Yet to be defined
Promoting Evolutionary Robotics as a Method in Cognitive Science

Marieke Rohde (m.rohde@sussex.ac.uk)
CCNR — Centre for Computational Neuroscience and Robotics
University of Sussex, UK

Abstract

Within the Cognitive Science community, voices claiming the importance of agent–environment interaction for cognition become louder and louder. This paper promotes evolutionary robotics, a technique for the automated generation of robots to exhibit intelligent behaviours, as a method for cognitive science that meets these demands for considering embodiment and situatedness of cognitive agents. Analysing research examples, the explanatory potential of evolutionary robotics is discussed and defeated against common prejudices.

Introduction

With the cognitive revolution in the 1950s, the study of the mind was liberated from the restrictive framework of radical behaviourism - the invention of the digital computer gave rise to the powerful metaphor of the mind as a computer: as formal languages could adequately describe the processes giving rise to intelligent behaviour in computers without relying on introspection, why would it not be possible to apply these same language to explain the processes inside a person or an animal, that give rise to its intelligent behaviour?

50 years down the road, we owe a lot of our understanding of cognition to research conducted under this paradigm. 50 years down the road, we have also realised difficulties associated with the methodology evolved from the mind as computer metaphor: Inference mechanisms relying on merely syntactic symbol manipulation have problems of discarding irrelevant information from the inference process, which leads to an explosion of computational complexity, as the symbol systems they work on become more complex. This problem is known as the frame problem [30]. The philosophical variant of the frame problem was termed the “problem of abduction” by Fodor [11], who notices that “there is typically no way to delimit a priori the considerations that may be relevant”[11], p. 37). He describes himself as “worried half to death” ([11] p. 39) and claims that “much of the field is in deep denial” ([11] p. 39) about the consequent limitations of the classical computational theory of mind.

In praxis, by constraining a symbol manipulation process with domain specific syntactic rules, those problems can be bypassed to a certain degree. But building what is relevant into an agent is somewhat unsatisfactory, when trying to understand *why we know what is relevant*, to explain our amazing capacity to *make sense* of the phletora of unpredictable situations we are confronted with. The problem of meaning in the classical computational paradigm is more fundamental, as adequately pointed out by Harndad: Elaborating on Searle’s famous Chinese room thought experiment[31], he formulated the symbol grounding problem [14]: given the arbitrariness of symbols, how can a closed symbol system ever *know* what it is, that it is doing? Only during interaction with the real world, a process will have the potential of being meaningful. To tackle questions of meaning and relevance, it is necessary to think outside the “cognitive box” depicted in Fig. 1 (B), considering both the given body and the given environment in which cognitive behaviour occurs.

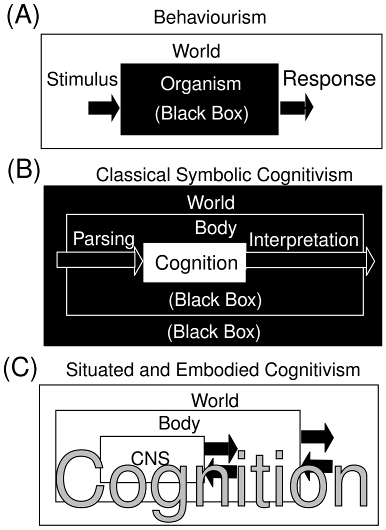

Figure 1: (A) The organism as a black box in behaviourism. (B) The body and environment as black boxes in symbolic cognitivism. (C) Embodied and Situated Cognitivism.

There are two principal approaches to such an enterprise:

- A hybrid approach, in which a symbol manipulation process is part of an embodied and situated agent with a defined sensorimotor system, which generates symbols for the computational process to work on and interprets the output symbols it generates. [1]

- An approach that does not feature a syntactic symbol manipulation device at all.

In both variants, meaning is a property emerging from the agent–environment interaction. The first variant relies on the hypothesis that perception, action and computational cognition are sufficiently separate processes to allow a functional and operational division. The alternative hypothesis, however, is that cognition, action and perception are tightly interlinked, and that no decoupled homuncular symbol manipulation process is part of the generation of cognitive behaviour at any stage. From such a viewpoint cognition by necessity becomes a process that extends the boundaries of brain and body (see Fig. 1, (C)).

