

Cognitive science in Romania?

Răzvan Florian

Center for Cognitive and Neural Studies
Cluj-Napoca, Romania
florian@coneur.org

Abstract. I review herein some of the research directions in cognitive science, with a focus on artificial intelligence, connectionism, computational neuroscience, and artificial life. Further, I argue for developing research in these areas in Romania, as it may be conducted without expensive equipment, may benefit from existing skills in computer science and physics, and may generate innovative, high return applications.

1 Introduction

Cognitive science has emerged a few decades ago as an interdisciplinary field of research dedicated to the study of intelligence and intelligent systems, either natural or artificial. Psychology, neuroscience, artificial intelligence, linguistics, philosophy, and anthropology interact all within cognitive science toward developing an integrated theory of the brain and cognition.

First, I will review the present situation in the field, after briefly describing its historical evolution. Limited by the scope of this paper, I will focus mostly on the computational approaches to the study of cognition. They are somewhat closer to yield practical applications, but other branches of cognitive science have equal scientific importance, and results from all of them must be integrated in the development of computational models. In the second part of the paper, I will argue for the potential importance of the development of cognitive science in Romania. References to bibliography and internet resources will be presented.

2 A review

2.1 The beginning

The birth of cognitive science is considered to be a conference at MIT in 1956, where the talks of H. A. Simon, A. Newell, N. Chomsky and G. Miller suggested that experimental psychology, theoretical linguistics, and computational simulations of cognitive processes are parts of a unified whole. A Center for Cognitive Studies was established afterwards at Harvard. In 1976 the newly-founded Cognitive Science

Society [1] started to publish the journal *Cognitive Science*, thus institutionalizing the emergence of the new research field.

2.2 Traditional artificial intelligence; symbolic systems

The first years of cognitive studies were particularly marked by the influence of the computer, that was a relatively new technology at that time. Intelligent behaviour was often viewed as computation. It was thought that human intelligence is achieved by symbolizing external and internal situations and events and by manipulating these symbols according to syntactic rules [2, 3, 4]. The supporters of this so-called cognitivist or functionalist approach sustained that once the good algorithms and ways of representing knowledge in symbols would be found, intelligence can be implemented in any kind of computing machines, like computer software, regardless of the hardware implementation. Until the 80's, most of the models in cognitive science and cognitive psychology were inspired by the functioning of the computer and phrased in computer science and information processing terminology; some of these models continue to be backed today by their supporters. Representational structures such as feature lists (for example, *CHAIR[seat, back, legs]*), schemata and frames (knowledge structures that contain fixed structural information, with slots that accept a range of values), semantic networks (lists and trees representing connections between words) and production systems (a set of condition-action pairs used as rules in the execution of actions) were used to explain and simulate on computers cognitive processes [5, 6]. It was proposed that problem solving is accomplished by humans through representing achievable situations in a branching tree and then searching in this problem space [7]. It was also proposed that objects are recognized by analysis of discrete features or by decomposing them in primitive geometrical components [8]. In robotics, the efforts were directed towards building internal models of the world, on which the program could operate to produce a plan of action for the robot. Plans of actions needed to achieve a goal were represented, for example, as *drink (move (mouth, pickup (cup, start-state)))*. Traditional research in natural language processing (NLP) is also based on symbolic approaches [9].

The methods of this so-called Good Old Fashioned Artificial Intelligence (GOFAI) had some impressive successes in certain domains; however, these successes are limited. Based on those methods, programs were built that solved problems and proved theorems from logic and geometry. However, they depend on humans for converting the problem in a representation suitable for them and are confined to domains where knowledge can be easily formalized. Expert systems are widely used in the industry for the planning of the processes, but once the situation gets out of their ontology, they have no capability of dealing with it. One of the best known expert systems is MYCIN [10], a program for advising physicians on treating bacterial infections of the blood and meningitis. An example of MYCIN's limitations is to tell MYCIN that that Cholerae Vibrio was detected in the patient's intestines. The system will recommend two weeks of tetracycline and nothing else. This would probably kill the bacteria, but most likely the patient will be dead of cholera long before the two weeks. However, the physician will presumably know that the diarrhea has to be treated as well [11].

