

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/221025750>

Embodied Artificial Intelligence: Trends and Challenges

Conference Paper · January 2003

DOI: 10.1007/978-3-540-27833-7_1 · Source: DBLP

CITATIONS

45

READS

2,615

2 authors:



Rolf Pfeifer

Osaka University, Toyonaka, Japan

250 PUBLICATIONS 9,907 CITATIONS

[SEE PROFILE](#)



Fumiya Iida

University of Cambridge

189 PUBLICATIONS 4,048 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



Amphibious legged robot locomotion [View project](#)



Locognition [View project](#)

Embodied Artificial Intelligence: Trends and Challenges

Rolf Pfeifer and Fumiya Iida

Artificial Intelligence Laboratory, Department of Informatics, University of Zurich
Andreasstrasse 15, CH-8050 Zurich, Switzerland
{pfeifer,iida}@ifi.unizh.ch

Abstract. The field of Artificial Intelligence, which started roughly half a century ago, has a turbulent history. In the 1980s there has been a major paradigm shift towards embodiment. While embodied artificial intelligence is still highly diverse, changing, and far from “theoretically stable”, a certain consensus about the important issues and methods has been achieved or is rapidly emerging. In this non-technical paper we briefly characterize the field, summarize its achievements, and identify important issues for future research. One of the fundamental unresolved problems has been and still is how thinking emerges from an embodied system. Provocatively speaking, the central issue could be captured by the question “How does walking relate to thinking?”

1. Introduction

This conference and this paper are about embodied artificial intelligence. If you search for “embodied artificial intelligence” or “embodied cognition” on the Internet using your favorite search engine, you will find a radically smaller number of entries than if you search for “artificial intelligence” or “cognition”. Trying to answer this question of why this might be the case, reveals a lot about the structure of this research field and uncovering its organization is one of the goals of this paper.

Over the last 50 years Artificial Intelligence (AI) has changed dramatically from a computational discipline into a highly transdisciplinary one that incorporates many different areas. Embodied AI, because of its very nature of being about embodied systems in the real physical and social world, must deal with many issues that are entirely alien to a computational perspective: as we will discuss later, physical organisms in the real world, whether biological or artificial, are highly complex and their investigation requires the cooperation of many different areas. The implications of this change in perspective are far-reaching and can hardly be overestimated. In this paper, we will try to outline some of them.

With the fundamental paradigm shift from a computational to an embodied perspective, the kinds of research topics, the theoretical and engineering issues, and the disciplines involved have undergone dramatic changes, or stated differently, the “landscape” has been completely transformed. In the first part of the paper we try to characterize these changes. In the second part, we will identify the grand challenges in

the field and discuss how far researchers have come towards achieving them. Given the enormous diversity, as discussed in the first part, this will necessarily be abstract, somewhat selective and reflect the authors' personal opinion, but we do hope that many people will agree with our description of how the field is now structured. We conclude with some general comments on the future of the field and applications.

2. The “landscape”

The landscape of artificial intelligence has always been rugged but it has become even more so over the last two decades. When the field started initially, roughly half a century ago, intelligence was essentially viewed as a computational process. Research topics included abstract problem solving and reasoning, knowledge representation, theorem proving, formal games like chess, search techniques, and – written – natural language, topics normal associated with higher level intelligence. It should be mentioned however, that in the 60s there was a considerable amount of research on robotics in artificial intelligence at MIT, SRI, and CMU. But later on the artificial intelligence research community has not paid much attention to this line of work.