Many cognitive scientists (most prominently Fodor and Pylyshyn [10]) are sceptical towards the second alternative, for they see a contradiction between the fact that humans exhibit behaviour that follows syntactic rules (e.g. language or logical reasoning) and the hypothesis that no symbol manipulation process is generating this behaviour. Denying such systematicities would be completely unreasonable, the great achievements of symbolic cognitive science give us an idea of the extent to which human behaviour follows syntactic regularities. However, behaviour being syntactic does not imply behaviour *just* being syntactic. Possibly, the symbols that we use do not stand in an arbitrary relation to their meaning in the way the symbols computers use do[2]. The use of symbols and representations is one of our most fascinating cognitive capacities, and very little understood. Explaining the nature of this capacity, however, will hardly be possible from within the symbolic cognitivist paradigm, as symbols are the explanans rather than the explanandum in this approach (I. Harvey, personal communication, March 2005).

This paper argues that evolutionary robotics is a methodological framework that allows an investigation of cognition that takes the embodiment and situatedness of a cognitive agent into consideration and thereby does not come short with respect to semantic grounding of modelled behaviour. Evolutionary robotics is neither relying on presuppositions on the separability of action, perception and cognition or the plainly syntactic character of internal representations. Research examples to illustrate the explanatory potential of this method are presented and issues in and possible objections to this novel framework for the study of cognitive behaviour are discussed.

Evolutionary Robotics

In their 2000 book ”Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines.”, Nolfi and Floreano define the subject of their publication as ”a new technique for automatic creation of autonomous robots [...] inspired by the darwinian [sic!] principle of selective reproduction of the fittest.” ([22], preface). A “fitness” measure is defined to rate the behaviour of real or simulated robots from an initially random population on a given task. Those robots performing best according to this measure produce “offspring” that is slightly changed (“mutated”) to be evaluated in the next generation. By iteration of this process, an optimisation search on the specified fitness function is performed.

Closing the Sensorimotor Loop

With respect to the preceding discussion, the important property of experiments in evolutionary robotics is the autonomous interaction of the evolving agents with their environment: What is rated is not the input or output of a computational process, but actual behaviour within a real or simulated environment. Cognitive properties emerging from the interaction of an agent with its environment, which cannot be located within the cognitive box of classical computational cognitive science, can be observed, analysed and explained in experiments in evolutionary robotics. For instance, Floreano et. Al. [9] have coevolved active vision and feature selection in robots to exhibit categorical perception, car driving and navigation behaviour: the remarkable point in these experiments is that the robots’ movement and change of attentional focus have a direct impact on the visual inputs, which again crucially impacts on the way visual information is exploited by the robot to achieve the required task. The external closure of the sensorimotor loop in these experiments is absolutely essential to realise and understand the cognitive behaviour exhibited by the evolved robots. Evidence that similar interactions of perception and action mediated through the environment take part in human and animal cognition comes e.g. from studies on how sensory adaptation in cats and in human subjects stands or falls on the possibility to interact with the environment [16].

Dynamical Explanations

If, as outlined in the previous section, cognition relies on real time interaction with the environment, dynamical systems theory (DST) is self–suggesting as mathematical framework for description and analysis. Every system that evolves in time can be described as dynamical system. Therefore, DST is a mathematical language in which an agent, its environment or even individual parts of the agent (e.g. its brain) can be described unitarily. Furthermore, the spatiotemporal coordination of behaviour is often essential for cognitive function in situated and embodied agents. Where classical architectures struggle to adequately describe and generate real time behaviour, e.g. to synchronise the time of inference and the time of the inferred [34], in DST, spatiotemporal variables are naturally and precisely accounted for. Ports and van Gelder [26] give an overview over the various ways and domains in which DST theory has been used to explain cognitive behaviour. A very detailed dynamical explanation of shape discrimination behaviour in an evolved robot can be found in [2].

---

[1] This solution was suggested by Harnad in [14].

[2] A very interesting renunciation of the classical concept of the Saussurian arbitrary sign/symbol is presented in Weber’s biosemiotic theory [35].

