The Homo Cyber Sapiens, the Robot Homonidus Intelligens, and the 'artificial life' approach to artificial intelligence.

Luc Steels
Artificial Intelligence Laboratory
Vrije Universiteit Brussel
Pleinlaan 2, B-1050 Brussels, Belgium
E-mail: steels@arti.vub.ac.be

January 30, 1995

Paper contributed to the Burda Symposium on Brain-Computer Interfaces. Muenchen, February 1995. Draft. Limited distribution.

1 Introduction

Intelligence is a capacity which has evolved incrementally in small steps during millions of years of evolution [7]. The hominid line split only 5 million years ago from the chimpanzee line. The *Homo Erectus*, which developed 1.5 million years ago, was the first to show elaborate tools, the use of fire, migration and thus the capability to deal with many different climatic variations. The last major brain size increase happened only 200,000 years ago. With it came an important evolution in the anatomy of the vocal tract necessary for speech and thus the full development of language. Before all this, many basic adaptations had evolved without which human intelligence would never have been possible: the development of complex sensory organs like eyes and ears, the development of nervous systems with steadily increasing complexity, the growth in the sophistication of communication, and so on.

The examination of the fossil record has shown that biological evolution exhibits periods of relative stability and then periods where there are sudden jumps [12]. These jumps may be due to increased pressures from climatic variations, new species invading the territory, and adaptations causing an avalanche of new capabilities that push up the complexity of a species. Also in the evolution of intelligence we find periods of (relatively) rapid progress, as in the emergence of *Homo Erectus* about 1,5 years ago or of *Homo Sapiens* about 200,000 years ago.

Another obvious conclusion from factual investigations of biological evolution is that it never stops. Species continue to evolve, adapt, change under the pressure of variations in the ecosystems or competitive developments within the species. Once the motor of complexity growth has started, it apparently does not stop. Indeed, the anthropological record shows important continuing evolutions in human intelligence towards more and more sophisticated and externalised representations [7]. The stable hierarchical Egyptian Farao culture with its emphasis on mythology and iconic communication [25] is quite different from the modern unstable dynamical societies with sophisticated scientific and technological tools based on abstract communication media.

This raises the following question: Is another major jump forward in intelligence possible? If so, would it be so substantial that it makes sense to talk about a new species? There are two possibilities. The first possibility starts from biology. The second one from technology.

Although there are no biological signs pointing in the direction of an increase in the anatomical basis of intelligence in humans or other species, there are awesome technological advances at the moment which may make it possible to extend biological capacities. The possible extensions include artificial sensory devices, electronic memory units, computer processors, and mechanical actuators. What would happen if this technology is applied to ourselves? Would it lead to a new species? This species might possibly be called the *Homo Cyber Sapiens*. Its members could start as extensions of ourselves but gradually become more independent from biological 'wetware' in order to continue their existence. Minsky [18] has suggested that this may lead to a form of immortality. These techniques could perhaps first be applied to make animals more intelligent. Primates are assumed to be intelligent to some degree and even capable of limited forms of language [9]. What would happen if they are given an artificial vocal apparatus capable of the articulation required for speech?

There is a second way towards other forms of intelligence. Efforts to build completely artificial humanoids, i.e. intelligent robots, seriously started in the late nineteen fifties and have been steadily making progress. In some optimistic views [19] robots with the capacity of human intelligence are only fifty years away. But most robot builders are much more pessimistic. Progress in artificial intelligence so far has been almost entirely in disembodied intelligence, focusing on the modeling and implementation of reasoning patterns. As I will discuss later in the paper, there remain many formidable obstacles. A new approach is now being tried which paradoxically wants to de-emphasise engineering in favor of biology. Its proponents argue that we must first build 'artificial life' before artificial intelligence becomes possible [26]. Maybe with this artificial life approach a new artificial species with human like intelligence might one day be possible. I propose to call this species the *Robot Homonidus Intelligens.*

This paper discusses briefly both lines of development from the viewpoint of somebody who has been conducting 20 years of research into artificial intelligence. It argues that both types of species may ultimately come about but are a long way off into the future. It also argues for fundamental research and experimentation to prepare this future emphasising the 'artificial life' approach to the development of intelligent humanoid robots.