Successes of the classical approach

By the mid 1980s, the classical, computational or cognitivistic approach, had grown into a large discipline with many facets and has brought forward many successes in terms of computer and engineering applications. If you start your favorite search engine on the Internet, you are, among many others, employing clever machine learning algorithms. Text processing system utilizes matching algorithms, or algorithms that try to infer user's intentions from the context of what have been done earlier. Controls for appliances using fuzzy logic, embedded systems (as they are employed in fuel injection systems, braking systems, air conditioners, etc.), control systems for elevators, and trains, natural language interfaces to directory information systems, translation support software, etc., are also among the successes of the classical approach. More recently, data mining systems have been developed that heavily rely on machine learning techniques, and chess programs have been realized that beat 99.99 percent of all humans on earth, a considerable achievement indeed! The development of these kinds of systems, although they have their origin in artificial intelligence, have now become indistinguishable from applied informatics in general: they have become a firm constituent of any computer science department.

Problems of the classical approach

However, the original intention of artificial intelligence was not only to develop clever algorithms, but also to understand natural forms of intelligence that have – as argued here – more to do with the interaction with the real world. Alas, as is now generally agreed, the classical approach has not contributed significantly to our understanding of,

for example, perception, locomotion, manipulation, everyday speech and conversation, social interaction in general, common sense, emotion, and so on.

Classical approaches to computer vision, for example, have been successful in factory environments, where there are constant lighting conditions, the geometry of the situation is precisely known (i.e. the camera is always in the same place, the objects appear on the conveyer belt always in the same position), and the types of potential objects are known and can therefore be modeled. However, when these conditions do not hold – and in the real world, they are never given, i.e. the distance of objects from the eyes always changes, which is one of the many consequences of moving around, and lighting conditions and orientation also vary continuously – these algorithms can no longer be used. Moreover, objects are often entirely or partially occluded, they move (e.g. cars, people), and they appear against very different and changing backgrounds. Artificial vision systems with capacities similar to human or animal vision, are far from being realized artificially.

A further example where the classical approach could not provide adequate answers is object manipulation. Indeed, animals and humans are enormously skilled at manipulating objects; even very simple animals like insects are masters at manipulation. Or watch a dog chew on a bone, how he controls it with his paws, mouth and tongue: unbelievable. Although there are specialized machines for virtually any kind of manipulation (driving a screw, picking up objects for packaging in production lines, lifting heavy objects in construction sites), the general purpose manipulation abilities of natural systems are to date unparalleled.

Locomotion is another case in point. Animals and humans move with an uncanny flexibility and elegance. We can walk with a bag in one hand, an arm around a friend, up and down the stairs, while looking around, something none of the existing robots can do. And building a running robot is still considered one of the great challenges.

In the classical approach, common sense has been treated at the level of “semantic content” and has been taken to include knowledge such as “cars cannot become pregnant”, “objects (normally) don’t fly”, “people have biological needs” (they get hungry and thirsty), etc. Building systems with this type of common-sense knowledge has been the goal of many classical natural language and problem solving systems like CYC (e.g. Lenat et al., 1986). But there is an important additional aspect of common-sense knowledge, which is to do with bodily sensations and feelings, and this aspect has its origin largely in our embodiment. Take, for example, the word “drinking” and freely associate what comes to mind. Perhaps being thirsty, liquid, cool drink, beer, hot sunshine, the feeling of wetness in your mouth, on the lips, and on your tongue when you are drinking, and the feeling of thirst disappearing as you drink, etc. It is this kind of common sense knowledge and common experience that everyone shares and that forms the basis of robust natural language communication, and it is firmly grounded in our own specific embodiment. And to our knowledge, there are currently no artificial systems, capable of dealing with this kind of knowledge in a flexible and adaptive way.

The last point that we would like to mention here concerns speech systems. While in restricted areas, speech systems can be used, e.g. as an interface to directory information systems, or systems where single word commands can be used (e.g. for robot control, or name databases for mobile phones), in most areas they have only been used

with limited success. Speech to text systems have to be tuned to the speaker's voice, and because of the high error rate, typically a lot of post-editing needs to be done on the text produced by the software. This may be one of the reasons why speech systems have not really taken off so far, even though the idea of not having to type any more, of producing text rapidly, is highly appealing. Although some of the systems may have a relatively impressive performance, the fact of the matter remains that there are to date no general purpose natural language systems whose performance even remotely resembles the one of humans in a free format everyday conversation.