Robots are typically controlled by digital computers, which, on the first glance seems to contradict the claim that their cognitive behaviour is not digital computation. This apparant contradiction seems even stronger for the case of computer simulation experiments. As discussed in van Gelder [12], it is important, not to confuse the abstract mathematical model and its realisation. Many processes, including insect population growth or flight trajectories of thrown rocks, can be simulated on digital computers. Drawing conclusions on their symbol processing properties from this fact, however, seems not reasonable.

"Cognition" Comes from Knowledge

"The word *cognitive* refers to perceiving and knowing" ([33], p. 1, orig. emph.), two concepts that we are familiar with from our own practice and conscious mental experience. The early cognitive science models were generally guided by such conscious introspection. But it soon became clear that conscious introspection will not give us the whole story; Damasio's work on somatic markers that unconsciously influence human decision making [7] is just one example of evidence for a cognitive mechanism that works unconsciously, and cannot even be localised in the brain. A clarification of what we mean by "cognitive" is inevitable: are animals cognitive? Are computers cognitive? Is balancing a cognitive capacity?

Experiments in evolutionary or behaviour based robotics frequently feature a high degree of idealisation to keep the models tractable and understandable. This practice has sometimes led cognitive scientists to the judgement that the behaviour they exhibit is not cognitive (e.g. Grush and Clark [4], Kirsh [17]). Both these cited criticisms come to the conclusion that the ultimate benchmark for cognitive behaviour is the use of internal representations. The problem with this criterion is that outside a symbolic cognitivist framework, "using internal representations" is not a well defined concept. If representation is meant as symbolic representation, such a claim is simply and inversion of the symbolic cognitivist program: from "cognition as computation" to "computation as cognition". A characterisation of a research area in terms of a research method seems little helpful and a priori biased against other explanatory frameworks. If, on the other hand, internal representation is not meant in such a restricted sense, one illdefined criterion is just substituted for by another. As mentioned before, the representational capacities of animals, let alone humans, are poorly understood and seem to be rather an area of scientific inquiry within cognitive science than a practical criterion to tell cognitive behaviour from non–cognitive behaviour. A number of tasks have been claimed to be "representation hungry"[5], but could be explained with simple dynamical models that did not make use of internal representations in a straight forward way (see e.g. [2], or the section "The A Priori Bias" of this paper). Such surprises illustrate how unreliable our intuitions about internal representations are[3].

Another common way to define cognition is with reference to rationality as some form of "doing the right thing" ([30], p.4), which opens yet another chapter: What is the right thing? As for living organisms, their behaviour shows what is right or wrong to them, we know it by mere virtue of being living organisms ourselves, be it a person banging her toe or the sithering of a plant (compare Weber [35], p. 118). As for chess playing computers or vacuum cleaning robots, we hesitate to make judgements about right and wrong for them. To them it is insignificant whether the floor is clean or the game is won, if we would have programmed them to lose games or to spread dirt in a room, it would not make a difference to them. Personally speaking, I think that the claim that "[programs], like people, are inherently goal oriented." ([6] p.3.) is plainly wrong, which may be taken as a general scepticism against strong AI. To claim that it is a program's goal to win a game of chess appears as little reasonable as claiming that it is a tsunami's goal to devastate a coastline, as both lack genuine subjectivity or intentionality. It is because of this monopoly on genuine rationality and subjectivity, that a useful definition of cognition and rationality will have to be rooted in biology. Candidate theories are e.g. the theory of autopoiesis[20], in which all that is done by living systems is conceived of as cognitive, or Moreno's et Al. theory of cognition as autonomous sensorimotor control that is relatively decoupled from metabolism [21]. This paper does not put forward a particular concept of cognition, but rather stresses that a useful criterion should be defined in the realm of the science of the living, not in the realm of descriptive formalisms. On the basis of such a criterion, it can then be assessed whether a computational model is indeed modelling cognition. With respect to the supposed simplicity of behaviour based robotics models, the mentioned evolved robot that discriminates objects using motion, vision and attentional focus [9], the evolution of sequential path finding behaviour in a maze [23] or the evolution of cooperative behaviour in a multi robot system to achieve a common goal [27] are just some examples of experiments that should qualify as models of cognitive behaviour according to all reasonable definitions of cognition.