2 Steps towards the *Homo Cyber Sapiens*.

The main idea behind the *Homo Cyber Sapiens* is that intelligence has been steadily evolving towards greater sophistication and power and that there is no reason to assume that this evolution has stopped. Jumps in evolution have always coincided with anatomical changes (increase in brain size and/or sensory and actuator modalities) and strong ecological pressures because they are the major driving force of natural selection.

For each significant jump forward in the evolution of human intelligence, there has been a dramatic *increase in brain size.* For example, the *Homo Sapiens* shows a 20 % volume increase compared with the *Homo Erectus.* There are no signs that the human brain is expanding (even at a slow rate), but technological extensions of the human brain might be possible in the not too distant future. There has been a trend in computer engineering towards artificial memories with ever increasing storage capacity and ever decreasing size as well as a trend towards faster and faster processors at a smaller size. There is no sign that the end is in sight yet, particularly when nanotechnology becomes fully developed [8]. So, if we could figure out a way to construct effective brain-computer interfaces with which a biological brain can tap additional memory and processing capacity, then the required significant increase in brain size can be realised.

The main question is of course what the nature of the extension should be.

- In one hypothesis the artificial brain extensions should mimic the operation of human neurophysiology. Much progress has been made the last years in neural modeling [], and various devices, some of them exploiting VLSI, have been constructed []. So far the performance of these artificial neural networks has fallen far short of natural systems and it could be argued that staying close to human neurophysiology will unnecessarily limit the potential power of artificial systems. For example, if we could extend our brain with a calculating device, we will undoubtely want one that performs calculations with a speed and precision comparable to current computers, as opposed to a calculator that is as slow and error-prone as the human brain. Similarly, if we extend our brain with new language capacities (for example 'plug-in' modules with vocabularies and grammars of a language) then we want this extension to be fast and accurate and not in need of continuous practice as in the case of natural linguistic skills.

- In another hypothesis, the artificial brain may be completely different from the natural brain. It suffices to build bridges between the two so that the contents and processing of one become accessible to the other. Such a hypothesis would rely more on results in knowledge-oriented artificial intelligence research (AI). AI engineering has yielded systems with impressive performance in areas like computer chess, expert problem solving or natural language processing, without mimicking the hardware of the brain. Such a solution would however have other drawbacks. So far AI systems are not situated nor adaptive. They are painstakingly built by human engineers and almost always need to be extended by hand. To make their use in an evolving context feasible, there would at least have to be a mechanism for regular updating. Nobody knows how to do that in an effective way today, despite intensive efforts in machine learning.

Each significant jump forward in the evolution of human intelligence has also coincided with *new sensory modalities and new actuators*. The most dramatic example is the development of the vocal tract two to three hundred years ago which must have triggered (or co-evolved) with the origins of language. Another earlier example is the development of bipedalism with its associated changes in the morphology of the limbs and the thorax. At the moment, highly sophisticated new sensory and actuator modalities are indeed feasible. For example, cameras, microphones, touch sensing devices, locomotion, as based on motors, wheels or legs, can now be constructed at any desired level of precision and sophistication.

The main unresolved problem here is to figure out the brain-computer interface, but significant progress is being made in this area. So far the emphasis is on replacing failing sensory modalities in handicaped persons, but the same technology can also be used to extend other modalities. For example, humans which already have two eyes could be equiped with additional cameras to extend the range of vision, or control directly the behavior of motors for locomotion. If direct connections could somehow be established between the brain and the electronic information highway, then there is the intriguing possibility that the brain has access to vast amounts of information and could in turn cause action at a distance by the intermediary of electronic devices. This idea is at this point so far out that we can hardly imagine its consequences. Will we in the future 'read' electronic mail directly or 'send messages' to other brains without the intermediary of our normal sensory apparatus, or maybe even without the intermediary of language? Will we travel and gain experiences in cyberspace once the appropriate brain-computer interfaces are possible? These possibilities would bring about a complete revolution in our perception of the world. For example, time and space, which are for us now extremely hard constraints, would no longer be limiting factors and thus be experienced in a quite different way.