Finally, it is interesting to note, It is interesting to note that these more natural kinds of activities (perception, manipulation, speech) are all activities that have, in some very essential ways, to do with complex, "high bandwidth" interaction with the real world. We will come back to this point later on.

Embodied Artificial Intelligence

These failures, largely due to the lack of rich system-environment interaction, have lead some researchers to pursue a different avenue, the one of embodiment. With this change of orientation, the nature of the research questions also began to change. Rodney Brooks, one of the first promoters of embodied intelligence (e.g. Brooks, 1991), started studying insect-like locomotion, building, for example, the six-legged walking robot "Ghengis". So, walking and locomotion in general became important research areas, topics typically associated with low-level sensory-motor intelligence. This is, of course, a fundamental change from studying chess, theorem proving, and abstract problem solving, and it is far from obvious how the two relate to one another, an issue that we will elaborate in detail later. Other subjects that people started investigating have been orientation behavior (i.e. finding one's way in only partially known and changing environments), path-finding, and elementary behaviors such as wall following, and obstacle avoidance.

The perspective of embodiment requires working with real world physical systems, i.e. robots. A crucial aspect of embodiment is that it requires working with real world physical systems, i.e. robots. Computers and robots are an entirely different ball game: computers are neat and clean, they have clearly defined inputs and outputs, and anybody can use them, can program them, and can perform simulations. Computers also have for the better part only very limited types of interaction with the outside world: input is via keyboard or mouse click, and output is via display panel. In other words, the "bandwidth" of communication with the environment is extremely low. Also computers follow clearly defined "input processing" output scheme that has, by the way, shaped the way we think about intelligent systems and has become the guiding metaphor of the classical cognitivist approach. Robots, by contrast, have a much wider sensory-motor repertoire that enables a tight coupling with the outside world and the computer metaphor of input-processing-output can no longer be directly applied.

Building robots requires engineering expertise, which is typically not present in computer science laboratories, let alone psychology departments. So, with the advent of embodiment the nature of the field, artificial intelligence, changed dramatically. While in the traditional approach, because of the interest in high-level intelligence, the

relation to psychology, in particular cognitive psychology was very prominent, the attention, at least in the early days of the approach of embodied intelligence, shifted more towards – non-human – biological systems, such as insects, but other kinds of animals as well.

Also, at this point, the meaning of the term “artificial intelligence” started to change, or rather started to adopt two meanings. One meaning stands for GOFAI (Good Old-Fashioned Artificial Intelligence), the traditional algorithmic approach. The other one designates the embodied approach, a paradigm that employs the synthetic methodology which has three goals: (1) understanding biological systems, (2) abstracting general principles of intelligent behavior, and (3) the application of this knowledge to build artificial systems such as robots or intelligent devices in general. As a result, the modern, embodied approach started to move out of computer science laboratories more into robotics and engineering or biology labs.

It is also of interest to look at the role of neuroscience in this context. In the 1970s and early 1980s, as researchers in artificial intelligence started to realize the problems with the traditional symbol processing approach, the field of artificial neural networks, an area that had been around since the 1950s, started to take off – new hope for AI researchers who had been struggling with the fundamental problems of the symbol processing paradigm. Inspiration was drawn from the brain, but only at a very abstract level. In the embodied approach, there was a renewed and much stronger interest in neuroscience because researchers realized that natural neural systems are extremely robust and efficient at controlling the interaction with the real world. As mentioned above, animals can move and manipulate objects with great ease, and they are controlled by – natural – neural networks. In addition, they can move very elegantly, with great speed and with little energy consumption. These impressive kinds of behaviors can only be achieved if the dynamical properties of the neural networks are exploited. This is quite in contrast to the traditional AI approach where mostly static feedforward networks were employed.