The defeat of the world chess champion, Garry Kasparov, by the Deep Blue computer in 1997 was widely publicized [12]. Another expert system, a recent computer program built on the FORR architecture [13] is capable of learning and successfully playing several types of games. However, a program that would beat a professional go player is still yet to be built, because the search space is much bigger in go than in other games. This is a good example where traditional methods fail.

Research in NLP have led to programs that are able to search and summarize text, to translate automatically, and to chat with a human partner. The state of the art programs in this field can be easily tested on the web [14, 15]: neither the word-by-word translation, nor the grammatical analysis of the phrase structure are enough to understand natural language. These problems point to the fact that understanding of the semantics and information about the context are crucial. The commercial Cyc project [16], still under development, struggled for more than ten years to build a huge semantic net that would cover the commonsense knowledge of an ordinary human. In spite of the huge quantity of information fed into computers, the results are well below expectations.

In general, most intuitive human knowledge still resists formalization, including that involved in comprehending simple stories or simple physical situations. Current artificial systems are usually brittle, in the sense that they are unable to adapt to situations unforeseen by their programmer. Many day-to-day human problems seem to have unmanageable computational complexities for systems designed in the framework of classic artificial intelligence (AI). Though it was proposed at times that the human brain

functions under similar principles, there is little direct evidence that symbol systems underlie human cognition [17]. Problems of those types of systems will be also discussed in section 2.5.

2.3 Connectionism and neural networks

The brain was an alternative source of inspiration for models of cognition. In 1943, McCulloch and Pitts [18] proposed a very simple model of a neuron as a binary threshold unit. They proved that groups of these units can compute logical functions and are capable in principle of universal computation. In the 60's, Rosenblatt [19] and Widrow and Hoff [20] studied the perceptrons, groups of units organized into layers with feed-forward connections between one layer and the next. For the simplest class of perceptrons without any intermediate layers, Rosenblatt showed the convergence of a learning algorithm, a way to change the weights of the connections between the units iteratively so that a desired computation was performed. This triggered hopes that such machines could be a basis for artificial intelligence. In 1969, Minsky and Papert [21] showed that some elementary computations, like the "exclusive or", cannot be performed by the one-layer perceptron, thus deflecting the attention of the mainstream scientific community from this direction [22]. However, important findings of Anderson, Kohonen, Hopfield and others [23] related to neural networks emerged in the meantime. Further, the backpropagation method proposed by Rumelhart and McClelland in 1986 [24] pushed this paradigm back into the mainstream research. The backpropagation method is an algorithm for adjusting the weights connecting units in successive layers of multi-layer perceptrons which solves many problems that are above the capabilities of the one-layer perceptron.

Artificial neural networks consist of large numbers of simple units (simulated on computers, or sometimes implemented in silicon electronic circuits) that vaguely resemble neurons (their activation is dependent on the weighted sum of the activations of the units that link to them). The activation function is usually nonlinear, corresponding to an activation threshold that is present in real neurons as well. Processing takes place through the propagation of activations between the units, via connections. Learning occurs by the adaptation of the connection weights [25].

The so-called connectionist or parallel distributed processing (PDP) models of cognition are widely used to explain a wide range of cognitive phenomena such as perception, memory and learning [24, 26], language [27], reading [28] and cognitive development [29]. Artificial neural networks are also used by engineers for tasks such as character recognition (OCR), speech recognition, interpretation of images, and even for controlling damaged military planes [30] or to control a robotic cat [31]. Connectionist models have had considerable success in many areas of cognition, but higher cognitive functions, such as reasoning and problem solving, are still to be developed.