The evolution of intelligence in humans and animals is however not only characterised by increases in brain capacity and sensory or actuator modalities. There have been each time also extreme *ecological pressures* to evolve towards higher complexity. For example, the evolution towards *Homo Erectus* coincides with the beginnings of migration and thus the need to work in group and handle many different new environmental circumstances. Do we see equal pressures today? It seems so.

1. There has been an exponential growth in the global human population, causing increasing complexity in the management of societies and enormous strains on the available resources. The population in most countries far exceeds what is manageable or viable, causing tensions, exploitation, subgroups that live in extremely difficult circumstances, etc. In the most developed countries the population growth has leveled off or is declining. But even in those cases the leveling off is compensated by immigration from countries with excess population, leading to a steady or slightly increasing population growth. In normal biological systems, excess population is regulated by processes of selection. But our species has become so good in circumventing natural selection that it supposes that resource limitations should no longer apply.

2. There has also been an enormous growth in the contacts between human individuals and subgroups causing a general unstability due to the rapid diffusion of information and the sudden confrontation of groups with very different cultures. Recent violent or non-violent conflicts such as in the former Yugoslavia, Chechenia, Algeria, etc. are almost all about clashes between cultures in which people are unable to make abstraction from relatively small cultural differences (such as difference in language or religion). At the same time, there is a fragmentation because of the availability of so many media. This causes societies to loose their coherence, sometimes leading to political instability.

3. There is an exponential growth in the amount and the availability of information. Whereas scientists or philosophers in the 17th century could still hope to have a reasonable view on the state of knowledge, this is no longer possible. Our brain no longer has the capacity to absorb all the new information, or learn all the skills that could be taught. The 'Renaissance Man' who managed to span the arts, the sciences and engineering, seems today almost out of the question despite the greatly expanded opportunities and powerful tools. This lack of a global perspective makes it much more difficult to further advance the state of our cultures or to prevent narrow-mindedness in tackling global problems. This happens for example when purely technological decisions lead to environmental disasters.

I am of course not suggesting that these problems can only be handled by the development of individuals with more brain power. I hope not, because the problems are so pressing that they need to be solved much sooner. But it does illustrate that the human species is today under just as much stress as it must have been in the past. Ways to make us collectively more intelligent or to have members of the species behave more intelligently seem to be no luxury for the survival of our species.

How realistic is the development of a *Homo Cyber Sapiens*? Today we know almost nothing how we could expand our sensory and actuator modalities nor how we could use plug-in devices to expand our memory and processing capacities possibly with ready-made modules providing knowledge and skills for specific tasks. There are small-scale experiments going on but employment seems to be far off in the future. But then, in how far is it ethically appropriate to push these developments? On the one hand, it appears frightening because the brain is not only the most complex but also the most delicate organ of the human body. It is frightening also because the new species might overpower us. At the same time, the expansion of our brain capacity appears a natural step because the evolution of intelligence has continuously happened in the past as well. Moreover the ecological pressures seem to force the further development of intelligence.

3 Steps towards the *Robot Homonidus Intelligens*.

Although current work on extending our brain capacities is in its infancy, this cannot be said from work on building robotic agents and artificial intelligence. The overall research efforts has so far been relatively small, compared to the research efforts in biology for understanding the functioning of the brain for example, or the research efforts for detecting the elementary particles. Nevertheless there has been constant work in the area since the beginnings of cybernetics and artificial intelligence in the fifties. The results so far are very impressive from various aspects, particularly technological: There have been many spinoffs for computer science, such as list processing, declarative programming, search algorithms, etc. A whole range of programs has been written that exhibit features of (human) intelligence. For example, chess programs now compete at grandmaster level, expert systems have demonstrated human-level performance in difficult problems like scheduling, diagnosis, or design, natural language programs of high complexity have been built for parsing and producing natural language, and some machine learning programs have been capable to extract compact representations from examples.