The scepticism about genuine purposes and goal orientation outlined above applies to embodied and situated cognitivism in the same way as it does to symbolic cognitivism: in both cases, the goals and purposes of behaviour are specified by the experimenter and meaningless to the artificial agent. However, an advantage that evolutionary robotics has with respect to modelling "pseudo goal oriented" behaviour is that the external value criterion is not explicitly built in either the agent or its environment; it is not part of the sensorimotor loop. Usually, in symbolic cognitivism, and indeed in some situated and embodied models of cognition as well,

---

[3] Situated, embodied and dynamical approaches to cognition are sometimes labelled "antirepresentationalist". This is not justified. The concept of "representation" is not a primitive within this explanatory framework, but that does not imply a refutation of the concept. In contrary, it opens up entirely new ways of using and refining it. Pasemann [25], for instance, explores the concept of representation in a very inspiring way.

value systems are modules built in the cognitive architecture (e.g. evaluation functions in machine learning, termination criteria in search). Such modules come with an a priori semantics of what is good and what is bad; Rutkowska [29] criticises this approach, as it pushes off a lot of the explanatory burden to evolution, the supposed designer of the value system, and, more crucially, as "[increased] flexibility requires some more general purpose style of value" ([29], p. 292). A 'ghost in the machine' is very vulnerable to change of the perceptual apparatus, at least if changes affect those channels that guide its judgement (e.g. reward channels). It is very questionable to what extend open ended developmental processes or sensory adaptation, such as the stunning human capacity to compensate for the loss of vision by tactile sensation [1] or to adapt to inversion of the visual field [18] can be explained with models that feature a localised value system with a simplistic a priori semantics. *Reinterpretation* of environmental input seems to require a much more complex and informed value system. In experiments in evolutionary robotics, a value system to extract information relevant to the task from the sensory, motor and internal state of the agent self–organises over the course of evolution. The form such an evolved value system takes is not predetermined, it can potentially make use of aspects of agent environment interaction that a human designer could probably not have come up with. For instance, Di Paolo [8] presents experimental results in which evolved robots do not feature a localised value system to steer their light seeking behaviour. Remarkable about these experiments is the robustness of the agents against sensor swapping, even though during the course of evolution, no such disruptions ocurred: the way the robots "judge" on the direction of light does not rely on the activation of individual light sensors, but on a more abstract and distributed system judging whether the behaviour is "successful". Possibly, those kinds of experiments are the first steps towards understanding of the cognitive capacity to make sense of the world around us.

The A Priori Bias

The property of being less biased about the forms cognitive mechanisms can take has now been mentioned several times in this paper, and is frequently pointed out as one of the most striking advantages of situated and embodied approaches in general and evolutionary robotics in particular (e.g. in [15, 22, 3]) over traditional computational models. However, it is impossible to design experiments that are completely devoid of a priori assumptions, and maybe, it is not even desirable.

Evolving agents is not putting cogs, legs, wheels and neurons in a bag and shaking it, such that an intelligent robot falls out. Instead, artificial evolution *parametrises a predefined set of possible agents.* The art of evolutionary robotics is to appropriately predefine this set, to clamp certain dimensions and leave the others free to evolve, and to phrase the right questions to be answered with a particular experiment. As a rule of the thumb, when investigating a certain cognitive capacity, *minimising the a priori assumptions about this capacity* is most likely to exploit the invaluable potential of evolutionary robotics to surprise you and cause you to refine your preconceptions. E.g. if the interest is in finding out the role of representation, a supposedly representation hungry task should be investigated without providing the robot with a representational system that biases the evolved solution to come up with a solution you expect. Di Paolo (personal communication, April 2005), for instance, evolved robots that had no sensors apart from a velocity sensor to stop in the middle of a variable length track. Even though very simplistic, this problem is representation hungry, as it requires some sort of memory: the agent first has to find out about the length of the track before it "knows" where to stop. The behaviour of the evolved agents is to run against one wall, bounce off, run against the other wall and then to go the middle of track and stop. The interesting finding here is that in the evolved neural control architecture, no variable correlates with track length. The only variable that correlates with track length is the physical velocity at the moment it bounces for the second time, i.e. an external, embodied variable that depends on the interactive history of the robot and its environment. This indicator is sufficient for the robot to apply constant deceleration after the second bounce, to then come to a stop in the middle of the corridor. The measuring and memorising of the track length, in this case, clearly is a cognitive capacity that transcends the limits of the cognitive box of symbolic cognitivism, in a way that a human designer could probably not have come up with. As humans and animals are inherently embodied and situated, how can we know that exceptiong the body and the environment from consideration does not result in discarding the most essential components of cognitive competence? Evolutionary robotics allows us test our intuitive preconceptions on what it takes to perform a certain task. Similar surprising results are e.g. the fact that decision making does not have to be discrete process [2] or the mentioned findings by Floreano et Al. [9] on how a perceptual mechanisms can crucially depend on action.