Diversification

So, in terms of research disciplines participating in the AI adventure, we see that in the classical approach it was mainly computer science, psychology, philosophy, and linguistics, whereas in the embodied approach, it is computer science and philosophy as before, but also engineering, robotics, biology, and neuroscience (with a focus on dynamics), whereas psychology and linguistics have lost their role as core disciplines. So we see somewhat of a shift from high-level (psychology, linguistics) to more low-level sensory-motor processes, with the neurosciences covering both aspects, sensory-motor and cognitive levels. With this shift, the terms used for describing the research area shifted: researchers working in the embodied approach no longer referred to themselves as working in artificial intelligence but more in robotics, engineering of adaptive systems, artificial life, adaptive locomotion, bio-inspired systems, and neuroinformatics. But more than that, not only have researchers in artificial intelligence moved into neighboring fields, but researchers that have their origins in these other fields started in natural ways to contribute to artificial intelligence. This way, the field

on the one hand significantly expanded, but on the other, its boundaries became even more fuzzy and ill-defined than before.

These considerations also provide a partial answer to the question of why we don't get many entries when we type "embodied intelligence" or "embodied artificial intelligence" into one of the search engines: Because the communities started to split and researchers in embodied intelligence started attending other kinds of conferences, e.g. "Intelligent Autonomous Systems, IAS", "Simulation of Adaptive Behavior – From Animals to Animats, SAB", "International Conference on Intelligent Robotics and Systems, IROS", "Adaptive Motion in Animals and Machines, AMAM", "European Conference on Artificial Life, ECAL", "Artificial Life Conference, ALIFE", "Artificial Life and Robotics, AROB", "Evolutionary Robotics, ER", or the various IEEE conferences (International Society of Electrical & Electronics Engineering), etc. Anecdotal speaking, I (Rolf Pfeifer) remember that initially, in the early 90s, when I tried to convince people at AI conferences such as International Joint Conference on Artificial Intelligence (IJCAI), the European Conference on Artificial Intelligence, ECAI, or the German annual AI Conference, that embodiment is not only interesting but essential to understanding intelligence, I mostly got very negative reactions and no real discussion was possible. So, together with many colleagues we turned to other conferences where people were more receptive to these new ideas. More recently, perhaps because of the stagnation in the field of classical AI in terms of tackling the big problems about the nature of intelligence, there has been a growing interest in embodiment and now AI conferences, at least some of them, have started workshops on issues in embodiment. But by and large, the two communities, the classical and the embodied one, are pretty much separate, and will probably remain so for a while.

Biorobotics

There are a number of additional interesting developments worth mentioning here. One is, in the field of embodiment, a renewed interest in high-level cognition. Rodney Brooks, at the time, had forcefully argued that getting insects to walk from scratch took evolution much longer than getting from insects to humans. This implies that creating insects was the really hard problem and after that, moving towards human level intelligence was relatively easy. Thus, so his conclusion, one should first work on insects rather than humans, one should do "biorobotics".

Many people started doing biorobotics and began cooperations with biology laboratories. An excellent example is the work by Dimitrios Lambrinos at the Artificial Intelligence Laboratory in Zurich, who started to cooperate with the world champion in ant navigation, Ruediger Wehner of the University of Zurich. Jointly, they built a series of robots, the Sahabot-Series that mimic long- and short-term navigation behavior of the desert ant *Cataglyphis* (e.g. Lambrinos et al., 2000). Rodney Brooks cooperated with the famous biologist Holk Cruse of the University of Bielefeld in Germany, who had been studying insect walking for many years and who had found that there is no central control for leg coordination in walking in ants. Brooks implemented Cruse's ideas on an MIT ant-like robot and termed the controller "cru(i)se" control, in honor of the designer, Holk Cruse. There are many examples of such cooperation which have all been very productive (for an excellent collection of

which have all been very productive (for an excellent collection of papers on biorobotics, see (Webb and Consi, 2000)).