But there are still very strong limits to the present state of the art which even raise the question whether we will ever get to intelligent robotic agents.

1. Most of the work in knowledge-oriented AI goes through the following cycle: (i) A particular, usually valuable, expertise is identified, (ii) this expertise is modeled at the knowledge level [21] and formalised, and (iii) the formal representation is coded in a computer-based form using symbolic programming techniques. The resulting system, often called a knowledge system, is then incorporated in a particular environment and used as a tool by humans in their work. A high degree of sophistication has been reached in this kind of approach and a wide variety of systems has been built and put into practice. But there are two important problems. The first one is that the effort required to build knowledge systems is very substantial. It takes several man-years of effort to capture moderately sized expertise and this may move up to tens of man-years for non-trivial tasks. A second, more important problem however is that such an effort captures instances of 'frozen intelligence' without the associated mechanisms and processes that gave rise to the intelligent behavior. This means that the maintenance and adaptation has to be done by hand which is quite expensive and in most cases unrealistic. This suggests that something is now known about the internal mechanisms of knowledge representation, reasoning, and problem solving but that we do not understand how these mechanisms develop in interaction with the environment.

2. Knowledge systems are examples of disembodied intelligence. They do not have any direct links to the real world through sensors or actuators. Instead, the link is made through the intermediary of humans. This works reasonably well for domains such as legal reasoning where the inputs and outputs take a form that is already symbolic. But if we are looking at robotic agents which have to operate independently in their environment, then the interfaces to the real world become essential. It has long been assumed that it would be straightforward to link up the symbols used by AI systems to the world through sensors and actuators, but the problems turn out to be enor-

mous [2]. They are mostly related to the fact that the signals obtained from sensors do not contain enough information to extract the detailed symbolic world models required by classical AI planning techniques or that it would take too much time to do so. Moreover actuators will never be full-proof. This means that it will never be possible to plan a particular course of action in advance and hope that the execution will be so perfect that it does not need to be adapted. These difficulties have lead in the last few years to a new 'bottom-up' approach to artificial intelligence which attempts to use minimal world models and emphasises direct coupling between sensing and effecting mediated by dynamical processes [30].

3. Current knowledge systems as well as robots are machines. They do not have any reason of existence on their own. They are not individuals, not even organisms. Although they have in many cases the possibility to choose the most appropriate action from a repertoire of actions, the goals and possible actions are supplied externally. When two robots cooperate, they do so because humans have programmed into them particular behaviors that cause the cooperation. The need for cooperation does not emerge from within the robots or the situations in which they find themselves. In this sense, current AI systems are tools for humans rather than independently existing autonomous entities in their own right. It is at this point fully unclear how a robotic agent could be given a sense of self, the right of initiative, the responsibility for its own action, and so on. Even worse, there is hardly any work going on within the context of artificial intelligence to build truely autonomous agents.

These three unresolved issues constitute formidable obstacles to the development of the *Robot Homonidus Intelligens*. The obstacles are not really technological. The state of the art in electronics, computer engineering, and mechanical engineering make it possible to build the body and brains of a humanoid and efforts in this direction are currently going on [3]. The real obstacles are a lack of a theory of intelligence and particularly of a theory that explains how intelligence grounded in a real world environment may come about.

4 The artificial life approach.

To overcome these fundamental bottlenecks, a new approach towards artificial intelligence has emerged recently which departs in certain ways radically from the knowledge-oriented approach dominating current work in AI [26]. The most characteristic feature of this approach is a move away from engineering and towards biology, not in the sense that attempts are made to develop realistic biological models (as some neural network researchers are trying to do [10]) but in the sense that researchers are trying to understand the *principles* with which biological systems operate and apply them to the construction of artificial systems. The word 'construction' is not even appropriate here, because one of the main ideas is that intelligent autonomous agents cannot be built but should evolve in a process similar to the way that intelligence evolved in nature: using a combination of evolution by natural selection and adaptivity and development as in the development of a biological individual [15], [4].