If variation of the parameters that are set by the evolutionary search process results in variation of the property or mechanism under investigation, such challenges to our preconceptions are most likely to occur. Hence, if the scientific interest is in aspects of embodiment, evolving the morphology [32] of robots is likely to lead to interesting insights. If the question asked is about the origins of communication, it is a good idea to not predefine communication channels, but rather to let them evolve freely[27]. If the interest is on neural dynamics, it is reasonable to evolve the control architecture with a minimal a priori bias on the dynamic potential of the evolved controller[24].

It has to be noted in this context that such models have to be seen as existential proofs, not as claims about universality of the evolved dynamical principles. It is not questioned that many of the described evolved behaviours could also be realised in a different way by symbolic cognitivist models. However, it should be obvious by now that our intuitions on how certain cognitive capacities *require* a certain mechanism (e.g. internal representations) can err, and that we should open our minds towards the possibility that bodily and environmental variables provide a vital contribution to cognition in general, or at least in particular cases, even if such a contribution does not seem obvious. Human and animal behaviour is and has always been tied to an actual body that interacts with an actual environment — if there are ways of exploiting this fact, why *would* it be benefitial to develop a mind that operates in an isolated box, syntactically manipulating symbolic tokens that are impoverished in their semantics? It is reasonable to ask this question, and maybe a situated and embodied investigation of the origins of mind would actually provide us with an answer that gives us a reason why this should be the case. However, presupposing that a decoupled syntactic manipulation process is essential to cognition without being able to convincingly answer this question seems intellectually precious.

Realising the power of a reduced a priori bias among a single dimension seduces to aim at a reduction of the a priori bias in general. However, in order to be able to understand the evolved solution, it is adviceable to look for an otherwise minimalistic task, a technique that Beer has termed the exploration of "minimally cognitive behavior" [3]. Apart from the tractability of the model, there is also a technical problem associated with poorly informing evolutionary search. If less assumptions are fixed, the search space becomes more complex, which can hamper the evolutionary search process that still relies on a simple scalar performance measure. An evaluation function will not automatically detect if mutation among a certain dimension is "on the right track" if the context of the other varied dimensions suppresses the expression of this advantage, thereby rendering large parts of the search space useless in terms of the simple evaluation criterion, a problem called the "bootstrap problem" ([22], p.13). Experienced evolutionary roboticists know techniques how to deal with the bootstrap problem, e.g. incremental evolution [22], seeding evolution with center crossing neural networks [19] or modular evolution [23, 25]. A good bootstrapping technique aids evolutionary search without interfering with its potential to surprise us and challenge our preconceptions. In Rohde and Di Paolo [28], it is shown that aiding evolutionary search is not simply a matter of delimiting or complifying the search space: in a directional pointing task, minimising the degrees of freedom in which an arm can move does not only restrict the repertoire of possible behaviours to solve the task, but also aggravates the search process. On the other hand, imposing the principle of linear synergy, a linear relation between the dynamic torques applied to different joints evident in human motor behaviour [13] on the evolved networks advances the evolution of behaviourally rich solutions of high performance. Presumably, as robotic platforms and tasks evolved become more complex, the development of good bootstrapping techniques will be one of the key concerns in evolutionary robotics in the years to come. The mentioned findings suggest that informing evolutionary search with findings from biology is a promising approach for this endeavour, which, in the reverse direction, will possibly also shed light on the nature of empirically evident mechanisms.