2.4 Computational neuroscience

The classic artificial neural networks, although originally inspired by their biological counterparts, differ in many ways from them. The supervised learning methods used in connectionist models, where the error of the network's output is "backpropagated" through the layers, are unrealistic because the real neurons can use only information that is available locally, i.e. the firing of the neurons that connect to them and the surrounding biochemical environment. More biologically plausible learning rules are those originally proposed by Hebb [32], where, for example, a synapse increases in strength when there is coincident presynaptic and postsynaptic activity. This is possible through the mechanisms of long term synaptic potentiation [33, 34]. Many models use the firing rate of the neurons as the basic parameters; however, in many situations the exact timing of the spikes proves to be important, or the reaction time of the brain to a stimulus corresponds to the propagation of a single spike [35]. These imply that there is a need for simulating the evolution of the neuron's potential on a finer timescale [36, 37]. Also, the dependence of a neuron's activation on its inputs is usually more complex than a simple addition [38, 39, 34]. Strictly feedforward architectures used in many connectionist models are not found in real brains. Recurrent neural networks are therefore considered to be more realistic. Moreover, the number of neurons in the brain

largely outnumbers the number of units in classical connectionist models and in any detailed simulation possible today.

The relatively recent field of computational neuroscience [40] proposes itself to build realistic computational models of the neural mechanisms, by closely following empiric findings from the neurosciences. It is assumed that the emergent cognitive capabilities of the brain are high-level effects that depend in a systematic way on the lower level phenomena. Computational neuroscience methods were successful in modeling some relatively low-level cerebral modules and functions. For example, a recent neural network model of the perirhinal cortex achieves accurate familiarity discrimination of stimuli, also showing activity patterns of the simulated neurons that are very similar to those recorded from the brains of primates [41]. This model predicts that the memory can discriminate the familiarity of about 10^8 patterns (while only about 10^5 patterns can be recalled at will, as other models have shown). Another recent model [42] explains the geometric visual hallucinations seen by many observers after taking hallucinogens, as emerging patterns of activity in the primary visual cortex.

2.5 Embodiment, situatedness, a-life

Results from a wide range of domains within cognitive science recently converged to express the need for “embodied” or “situated” theories of cognition, as opposed to the approaches that consider that all the aspects of human intelligence can be captured by symbol processing. It is considered that a genuine intelligent artificial agent must have a body, existing as a physical entity in the real world, and must be embedded in the situations it can conceptualize, by acquiring information about its environment through its sensors and interacting with the environment through effectors. Abstract concepts are also grounded in this interaction.

It was argued by Bickhard that symbolic representations only cannot provide representational content for the cognitive systems that use them [43]. Symbolic representations are the main components of classical cognitivist models, but also used in many connectionist models where the inputs or the outputs of the networks are symbols. They represent actual content for external observers, like the designer or the user of classic artificial intelligence programs, but not for the programs themselves. Other problems of symbolic accounts of cognition were raised by Harnad [44] and Clancey [45].

Barsalou [17] also recently proposed a perceptual approach to higher functions of cognition. It is argued, based on experimental facts, that even abstract reasoning does not need the assumption of an amodal symbol system, but that all concepts, even the most abstract ones, are ultimately grounded in perceptual and motor mechanisms. This is also supported by brain imagery studies of the activations produced by words representing objects or animals [46]. Indurkha [47] proposed a detailed model of the grounding of concepts in sensorimotor input that explains the understanding of metaphors, models and analogies. Some other recent studies focused on the embodied nature of mathematical concepts [48, 49, 50, 51].

Bickhard [43] proposes, with philosophical arguments, that genuine representation that overcomes the problems of symbolism can emerge in goal directed embodied cognitive agents that are in dynamical interaction with the environment that is represented, via sensors and effectors. This interactivist proposal converges with an independent trend in robotics, motivated by the problems of the classical approaches, which considers that intelligent behaviour can be generated without explicit (symbolic) reasoning systems [52], with Varela's enactivist proposal [53], based on philosophical and biological arguments, and with O'Regan's [54] proposal suggested by experiments in human vision and sensorimotor adaptation. It seems then that research in artificial life (a-life), a research field that builds models of agents inspired by biological reality, embodied either as robots that interact with the real world or in simulated environments [55, 56, 57] and usually controlled by neural networks may be the way to get genuine artificial intelligence [58]. More developed reviews of the evolution of constructivist and embodied theories of cognition and their recent convergence can be found in [59] and [60].