In the work of our own group we have adopted a number of key hypotheses and we have also set up in the laboratory an experiment with physical robots in which we can experimentally investigate important steps in the evolution of intelligence, such as the origin of behavioral diversity, the origin of communication, the origin of self, the origin of cooperation, etc. The hypotheses and nature of the experiments is now briefly reviewed.

Major hypotheses

1. Against reductionism

For a long time, science has made progress by reducing the complexity at one level by looking at the underlying components. Behavior at a particular level is explained by clarifying the behavior of the components at the next level down. Also in the case of intelligence, we see that many researchers hope that an understanding of intelligence will come from understanding the behavior of the underlying components. For example, most neurophysiologists believe that a theory of intelligence will result from understanding the behavior of neural networks in the brain. Some physicists go even so far as to claim that only a reduction of the biochemical structures and processes in the brain to the quantum level will provide an explanation of intelligence [23].

At the moment there is however a strong opposing tendency to adopt a wholistic point of view, also in the basic sciences [5]. This means that it is now understood that there are properties at each level which cannot be reduced to the level below, but follow from the dynamics at that level, and from interactions (resonances) between the dynamics of the different levels [22]. In the case of intelligence, this means that it will not be possible to understand intelligence by only focusing on the structures and processes causally determining observable behavior. Part of the explanation of intelligence will have to come from the internal dynamics, the interaction with the structures and processes in the environment, and the coupling between the different levels.

2. Foundation in dynamics

The classical AI approach constructs theories of intelligence based on logic (see e.g. [11]). But the search for a theory of intelligence which is compatible with physics and biology and which sees intelligence as a universal phenomenon present at many different levels of biological systems, pushes us into another direction. Most theories of complex natural phenomena are phrased in terms of the recently developed theory of complex dynamical systems, which includes theories of chaos and self-organisation [22]. A similar approach seems necessary. Many steps in this direction have already been set, particularly by researchers exploring neural networks [14]. They may lead to whole new theories, including dynamical theories of consciousness [29].

3. Biological grounding

Living systems are systems which actively maintain themselves. They therefore are selfish in the sense that their internal mechanisms and their behavior are geared towards their own survival [6]. This selfishness can be seen at each level of biological complexity from genes to societies. The stronger a level establishes itself, the more it will resist to be sacrificed for the selfishness of one of the other levels. When we want autonomous robotic *agents* as opposed to machines, they will have to be embodied with the same drive towards maintaining themselves, and this drive should be the main force towards the development of more complexity. For example, the primary need for a robot is to have sufficient energy to continue its existence. But a robot may also be in a position that the survival of other robots is a condition for its own survival, forcing the robot towards cooperation despite competition. The internal structures of a robot can in turn be seen to be in competition for existence and the evolution of new behavioral complexity can be driven by selectionist mechanisms [27].

Experimental approach

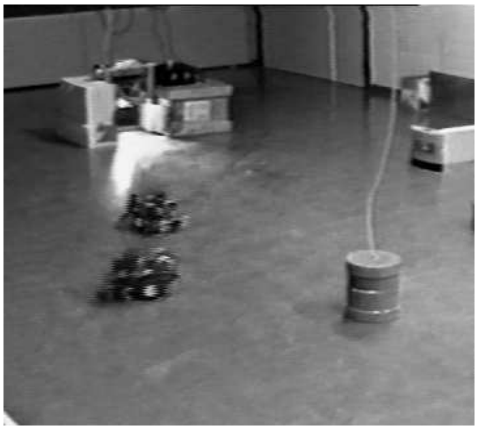

Figure 1: Robotic ecosystem constructed at the VUB AI laboratory. There is a charging station which robots can use to recharge their batteries. There are also 'parasites' in the form of lamps which take energy away from the charging station. Robots temporarily kill off parasites by pushing against the boxes.

