkin, S. (1999). Sources of mathematical thinking: Behavioral and brain-imaging evidence. *Science*, 284:970–974.
- Feynman, R. P. (1965 / 1992). *The character of physical law*. Penguin Books, London, UK.
- Fuster, J. M. (1995). *Memory in the cerebral cortex: An empirical approach to neural networks in the human and nonhuman primate*. MIT Press, Cambridge, MA.
- Gallistel, C. R. and Gelman, R. (2000). Non-verbal numerical cognition: from reals to integers. *Trends in Cognitive Science*, 4 (2):59–65.
- Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin, Boston, MA.
- Guay, R. and McDaniel, E. (1977). The relationship between mathematics achievement and spatial abilities among elementary school children. *Journal for Research in Mathematics Education*, 8 (3):211–215.
- Harnad, S. (1990). The symbol grounding problem. *Physica D*, 42:335–346.
- Hirose, N. (2002). An ecological approach to embodiment and cognition. *Cognitive Systems Research*, in press.
- Holyoak, K. J. and Thagard, P. (1996). *Mental leaps: Analogy in creative thought*. MIT Press, Cambridge, MA.
- Hughes, H. C. (1999). *Sensory exotica: A world beyond human experience*. MIT Press, Cambridge, MA.
- Indurkha, B. (1992). *Metaphor and cognition: An interactionist approach*. Kluwer Academic Publishers, Dordrecht, the Netherlands.
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behavioral and Brain Sciences*, 17 (2):187–245.
- Jeannerod, M. (1999). *The cognitive neuroscience of action*. Blackwell.
- Jirenhed, D.-A., Hesselow, G., and Ziemke, T. (2001). Exploring internal simulation of perception in mobile robots. *Lund University Cognitive Studies*, 86:107–113.
- Kandel, R., Schwartz, J. H., and Jessell, T. M., (Eds.) (2000). *Principles of neural science*. McGraw-Hill.
- Kosslyn, S. M. and Thompson, W. L. (2000). Shared mechanisms in visual imagery and visual perception: Insights from cognitive neuroscience. In Gazzaniga, M. S., (Ed.), *The new cognitive neurosciences, 2nd edition*, pages 975–985. MIT Press, Cambridge, MA.
- Kuhn, T. (1962). *The structure of scientific revolutions*. University of Chicago Press, Chicago.
- Lakoff, G. and Nunez, R. (1999). The metaphorical structure of mathematics: sketching out cognitive foundations for a mind-based mathematics. In English, L., (Ed.), *Mathematical Reasoning: Analogies, Metaphors and Images*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Lakoff, G. and Nunez, R. (2000). *Where mathematics comes from: How the embodied mind brings mathematics into being*. Basic Books.
- Langley, P. (1998). The computer-aided discovery of scientific knowledge. In *Proceedings of the 1st International Conference on Discovery Science*. Springer-Verlag, New York, NY.
- Langley, P. (2000). The computational support of scientific discovery. *International Journal of Human-Computer Studies*, 53:393–410.
- Langley, P., Simon, H. A., Bradshaw, G. L., and Zytkow, J. M. (1993). Scientific discovery. In Goldman, A. I., (Ed.), *Readings in philosophy and cognitive science*. MIT Press, Cambridge, MA.
- MacDorman, K. F., Tatani, K., Miyazaki, Y., and Koeda, M. (2000). Proto-symbol emergence. *Proceedings of IROS-2000: IEEE/RSJ International Conference on Intelligent Robots and Systems, Kagawa University, Takamatsu, Japan*.

- Martin, A., Ungerleider, L. G., and Haxby, J. V. (2000). Category specificity and the brain: The sensory/motor model of semantic representations of objects. In Gazzaniga, M. S., (Ed.), *The new cognitive neurosciences, 2nd edition*, pages 1023–1036. MIT Press, Cambridge, MA.
- McGinn, C. (1994). The problem of philosophy. *Philosophical Studies*, 76:133–156.
- Mitchelmore, M. (1980). Prediction of developmental stages in representation of regular space figures. *Journal for Research in Mathematics Education*, 11 (2):83–93.
- Nersessian, N. J. (1992). How do scientists think. In Giere, R., (Ed.), *Cognitive models of science*. University of Minnesota Press.
- O'Regan, J. and Noe, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioural and Brain Sciences*, 24:5.
- Pallrand, G. and Seeber, F. (1984). Spatial ability and achievement in introductory physics. *Journal of Research in Science Teaching*, 21 (5):507–516.
- Peschl, M. F. (1999). The development of scientific concepts and their embodiment in the representational activities of cognitive systems: Neural representation spaces, theory spaces, and paradigmatic shifts. In Loocke, P. V., (Ed.), *The nature of concepts: Evolution, structure and representation*. Routledge, London, UK.
- Peschl, M. F. (2001). Constructivism, cognition, and science - an investigation of its links and possible shortcomings. *Foundations of Science*, 6 (1-3):125–161.
- Pesenti, M., Thioux, M., Seron, X., and Volder, A. D. (2000). Neuroanatomical substrates of arabic number processing, numerical comparison, and simple addition: A PET study. *Journal of Cognitive Neuroscience*, 12:461–479.
- Pfeifer, R. and Scheier, C. (1999). *Understanding intelligence*. MIT Press, Cambridge, MA.
- Poole, C. and Stanley, G. (1972). A factorial and predictive study of spatial abilities. *Australian Journal of Psychology*, 24 (3):317–320.
- Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson, London, UK.
- Pulvermuller, F. (1999). Words in the brain's language. *Behavioral and Brain Sciences*, 22:253–336.
- Quartz, S. R. and Sejnowski, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences*, 20 (4):537–596.
- Riegler, A. (1998). "The end of science": Can we overcome cognitive limitations? *Evolution and Cognition*, 4 (1):37–50.
- Rolls, E. T. and Treves, A. (1998). *Neural networks and brain function*. Oxford University Press, Oxford, UK.
- Siemankowski, G. and MacKnight, F. (1971). Spatial cognition: Success prognosticator in college science courses. *Journal of College Science Teaching*, 1 (1):56–59.
- Steels, L. and Brooks, R., (Eds.) (1995). *The artificial life route to artificial intelligence: Building embodied, situated agents*. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Stojanov, G. (2001). Petitage: A case study in developmental robotics. *Lund University Cognitive Studies*, 85.
- Stojanov, G. and Gerbino, W. (1999). Investigating perception in humans inhabiting simple virtual environments. *European Conference on Visual Perception, Trieste*.
- Tolman, E. C. (1932). *Purposive behaviour in animals and men*. Century, New York, NY.
- Valdes-Perez, R. E. (1999). Discovery tools for science apps. *Communications of the ACM*, 42 (11):37.
- Wexler, M., Kosslyn, S. M., and Berthoz, A. (1998). Motor processes in mental rotation. *Cognition*, 68:77–94.
- Wynn, K. (1993). Evidence against empiricist accounts for the origins of numerical knowledge. In Goldman, A. I., (Ed.), *Readings in philosophy and cognitive science*. MIT Press, Cambridge, MA.
- Ziemke, T. (2001). The construction of 'reality' in the robot: Constructivist perspectives on situated artificial intelligence and adaptive robotics. *Foundations of Science*, 6:163–233.