3 What about Romania?

In Romania, psychology and related fields were banished by the former communist regime. After the fall of communism, the newly-established national Cognitive Science Association [61] has organized several summer schools to promote cognitive science. Some psychology university departments continue to actively promote a cognitive approach to psychology while some philosophy departments are also interested in issues relevant to cognition. However, research in cognitive neuroscience seems not to be developed due to lack of expensive equipment needed for neuronal recordings or brain imagery.

Research in artificial intelligence, including neural networks and genetic algorithms, has been conducted in computer science and engineering departments, which have been traditionally competitive in Romania. However, these studies were usually concerned with engineering or formal approaches, paying little attention to results from cognitive science or neuroscience. More generally, the lack of flexibility of the Romanian academic system, that discourages rather than promotes interdisciplinary research, and the relatively scarce incentives for internationally competitive research has hindered interdisciplinarity, a definitory element of cognitive science. Moreover, the government does not promote cognitive science research, which is not mentioned at all in the programs of the governmental department concerned with the research [62], in harsh contrast with international efforts in support of cognitive studies and neuroscience (e.g., [63]). This situation has obviously lead to underdevelopment of cognitive science in Romania.

However, out of the three research areas that currently seem to be the most promising in the light of their potential applications - artificial intelligence, genomics / proteomics and nanotechnology, only the former may be conducted without expensive laboratory equipment. Computers, internet access to information and specialty journals, eventually some cheap sensors and motors to test real implementations in robots, is all that is needed besides sustained work and creativity. Computer skills are not lacking in Romania, which is a big advantage in comparison to the situation from many western countries. The domain may also benefit from people trained in physics, as many modeling methods in computational neuroscience and neural networks are borrowed from physics.

The information technology (IT) sector is already a major industry in Romania, generating relatively important revenues and profits. The government declaratively supports its development. However, with a few notable exceptions, most of the current revenues come from outsourcing activities performed for western companies. This situation, although temporary beneficial, may yield long term problems generated by lack of capitalisation of intelectual property and market share, like overdependence, unpredictibility and instability [64, 65]. Innovative applications generated by computational intelligence, such as real-world robotics, object recognition, and information search and production techniques may provide means to overcome this situation. With relatively little seed capital for development combined with professional marketing and management, these applications may yield products with high profit margins that could be marketed internationally without the need for expensive infrastructure. Research in cognitive science may yield not only computational applications, but also applications in domains like new educational methods and principles, user interface design, improvement of human factors or the treatment of neurologically impaired patients. Paradoxically, for a country without many financial resources, plenty of capital for high potential companies is available in Romania, but the venture capitalists are not currently able to find enough competitive and well managed projects [66, 67]. The lack of qualified commercial and management specialists in IT is indeed a problem that could be overcome on short term through intensive training abroad, and on long term by improvements in the quality of business studies in Romania.

We have seen that straightforward modeling in traditional artificial intelligence has fundamental limits. Results from the other areas of cognitive science, such as neuroscience and psychology, should be considered in detail within an interdisciplinary framework. Because computational intelligence can lead to highly profitable applications with relatively little investment, we consider that cognitive science development in Romania deserves better support, both through promotion of interdisciplinarity among scientists, as well as financial or administrative support from government or private companies.

4 Conclusion

Cognitive science is a diverse and dynamic research field, where multiple paradigms interact in interdisciplinary synergy. Computational approaches to the modeling and study of cognitive functions have led to several research directions, such as classical artificial intelligence and robotics, connectionism, artificial neural networks, computational neuroscience, a-life or the emerging computational neuropsychology [68]. Development of computational intelligence research in Romania, in close connection with cognitive science as an interdisciplinary field of research, should be supported as one way to enhance the country's international scientific and economic competitiveness.

5 Further readings

5.1 General cognitive science

A history of the early days of cognitive science is presented in [69]. Introductions to cognitive science are [70, 71] (though somehow biased by the symbolicist computational perspective on cognition). [60] is a textbook of embodied cognitive science. A comprehensive review of all the domains of cognitive science is [72], also available online. About 250 cognitive science books from MIT Press are available online [73] for a relatively small fee. Some relevant journals are [Behavioural and Brain Sciences](#)¹, [Cognitive Science](#), [Trends in Cognitive Science](#), [Cognition](#)². [CogPrints](#) is an open e-print archive for cognitive science.