Developmental robotics

However, after a few years of working on insect like behavior, Brooks started changing research topics. He argued that we have to “think big” and should work towards human level intelligence, and the project “Cog” for the development of a humanoid robot, was born (Brooks and Stein, 1993). He neatly mapped out the necessary steps and stages for achieving human-level intelligence, but due to many problems, after less than 10 years, changed topics again. But the Cog project generated a lot of excitement and many researchers were attracted by the idea of moving towards human-level intelligence, which had also been the target of classical artificial intelligence, and the field of developmental robotics emerged. The term developmental robotics designates the attempt to model aspects of human or primate development using real robots. Its pertinent conferences come under many labels, “Emergence and Development of Embodied Cognition, EDEC”, “Epigenetic robotics”, “Development of Embodied Cognition, DECO”, “International Conference on Development and Learning”, etc. This was, of course, a happy turn for those who might have been slightly sad or disappointed by the direction the field took – insects simply are not as sexy as humans! And human intelligence happens to be the most fascinating type of intelligence that we know. But once again, this strand of conferences is separate from the traditional ones in artificial intelligence, and they do not contain the term “embodied intelligence”.

Ubiquitous computing

Another line of development that should be introduced here is the one of ubiquitous computing (Weiser 1993). Computer science has undergone dramatic changes as well. Computing as such, software engineering, the development of algorithms, operating systems, the virtual machine, etc. are topics that we now understand relatively well and it is not clear whether there will be big innovations in these areas in the near future. Rather, it seems that the new challenges are seen in the interaction with the real world. Initially, the field was characterized by the idea of putting sensors everywhere, into rooms (mostly cameras, motion detectors), floors (e.g. pressure sensors to detect the position of individuals) objects such as cars, chairs, beds, but also cups, or any kind of devices such as mobile telephones, clothes (e.g. t-shirts, shoes) to measure physiological data of the individual wearing them for sports or medical reasons (the list is in fact endless). More recently, ubiquitous computing has also been investigating actuation, i.e. ways in which systems can influence their environments: control systems for buildings for temperature, humidity, windows, and blinds; cars that automatically apply their breaks when the distance to the car in front gets too small, or – in the medical domain – systems that monitor physiological variables (pulse rate, skin resistance, level of dehydration) and send a message to a physician if necessary. The

field of ubiquitous computing is closely related to user interfaces or generally to human-machine interaction.

Even though user interfaces have always been an important topic in computers, the problem, in contrast to robotics, has been the low “bandwidth of communication”, as pointed out earlier. In order to increase this “bandwidth”, there has been a lot of work on speech, spoken language, to interact with computers, but these efforts, for various reasons, have only been met with very limited success (see our discussion above). Just recently have there been projects for developing more interesting and richer interfaces using, for example, touch, and to some extent vision. There is also work on smell but that has – although very exciting – not yet advanced significantly. The research on wearables should be pointed out here as well. What is interesting about these “movements”, human-machine interface, wearables, ubiquitous computing, is that now virtually all computer science departments start moving into the real world. They are not doing robotics per se, but many have started hiring engineers and establishing mechanical and electronics workshops where they can build hardware, because now real-world devices with certain sensory-motor abilities need to be constructed, devices that could be called “robotic devices”. So far as we can tell, there has been little theory development, but there is a lot of creative experimentation going on. We feel that the set of design principles that we have developed for embodied systems will be extremely useful in designing such systems (e.g. Pfeifer and Scheier, 1999). For example, the principle of sensory-motor coordination which states that through the – active – interaction with the environment, patterns of sensory stimulation are induced that are correlated across sensory modalities, is an important guiding principle, but has, to date, not been applied. We might also say that computer science has now come full circle, from disembodied algorithm to embodied real-world computing, or rather real-world interaction, with embodied artificial intelligence as the fore-runner.