Discussion

Hopefully, this paper succeeds in demonstrating the explanatory potential of evolutionary robotics as a method in cognitive science. Starting with a general discussion of the importance of an embodied, situated and dynamical cognitive science, evolutionary robotics was proposed as a method that meets the described challenges to the dominant methodology of symbolic cognitivism. The natural account this method provides for real time dynamical interaction of cognitive agents in a closed sensorimotor loop has been pointed out as a benefit, as well as its potential to surprise us and make us refine our preconceptions about how a cognitive mechanism works. Common scepticisms, such as whether evolutionary robotics conflicts with syntactic and representational aspects of human cognition or whether experiments in evolutionary robotics can be seen as truly cognitive models, have been discussed, to come to the conclusion that these scepticisms are not justified.

As a personal remark, I am convinced that our cognition is essentially dependent on how our physiology allows us to interact with the world. To put it in another way, I believe that if we had sonar sense instead of eyes or if our environment would not feature gravity, the way we think would be radically different. Also, I do not believe that a pure symbol manipulation process is part of our mental processing at any stage. Whatever it is we humans do to achieve our extraordinary cognitive capacities, it does not involve the generation of internal representations, deprived of their genuine meaning, so they can be manipulated and composed on the sole basis of their syntax. These convictions have lead me to adapt evolutionary robotics as a methodology for the study of mind, and I hope that I could make conceivable why evolutionary robotics is a reasonable methodology for someone holding those beliefs. I also hope that I could stimulate other cognitive scientists to have a benevolent eye on findings from evolutionary robotics, to inspire them for their own work or even to try to apply this method themselves.

As a final clarification, it should perhaps be added that the use of evolutionary search is not typically interpreted as a model of biological evolution, and the use of neural networks does not normally imply an analogy to biological brains. Evolutionary search is applied simply because it allows for a parametrisation of the predefined set of possible robots that goes beyond the cognitive limits and imaginatory capacity of a human designer, other search algorithms serving this purpose would be as suitable to fulfill that role. Neural networks are frequently used to control robots because they allow a general and unbiased representation and parametrisation of mathematical functions or dynamical systems. However, in individual cases, such analogies may be drawn.

References

[1] Bach–y–Rita, P., M. E. Tyler, and K. A Kaczmarek: *Seeing With the Brain.* Int. J. Human-Computer Interaction 15(2) 2003. 285-295.

[2] Beer, R.D.: *The Dynamics of Active Categorical Perception in an Evolved Model Agent.* Adaptive Behavior 11(4) 2003: 209-243.

[3] Beer, R.D.: *Toward the Evolution of Dynamical Neural Networks for Minimally Cognitive Behavior.* In: P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (eds.): *From Animals to Animats 4: Proceedings of SAB 1996* MIT Press Cambridge MA 1996. 421-429.

[4] Clark, A. & R. Grush: *Towards a Cognitive Robotics.* Adaptive Behavior, 7(1) 1999. 5-16.

[5] Clark, A. & Toribio, J.: *Doing without Representing?* Synthese 101 1994. 401-431.

[6] Collins, M.A. & E. E. Smith (eds.):*Readings in Cognitive Science: A Perspective from Psychology and Artificial Intelligence.* (Editorial) Morgan Kaufmann, San Mateo CA 1988. 1–.

[7] Damasio, A.R: *Descartes' Error: Emotion, Reason, and the Human Brain.* Putnam, New York 1994.

[8] Di Paolo, E.:*Organismically-inspired robotics: Homeostatic adaptation and natural teleology beyond the closed sensorimotor loop.* In: K. Murase and T. Asakura (eds.):*Dynamical systems approach to embodiment and sociality.* Adelaide, Australia: Advanced Knowledge International 2003. 19-42.

[9] Floreano, D., Kato, T., Marocco, D. and Sauser, E.: *Coevolution of Active Vision and Feature Selection.* Biological Cybernetics, 90(3) 2004. 218-228.

[10] Fodor, J. A. & Z. W. Pylyshyn: *Connectionism and Cognitive Architecture: A Critical Analysis.* in: S. Pinker and J. Mehler (eds): *Connections and Symbols.* (*Cognition* Special Issue) MIT Press, Cambridge MA 1988.

[11] Fodor, J. A.: *The Mind Doesn't Work That Way: The Scope And Limits Of Computational Psychology.* MIT Press, Cambridge MA 2001.

[12] van Gelder, T. J.: *The Dynamical Hypothesis in Cognitive Science.* Behavioral and Brain Sciences 21 1998. 1-14.