5.2 Neural networks, computational neuroscience

[23, 74] are collections of the papers that established the field of neural networks or significant directions of research within the field. [22, 75, 76, 77] are introductions to classical artificial neural networks. [25] is a comprehensive review of the field of neural networks. [40] is a discussion of the foundational ideas of computational neuroscience. [33] discusses firing rate neural network models that agree with biological data. [78] exposes the neural network models inspired from statistical physics and attractor networks as memories. An introduction to spiking models of neurons may be found in [36]. Quantitative descriptions of the structure of the cortex can be found in [79]. [80] also presents descriptions and models of the cortex. [81] is a collection of recent results from computational neuroscience. The main journals of the field are [Neural Computation](#), [Neurocomputing](#), [Neural Networks](#).

5.3 Cognitive neuroscience

Michael Gazzaniga has edited a series of comprehensive reviews of cognitive neuroscience: [82, 83, 84, 85, 86]. Other specialized references are [87, 88, 89]. Relevant journals are [Journal of Cognitive Neuroscience](#), [Learning and Memory](#)³.

5.4 Neuroscience

Comprehensive, up-to-date information from neuroscience may be found in [34] and [90]. Relevant journals are [Journal of Neuroscience](#)³, [Neuroscience](#), [Nature Neuroscience](#), [Nature Reviews Neuroscience](#), [Annual Review of Neuroscience](#), [Brain](#)⁴, [Cerebral Cortex](#), [Trends in Neurosciences](#), [PNAS](#)³.

¹ Target papers are available online free of charge

² Brief articles are available online free of charge.

³ Articles older than one year are available online free of charge.

⁴ Full text is available online free of charge.

5.5 A-life

[91] offers a comprehensive online collection of a-life resources. [55, 56, 57, 92] are collections of papers from a-life conferences. Some relevant journals are [Artificial Life](#), [Artificial Life and Robotics](#), [Adaptive Behaviour](#).

References

1. Cognitive Science Society (home page). <http://www.cognitivesciencesociety.org/>, 2001.
2. Jerry Fodor. *The Language of Thought*. Crowell, New York, 1975.
3. Zenon Pylyshyn. Computation and cognition: Issues in the foundation of cognitive science. *Behavioral and Brain Sciences*, 3, pp. 111-132, 1980.
4. Herbert A. Simon and Craig A. Kaplan. *Foundations of cognitive science*, p. 40. In Posner [71], 1989.
5. J. R. Anderson. *The adaptive character of thought*. Erlbaum, Hillsdale, NJ, 1993.
6. A. Newell. *Unified theories of cognition*. Harvard University Press, Cambridge, MA, 1990.
7. A. Newell and H. A. Simon. *Human problem solving*. Prentice - Hall, Englewood Cliffs, NJ, 1972.
8. I. Biederman. Recognition by components: A theory of human image understanding. *Psychological review*, 94, pp. 115-147, 1987.
9. Doina Tătar. *Inteligența artificială: demonstrarea automată a teoremelor; prelucrarea limbajului natural [Artificial intelligence: automatical theorem proving; natural language processing]*. Editura Albastră, Cluj-Napoca, Romania, 2001.
10. Edward H. Shortliffe. *Computer-Based Medical Consultations: MYCIN*. American Elsevier, New York, NY, 1976.
11. John McCarthy. Some expert systems need common sense. *Annals of the New York Academy of Sciences*, 426; Computer culture: The scientific, intellectual and social impact of the computer, 1984. Available from: <http://www-formal.stanford.edu/jmc/someneed.html>.
12. IBM Corporation, Deep Blue home page. <http://www.chess.ibm.com/>, 2001.
13. Susan L. Epstein. For the right reasons: The FORR architecture for learning in a skill domain. *Cognitive science*, 18, pp. 479-511, 1994.
14. Babelfish, automatic translator. <http://babelfish.altavista.com/>, 2001.
15. List of chat bots. <http://www.botspot.com/search/s-chat.htm>, 2001.
16. Cyc, corporate home page. <http://www.cyc.com/>, 2001.
17. L. W. Barsalou. Perceptual symbol systems. *Behavioural and Brain Sciences*, 2, pp. 577-660, 1999. Available from: <http://www.bbsonline.org/Preprints/OldArchive/bbs.barsalou.html>.
18. Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, pp. 115-133, 1943.
19. F. Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65, pp. 386-408, 1958.
20. Bernard Widrow and Marcian E. Hoff. Adaptive switching circuits. *1960 IRE WESCON Convention Record, New York*, pp. 96-104, 1960.
21. Marvin Minsky and Seymour Papert. *Perceptrons*. MIT Press, Cambridge, MA, 1969.
22. John Hertz, Anders Krogh, and Richard G. Palmer. *Introduction to the theory of neural computation*. Perseus Books, Cambridge, MA, 1991.
23. James A. Anderson and Edward Rosenfeld, editors. *Neurocomputing: Foundations of research*. MIT Press, Cambridge, MA, 1988.
24. D. E. Rumelhart and J. L. McClelland, editors. *Parallel distributed processing: Explorations in the microstructures of cognition*. MIT Press, Cambridge, MA, 1986.
25. M. Arbib, editor. *Handbook of brain theory and neural networks*. MIT Press, Cambridge, MA, 1998.
26. P. McLeod, K. Plunkett, and E. T. Rolls. *Introduction to connectionist modelling of cognitive processes*. Oxford University Press, Oxford, UK, 1998.
27. Morten H. Christiansen and Nick Chater. Connectionist natural language processing: The state of the art. *Cognitive Science*, 23 (4), pp. 417-437, 1999.
28. T. Sejnowski and C. Rosenberg. Parallel networks that learn to pronounce English text. *Complex Systems*, 1, pp. 145-168, 1987.