Artificial life and multi-agent systems

Another interesting line of development has its origins in the field of Artificial Life, also called Alife for short. The classical perspective of artificial intelligence had a strong focus on the individual, just as psychology, and psychology was the major discipline with which artificial intelligence researchers cooperated at the time. ALife research which has strong roots in biology – rather than psychology – has been focusing on emergence of behavior in large populations of agents, in other words it is interested in what some call multi-agent systems. We deliberately say “that some call multi-agent systems” because normally, in Alife research, the term complex dynamical system is preferred, as it encompasses also physical systems where the individual components only have limited “agent character”, e.g. the molecules in the famous Bénard experiment. An agent typically has certain sensory-motor abilities, i.e. it can perceive aspects of the environment, and depending on this information and its own state, performs a particular behavior. Molecules, rocks, or other “dead” physical objects do not have this ability. One point of interest has been the emergence of complex global behavior from simple rules and local interactions. (Langton, 1995)

Modular robotics, a research area that has drawn inspiration from artificial life research, also relates to multi-agent systems, where the individual agents are robotic modules capable of configuring into different morphologies (see the volume by Hara and Pfeifer (2003) for examples of modular robotic systems). One of the goals of this research is to design systems capable of self-repair, a property that all living systems have to some extent. Self-assembly and self-reconfiguration are fascinating topics that will become increasingly important as systems have to operate over extended periods of time in remote, hostile environments. The seminal work by Murata and his co-workers (Murata et al., 2004) demonstrates, how self-reconfiguration can be achieved not only in simulation but with real robotic systems. It should be mentioned, however, that to date, much of the research on self-repair and self-reconfiguration is tightly controlled, rather than being emergent from local interactions.

Evolutionary systems are another example of “population thinking”, where the adaptivity of entire populations is studied rather than that of individuals. Because of its close relation to biology, economics has also taken inspiration from multi-agent systems and created the discipline of agent-based economics (e.g. Epstein and Axtell, 1996). Work on self-organization in insect societies, for example, by Jean-Louis Deneubourg of the Université libre de Bruxelles, has attracted many researchers from different fields: “ant intelligence” was one of their slogans (e.g. Bonabeau et al., 1999).

Interestingly, the term multi-agent systems has quickly been adopted by researchers in classical artificial intelligence. However, rather than looking for emergence, they endowed their individual agents with the same types of centralized control that they used for individuals (e.g. Ferber, 1999). As a consequence they could not study emergent phenomena, and a look into the journal “Autonomous Agents and Multi-Agent Systems” shows that the research under the heading “multi-agent systems” typically has different goals and does not focus on emergence. For the better part, the research is geared towards internet applications using software agents.

In robotics there has also been an interest in multi-agent systems. There the problem has been that often only relatively few robots have been available so that it has proved difficult to investigate emergence phenomena in populations. This is illustrated by the rapidly growing “Robocup” or robot soccer community. Initially the robots, for the better part, were programmed directly by the designers in order to win the game. More recently there has been growing interest and significant results in producing scientifically compelling and elegant solutions by incorporating ideas of emergence, but this still remains a big challenge.

One of the important research problems and limitations so far has been the achievement of higher levels of intelligence by the multi-agent community: typically, as in the work of ethologist and Alife researcher Charlotte Hemelrijk, the interest is in emergent hierarchies, group size formation, or migration patterns. Thinking, reasoning, or language, have typically not been topics of interest here. An exception is the work of the group of researchers interested in evolution of communication and evolution of language. An excellent example of this type of research that tries to combine population thinking or multi-agent systems with higher-level processes such as language is the “Talking Heads” experiment by Luc Steels (e.g. Steels, 2001, 2003). In an ingenious experiment he could demonstrate how, for example, a common vocabulary emerges through interaction of agents with their environment and with other agents via

a language game. He has also been working on emergence of syntax, but in these experiments many assumptions have to be made to bootstrap the process. In this research strand, many insights have been gained into how communication systems establish themselves and how something like grammar could emerge. Although fascinating and highly promising, the jury is still out on whether this approach will indeed lead to something resembling human natural language.

Because of the fundamental differences in goals, the distributed agents community, artificial life style, and the artificial intelligence and robotics community, individual style, have to date remained largely separate.