[13] Gottlieb, G. L., Q. Song, G. L. Almeida, D. Hong, and D. Corcos: *Directional Control of Planar Human Arm Movement.* Journal of Neurophysiology 78 (1997) 2985–2998.

[14] Harnad, S.: *The Symbol Grounding Problem.* Physica D, Vol. 42, 1990. pp. 335-346.

[15] Harvey, I., Di Paolo, E. A., Tuci, E., Wood, R., Quinn, M.: *Evolutionary Robotics: A New Scientific Tool for Studying Cognition.* Artificial Life, Forthcoming.

[16] Held, R.: *Plasticity in Sensory-Motor Systems.* Scientific American, 213 (5) 1965. 84-94.

[17] Kirsh, D.: *Today the Earwig, Tomorrow Man?* Artificial Intelligence, 47, 1991. 161–184.

[18] Kohler, I.: *Experiments with Goggles.* Scientific American, May 1962.

[19] Mathayomchan, B. and R.D. Beer: *Center-crossing recurrent neural networks for the evolution of rhythmic behavior.* Neural Computation 14 2002. 2043-2051.

[20] Maturana, H.R. & F.J. Varela: *Autopoiesis and Cognition: the Realization of the Living.* Reidel, Dordrecht 1980.

[21] Moreno A., J. Umerez & J. Ibañez: *Cognition and Life. The Autonomy of Cognition.* Special Issue Brain & Cognition 34 (1) 1997. 107-129

[22] Nolfi, S and D. Floreano: *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self–Organizing Machines.* MIT Press, Cambridge MA 2000.

[23] Paine, R. W. and J. Tani: *Adaptive Motor Primitive and Sequence Formation in a Hierarchical Recurrent Neural Network.* in: Schaal, S., et al. (eds.): *From Animals to Animats 8. Proceedings of the Eighth International Conference on the Simulation of Adaptive Behavior (SAB 2004).* MIT Press Cambridge, MA 2004. 274-283.

[24] Pasemann, F., U. Steinmetz, M. Hülse and B. Lara: *Robot Control and the Evolution of Modular Neurodynamics.* Theory in Biosciences, 120 2001. 311–326.

[25] Pasemann, F.: *Repräsentation ohne Repräsentation - Überlegungen zu einer Neurodynamik modularer kognitiver Systeme.* in: Rusch, G.; Schmidt, S. J.; Breidbach, O. (eds.): *Interne Repräsentationen - Neue Konzepte der Hirnforschung.* Suhrkamp, Frankfurt 1996. 42-91.

[26] Ports, R.F. & T. van Gelder (eds): *Mind as Motion. Explorations in the Dynamics of Cognition.* MIT Press, Cambridge MA 1995.

[27] Quinn, M.: *Evolving Communication without Dedicated Communication Channels.* In J. Kelemen and P. Sosík (eds.): Proceedings of ECAL 2001. 357–366.

[28] Rohde, M. & E. Di Paolo: *t for Two. Linear Synergy Advances the Evolution of Directional Pointing Behaviour.* submitted to: ECAL 2005 - VIIIth European Conference on Artificial Life 5.– 9.9.2005. Kent, UK.

[29] Rutkowska, J.: *What's value worth? Constraining Unsupervised Behaviour Acquisition.* In: Proc. of the Fourth European Conference on Artificial Life 1997. 290–298.

[30] Russel, S & P. Norvig: *Artificial Intelligence. A Modern Approach.* Prentice Hall, New Jersey 1995.

[31] Searle, J. R.: *Minds, Brains and Programs.* Behavioral and Brain Sciences 3 1980. 417-457.

[32] Sims, K.: *Evolving 3D Morphology and Behavior by Competition.* Artificial Life IV Proceedings, ed.by Brooks & Maes, MIT Press, 1994, 28-39

[33] Stillings, N.E. et Al.: Cognitive Science. An Introduction. Second Edition, MIT Press, Cambridge MA 1995.

[34] Smith, B. C. : *On the Origin of Objects.* MIT Press, Cambridge MA 1996.

[35] Weber, A.:*Natur als Bedeutung. Versuch einer Semiotischen Theorie des Lebendigen.* Königshausen & Neumann, Würzburg 2003.