29. Michael S. C. Thomas and Anne Karmiloff-Smith. Connectionist models of development, atypical development and individual differences. In R. J. Sternberg, J. Lautrey, and T. Lubart, editors, *Models of Intelligence for the Next Millennium*. American Psychological Association, 2001.
30. Duncan Graham-Rowe. Crash course. *New Scientist* 4, 1999.
Available from: <http://www.newscientist.com/hottopics/ai/crashcourse.jsp>.
31. Robokoneko project home page. <http://www.cs.usu.edu/~degaris/robokoneko/>, 2001.
32. D. O. Hebb. *The organization of behavior*. Wiley, New York, NY, 1949.
33. Edmund T. Rolls and Alessandro Treves. *Neural networks and brain function*. Oxford University Press, Oxford, UK, 1998.
34. R. Kandel, J. H. Schwartz, and T. M. Jessell, editors. *Principles of neural science*. McGraw-Hill, 2000.
35. S. Thorpe, D. Fize, and C. Marlot. Speed of processing in the human visual system. *Nature*, 381:520-522, 1996.
36. Wolfgang Maas and Christopher M. Bishop, editors. *Pulsed neural networks*. MIT Press, Cambridge, MA, 1999.
37. F. Rieke, D. Warland, R. de Ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the neural code*. MIT Press, Cambridge, MA, 1996.
38. C. Koch and I. Segev. *Methods in neuronal modeling: From synapses to networks*. MIT Press, Cambridge, MA, 1997.
39. Christof Koch and Idan Segev. The role of single neurons in information processing. *Nature Neuroscience*, 3, pp. 1171-1177, 2000.
40. Patricia Churchland and Terrence J. Sejnowski. *The computational brain*. MIT Press, Cambridge, MA, 1994.
41. R. Bogacz, M.W. Brown, and C. Giraud-Carrier. Model of familiarity discrimination in the perirhinal cortex. *Journal of Computational Neuroscience*, 10, pp. 5-23, 2001.
See also: <http://www.cs.bris.ac.uk/~bogacz/bbc/>.
42. P. C. Bressloff, J. D. Cowan, M. Golubitsky, P. J. Thomas, and M. Wiener. What geometric visual hallucinations tell us about the visual cortex. *Neural Computation*, 13, 2001.
See also: <http://www.math.utah.edu/~bressloff/publications/01-3abs.html>.
43. M.H. Bickhard. Representational content in humans and machines. *Journal of Experimental and Theoretical Artificial Intelligence*, 5, pp. 285-333, 1993.
Available from: <http://www.lehigh.edu/~mhb0/repcnpage.html> ;
See also: <http://www.lehigh.edu/~mhb0/pubspage.html>.
44. Stevan Harnad. The symbol grounding problem. *Physica D*, 42, pp. 335-346, 1990.
Available from: http://cogsci.soton.ac.uk/harnad/Papers/Harnad/harnad90_sgproblem.html.
45. W. J. Clancey. *Situated cognition - On human knowledge and computer representations*. Cambridge University Press, Cambridge, UK, 1997.
46. F. Pulvermuller. Words in the brain's language. *Behavioural and Brain Sciences*, Submitted, 2001.
Available from: <http://www.bbsonline.org/Preprints/OldArchive/bbs.pulvermueller.html>.
47. B. Indurkha. *Metaphor and cognition: an interactionist approach*. Kluwer Academic Publishers, Dordrecht, the Netherlands, 1992.
48. R. Nunez and G. Lakoff. What did Weierstrass really define? *Mathematical Cognition*, 4, p. 2, 1998.
49. G. Lakoff and R. Nunez. The metaphorical structure of mathematics: sketching out cognitive foundations for a mind-based mathematics. In Lyn English, editor, *Mathematical Reasoning: Analogies, Metaphors and Images*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1999.
50. S. Dehaene. *The number sense: How the mind creates mathematics*. Oxford University Press, Oxford, UK, 1997.
51. S. Dehaene, E. Spelke, P. Pinel, R. Stănescu, and S. Tsivkin. Sources of mathematical thinking: Behavioral and brain-imaging evidence. *Science*, 284, pp. 970-974, 1999.
52. Rodney A. Brooks. *Cambrian intelligence: The early history of the new AI*. MIT Press, Cambridge, MA, 1999. Partially available from: <http://www.ai.mit.edu/people/brooks/papers.html>.
53. Francisco J. Varela, Evan Thompson, and Eleanor Rosch. *The embodied mind: Cognitive science and human experience*. MIT Press, Cambridge, MA, 1992.
54. J.K. O'Regan and A. Noe. A sensorimotor account of vision and visual consciousness. *Behavioural and Brain Sciences*, 24, p. 5, 2001. Available from: <http://www.bbsonline.org/Preprints/ORegan/>.