In our laboratory, we have created a complete robotic ecosystem (figure 4) which involves an environment with different pressures for the robots (e.g. the need to collect energy and ensure that it is available), different robotic agents which have to cooperate but are also in competition with each other, and a growing repertoire of adaptive structural components (called behavior systems) which are causally responsible for behavior. (see [28], [17]).

Such an integrated experimental environment ensures that all the different levels (genetic, structural, individual, group) are present at the same time, each with strong interactions to the environment. This way a wholistic approach to the study of intelligence is possible. Our objective is to come up and test out scenarios that show the progressive steps towards intelligent agents, similar to the way biologists and chemists are investigating scenarios for the origin of life [13] or physicists are researching scenarios for the origin of the universe [32]. The big challenge is to use only principles compatible with the laws of physics and biology and to avoid programming specific mental capabilities. Instead, intelligent behavior, including individuality, linguistic communication, cooperation, etc. must emerge as a result of pressures in the ecosystem in which the agents find themselves and structure-forming processes such as self-organisation or selectionism.

Here are two more concrete example experiments:

1. The first one focuses on the emergence of diversity in robotic agents and is described in more detail in [31]. The experiment starts with identical agents which are cooperating in the exploitation of energy in the ecosystem and the elimination of competition (the parasites). The robots do not know initially how much work they have to do before going to the recharging station and so this behavior has to be adaptive. Because the behavior of the agents is coupled, a small difference of behavior, in which one agent works slightly more than another one, may get re-enforced leading to two types of agents where one works on average significantly more than the other one. Thus a social differentiation occurs which gives the less hard working group more time for other activities.

2. The second experiment focuses on the emergence of identity and coupled to it communicative signalling. The experiment again starts with identical agents (at least 4) which are in competition for energy but also can and must cooperate. A robot on the charging station has control over the charging station because it can turn off the light that other robots need to use in their phototaxis towards the station. Pairs of robots may develop a strategy where one robot only releases the charging station to another robot. This strategy requires that one robot recognises the other one, for example by distinguishing the sound pattern that they produce. Another robot can only break into a relation by imitating this sound pattern which will in turn lead to more sophisticated signalling. This phenomenon is probably responsible for the complexity of bird song. Such signalling is one of the key steps towards the development of language. More importantly, the experiment demonstrates how pressures in an ecosystem and 'selfish' behavior with only a local view of the ecosystem may lead to individuality.

These and other experiments are extremely hard to carry out and there are as yet very little theoretical ideas on what mechanisms could explain the emergence of intelligent behavior. Evolution by natural selection is undoubtely one of them, but there are probably many other mechanisms at work. Nevertheless, we believe that it is only by probing the basic questions of the origins of intelligence that we will one day be able to evolve autonomous robotic agents with the complexity of humans. Our own research suggests that this day is way off into the future.

5 Conclusions

The paper discussed possible ways in which new forms of intelligence comparable, or even more powerful, than current human intelligence may come about. One way, the *Homo Cyber Sapiens*, is rooted in biology and based on extensions by technological artefacts. The other way, the *Robot Homonidus Intelligens*, is purely technological. Both developments may happen, driven by new technological advances and the increasing pressures occuring in human societies, although their realisation is probably far into the future. The paper argued that for robotic agents, a biological approach in which such agents are evolved rather than completely designed and programmed, is the only viable way.

References

[1] Babloyantz, A. (1991) Self-Organisation, Emerging Properties, and Learning. NATO ASI Series. New York: Plenum Press.

[2] Brooks, R. (1991) Intelligence without reason, IJCAI-91, Sydney, Australia, pp 569–595.

[3] Brooks, R. (1994) Coherent Behavior from Many Adaptive Processes. In: Cliff, D. et.al. (1994) eds. From animals to animats 3. Cambridge: MIT Press. p. 22-29.

[4] Cliff, D., P. Husbands, and I. Harvey (1993) Evolving Visually Guided Robots. In: Meyer, J-A., H.L. Roitblatt, and S.W. Wilson (1993) From Animals to Animats2. Proceedings of the Second International Conference on Simulation of Adaptive Behavior. MIT Press/Bradford Books, Cambridge Ma. p. 374-383.