Summary

In summary, we can see that the landscape has changed significantly: while originally artificial intelligence was clearly a computational discipline, dominated by computer science, cognitive psychology, linguistics, and philosophy, it has turned into a multidisciplinary field requiring the cooperation and talents of many other fields such as biology, neuroscience, engineering (electronic and mechanical), robotics, biomechanics, material sciences, and dynamical systems. And this exciting new transdisciplinary community is now called “embodied artificial intelligence.” While for some time, psychology and linguistics have not been at center stage, with the rise of developmental robotics, there has been renewed interest in these disciplines. The ultimate quest to understand and build systems capable of high-level thinking and natural language, and ultimately consciousness, has remained unchanged. Only the path on how to get there is fundamentally different. Although the emergence of ideas of embodiment can be traced back to pre-Socratic thinking and can be found throughout the history of philosophy, the recent developments in artificial intelligence that enable not only the analysis but also the construction of embodied systems, are supplying ample novel intellectual fodder for philosophers. As we will show later, these developments significantly change the image we have of ourselves and our society.

In spite of the multifaceted nature, there is a unifying principle and that is the actual agent to be designed in the context of the synthetic methodology, be it physical in the real world, or simulated in a realistic physics-based simulation. Such agents have a highly integrating function by bringing together results from all these different areas, and allowing concrete testing in an objective way. Moreover, they serve as excellent platforms for transdisciplinary research and communication.

3. State-of-the-art and challenges

Given the diversity of embodied artificial intelligence and the ruggedness of the landscape it will be next to impossible to come up with a set of challenges and a characterization of the state-of-the-art that everybody will agree on.

In characterizing the state-of-the-art we will start from the overall challenges that we will organize according to the three time scales (“here and now”, ontogenetic, phylogenetic) (see Table 1). These time scales, although clearly identifiable, have

important interactions, a point that we will also take into account. Moreover, we will divide our discussion into two parts, theoretical/ conceptual, and engineering. In identifying the challenges and research issues we tried to do a comprehensive survey of the literature and we, in particular, consulted the papers in this volume in order to assess the important trends. By the very nature of this endeavor of identifying challenges, this will be rather subjective and mirrors the personal research interests of the authors.

Table 1. Time scales for understanding and designing agents

<i>time scale</i>	<i>designer commitments</i>
state-oriented "here and now"	"hand design"
learning and development "ontogenetic"	initial conditions; learning and developmental processes
evolutionary "phylogenetic"	evolutionary algorithms; morphogenetic processes

However, we do believe that they reflect, one way or other, the important directions in the field. Nevertheless, we do not expect everyone to agree.

We propose the following "grand challenges" for future research, theoretical understanding of behavior; achieving higher level intelligence; automated design methods (artificial evolution and morphogenesis), and "moving into the real world".

Theoretical understanding of behavior

By theoretical understanding of behavior we mean an understanding of how particular behaviors in the real world can be achieved in artificial agents. This may also shed light on how particular behaviors that we observe in nature come about, which is also one of the goals of artificial intelligence research. This goal is mainly to do with the "here and now" time scale, i.e. with the question of the mechanisms underlying behavior. Although a vast body of knowledge has been accumulated this still remains one of the big conundrums.

As outlined in the previous section, many research areas and a host of studies have contributed to this understanding. However, we still don't have, for example, general purpose perceptual systems – human or primate vision is still unparalleled, and we still have an insufficient understanding of motor control, e.g. how we can achieve rapid legged locomotion. For example, there has been a lot of progress in research on humanoid walking robots, especially in Japan (e.g. Sony's QRIO, Honda's Asimo, Kawada's HRP, the University of Tokyo's H-7, to mention but a few). However, although most of these robots show impressive performance, they still walk slower than humans, their walking style looks somewhat unnatural, and research on running is still in its infancy.