55. Jean-Arcady Meyer and Stewart W. Wilson, editors. *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, MA, 1991.
56. Jean-Arcady Meyer, Herbert L. Roitblat, and Stewart W. Wilson, editors. *From animals to animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, MA, 1993.
57. Dave Cliff, Philip Husbands, Jean-Arcady Meyer, and Stewart W. Wilson, editors. *From animals to animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. MIT Press, Cambridge, MA, 1994.
58. Luc Steels and Rodney Brooks, editors. *The artificial life route to artificial intelligence: Building embodied, situated agents*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1995.
59. Tom Ziemke. The construction of 'reality' in the robot: Constructivist perspectives on situated artificial intelligence and adaptive robotics. *Foundations of Science*, 6, pp. 163-233, 2001. Available from: <http://researchindex.com/ziemke00construction.html>, <http://www.ida.his.se/ida/~tom/>.
60. Rolf Pfeifer and Christian Scheier. *Understanding intelligence*. MIT Press, Cambridge, MA, 1999. See also: <http://www.ifi.unizh.ch/~pfeifer/mitbook/>.
61. ASCR home page (Asociația de Științe Cognitive din România [Romanian Cognitive Science Association]). <http://www.psychology.ro/Psiho/ascr.htm>, 2001.
62. Ministry of Education and Research, Romania, department of research; home page. <http://www.mct.ro/>, 2001.
63. Cognitique, a program of Ministry of Research, France; home page. <http://www.recherche.gouv.fr/recherche/aci/cognib.htm>, 2001.
64. Ziarul financiar. România poate trage o carte bună în competiția outsourcing [Romania can get a good card in the outsourcing competition]. *Ziarul financiar*, 720, p. 12, 2001.
65. Ziarul financiar. Când america suspină, nici CornerSoft Technologies nu se simte prea bine [When America sighs, CornerSoft Technologies doesn't feel well, neither]. *Ziarul financiar*, 728, p. 11, 2001.
66. Răzvan Voican. Capitalul de risc, atras de creșterea economică rapidă [Venture capital, attracted by rapid economic growth]. *Ziarul financiar*, 693, p. 1, 2001.
67. Laurențiu Ispir and Emil Lazăr. Românii nu i-au scos pe investitorii străini din poziția de așteptare [Romanians didn't get foreign investors out of a waiting stance]. *Ziarul financiar*, 656, p. 8, 2001.
68. NIPS*2001 workshop on computational neuropsychology home page. <http://www.cs.colorado.edu/~mozer/nips2001workshop.html>, 2001.
69. H. Gardner. *The mind's new science: A history of the cognitive revolution*. Basic Books, New York, NY, 1985.
70. Neil A. Stillings, Steven E. Weisler, Christopher H. Chase, Mark H. Feinstein, Jay L. Garfield, and Edwina L. Rissland. *Cognitive science: An introduction*. MIT Press, Cambridge, MA, 1995.