[5] Cohen, J. and I. Stewart (1994) The Collapse of Chaos. London: Viking Press.

[6] Dawkins, R. (1976) The Selfish Gene. Oxford University Press. Oxford.

[7] Donald, M. (1991) Origins of the modern mind. Three stages in the evolution of culture and cognition. Cambridge Ma: Harvard University Press.

[8] Drexler, K. (1994) Nanosystems. Molecular Machinery, Manufacturing, and Computation. New York: John Wiley.

[9] Dunbar, R. (1988) Primate Social Systems. London: Croom Helm.

[10] Eeckman, F. and J. Bower (1993) (eds.) Computation and Neural Systems. Boston: Kluwer Academic Publishers.

[11] Genesereth, M. and N. Nilsson (1987) Logical Foundations of Artificial Intelligence. Morgan Kaufmann Pub. Los Altos.

[12] Gould, S. and N. Eldredge (1977) Punctuated Equilibria: The tempo and mode of evolution reconsidered. Paleobiology, 3: 115-151.

[13] Kauffman, S.A. (1993) The origins of order: self organization and selection in evolution. Oxford University Press, Oxford.

[14] Kosko, B. (1992) Neural Networks and Fuzzy Systems. A Dynamical Systems Approach to Machine Intelligence. Prentice Hall, Englewood Cliffs, NJ.

[15] Koza, J. (1992) Genetic Programming. MIT Press, Cambridge Ma.

[16] Lieberman, P. (1984) The Biology and Evolution of Language. Cambridge, Ma: Harvard University Press.

[17] McFarland, D. (1994) Towards Robot Cooperation. Proceedings of the Simulation of Adaptive Behavior Conference. Brighton. MIT Press.

[18] Minsky, M. (1994) Will Robots Inherit the Earth? In: Scientific American October 1994. p. 87-91

[19] Moravec, H. (1988) Mind Children. The Future of Robot and Human Intelligence. Cambridge Ma: MIT Press.

[20] Nadel, L. et.al. (1989) Neural Connections, Mental Computations. Cambridge Ma: MIT Press.

[21] Newell,A. (1981). The knowledge level, *Journal of Artificial Intelligence*, vol 18, no 1, pp 87–127.

[22] Nicolis, G. and I. Prigogine (1985) Exploring Complexity. Piper, Munchen.

[23] Penrose, R. (1990) The Emperor's New Mind. Oxford University Press. Oxford.

[24] Smithers, T. (1994) Are autonomous agents information processing systems? In: Steels, L. and R. Brooks, (Eds.), *The 'artificial life' route to 'artifical intelligence': building situated embodied agents*, Lawrence Erlbaum Associates, New Haven.

[25] Spencer, A.J. (1982) Death in Ancient Egypt. London: Penguin Books.

[26] Steels, L. (1994) The Artificial Life roots of Artificial Intelligence. Artificial Life Journal, Vol 1:1/2. p. 89-125.

[27] Steels, L. (1994a) Emergent Functionality in Robotic Agents through on-line Evolution. In: Proceedings of the Artificial Life Conference MIT. Cambridge: MIT Press.

[28] Steels, L. (1994b) A case study in the behavior-oriented design of autonomous agents. Proceedings of the Simulation of Adaptive Behavior Conference. Brighton. Cambridge, Ma: MIT Press.

[29] Steels, L. (1994c) Is artificial consciousness possible? In: Traitteur, G. (ed). Consciousness and Cognition. Kluwer Academic Publishers, Dordrecht. 1993.

[30] Steels, L. and R. Brooks (eds.) (1995) The 'artificial life' route to 'artificial intelligence'. Building Situated Embodied Agents. New Haven: Lawrence Erlbaum Ass.

[31] Steels, L. (1995) Intelligence - dynamics and representation. In: Steels, L. (1995) The Biology and Technology of Intelligent Autonomous Agents. NATO ASI Series. Berlin: Springer Verlag.

[32] . Weinberg, S. (1990) The first three minutes. New York: Vintage Books.