71. Michael I. Posner, editor. *Foundations of cognitive science*. MIT Press, Cambridge, MA, 1989.
72. Robert A. Wilson and Frank Keil, editors. *The MIT encyclopedia of cognitive sciences*. MIT Press, Cambridge, Massachusetts, 1999. Available from: <http://cognet.mit.edu/MITECS/>.
73. CogNet home page. <http://cognet.mit.edu/>, 2001.
74. James A. Anderson, Andras Pellionisz, and Edward Rosenfeld, editors. *Neurocomputing 2: Directions of research*. MIT Press, Cambridge, MA, 1993.
75. Simon S. Haykin. *Neural networks: A comprehensive foundation*. Prentice Hall, 1998.
76. Dan Dumitrescu. *Modele conexioniste în inteligența artificială [Connectionist models in artificial intelligence]*. Universitatea Babeș-Bolyai, Cluj-Napoca, Romania, 1995.
77. D. Dumitrescu and H. Costin. *Rețele neuronale [Neural networks]*. Teora, Bucharest, Romania, 1996.
78. Daniel Amit. *Modeling brain function: The world of attractor neural networks*. Cambridge University Press, Cambridge, UK, 1992.
79. V. Braitenberg and A. Schuz. *Anatomy of the cortex: Statistics and geometry*. Springer Verlag, Berlin, 1991.
80. Moshe Abeles. *Corticonics: Neural circuits of the cerebral cortex*. Cambridge University Press, Cambridge, UK, 1991.
81. Lawrence Abbott and Terrence J. Sejnowski, editors. *Neural codes and distributed representations: Foundations of neural computation*. MIT Press, Cambridge, MA, 1999.
82. Michael S. Gazzaniga, editor. *Handbook of cognitive neuroscience*. Plenum, 1984.
83. Michael S. Gazzaniga, editor. *The new cognitive neurosciences*. MIT Press, Cambridge, MA, 2000.

84. Michael S. Gazzaniga, Richard B. Ivry, and George Mangun. *Cognitive neuroscience: The biology of the mind*. W W Norton & Co, 1998.
85. Michael S. Gazzaniga, editor. *Cognitive neuroscience: A reader*. Blackwell, 2000.
86. Michael S. Gazzaniga, editor. *Conversations in the cognitive neurosciences*. MIT Press, Cambridge, MA, 1996.
87. Martha J. Farah. *The cognitive neuroscience of vision*. Blackwell, 2000.
88. Marc Jeannerod. *The cognitive neuroscience of action*. Blackwell, 1999.
89. Joaquin M. Fuster. *Memory in the cerebral cortex: An empirical approach to neural networks in the human and nonhuman primate*. MIT Press, Cambridge, MA, 1995.
90. M. J. Zigmond, F. E. Bloom, S. C. Landis, J. L. Roberts, and L. R. Squire, editors. *Fundamental neuroscience*. Academic Press, 1998.
91. alife.org. <http://www.alife.org/>, 2001.
92. Rodney Brooks and Pattie Maes, editors. *Artificial life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, Cambridge, MA, 1994.