One of the issues, and this is one of the challenges, is the fact that most of the research has been focused on control, which has been, and still is, the standard perspective in robotics. Recent work in the area of biomechanics seems to suggest that material and morphological properties, i.e. the intrinsic dynamical properties of the muscle-tendon systems and the specific shapes and material properties of the limbs and the body play an essential role in locomotion (e.g. Blickhan et al., 2003; Kubow and Full, 1999), but also in behavior in general, e.g. object manipulation, posture control, gesturing, etc. These ideas are captured in the theoretical principle of “ecological balance”, as outlined by Pfeifer et al., (in press), Hara and Pfeifer (2000), Ishiguro et al., (2003) and earlier in Pfeifer and Scheier (1999), which states that there is a balance or task distribution between morphology, materials, control, and interaction with the environment: Some tasks, e.g. the elastic movement of the knee joint when the foot hits the ground in running, can be taken over by the – elastic – materials, and their trajectories do not need to be explicitly controlled. By morphology we mean the form and structure of an organism and its parts, including the physical nature of the sensors and their distribution. We discuss materials separately, as they play an extraordinary role in agent design.

There is another aspect of ecological balance, namely that there should be a match in the complexity of the sensory, motor and (neural) control systems. Many robotic systems are “unbalanced” in the sense that they are built of hard materials and electrical motors, and thus the control requires an enormous amount of computation. Robot vision systems are also often unbalanced as they are largely algorithmic and do not exploit morphological properties. For example, natural systems don’t have cameras but retinas that perform some kind of morphological computation by their non-homogeneous arrangement of the light-sensitive cells. Moreover, generally speaking retinas perform an enormous amount of computation right at the periphery so that the signals that are passed on, are already highly processed. Artificial retinas have been around since the mid-80s (e.g. Mead, 1989), but they are still not widely used in the field. Moreover, vision or perception in general is not a matter of mapping inputs to internal representation, but of sensory-motor coordination, requiring a complex motor system as well. While initially it might seem that taking the motor system into account as well in perception would make the problem harder, when viewed in an ecological context, many problems might in fact be simplified, as demonstrated by the field of active vision or animate vision (e.g. Ballard, 1991). In animate vision, the ability of the agent (the vision system) to move is exploited to make the vision task easier. The development of vision systems, which includes the development of retinas, remains a big challenge. And these vision systems must not be developed in isolation, but in the context of multi-modal systems (see also below, achieving higher level intelligence).

Recently, it has been demonstrated that by exploiting the intrinsic dynamics of an agent, the complexity of the control system can be substantially reduced (e.g. Collins et al., 2001; Iida and Pfeifer, 2004a, b; Wisse and Frankenhuyzen, 2003; Yamamoto and Kuniyoshi, 2001), as articulated in the principle of ecological balance. Thus, in order to achieve rapid locomotion, but also motion in general, material properties must be exploited. In order to achieve real progress, artificial muscles, tendons, and flexible joints must be developed which represents a big engineering challenge. Big

strides in this direction have been made by Rudolf Bannasch and his colleagues (Boblan et al., 2004).

Behavior in general requires sensory-motor coordination that again, in natural systems, is achieved by a subtle interplay of morphology (of the sensory and motor systems), materials, control, and interaction with the environment. While the design principles of Pfeifer et al. (in press) do provide intuitions, they are only qualitative in nature. What is needed now, and this is a big challenge, is a more quantitative approach. While it is relatively straightforward to quantify sensory data and to estimate the amount of computation in a controller, little research has been done on quantifying morphology and materials in computational terms. Finding a common currency which is required for a theoretical and quantitative understanding, is an important research issue as it will connect the computational effort (or control) with the contributions of physical, i.e. non-computational aspects of the system (for quantitative research in the field of sensory-motor coordination that will be relevant for these issues using methods from information theory and statistics, see, e.g. Sporns and Pegors, 2004; te Boekhorst et al., 2003) (Lungarella and Pfeifer, 2001). Lichtensteiger (2004), for example, demonstrated how the pre-processing function performed by the morphological arrangement of facets in an insect (or robot) eye, can be measured quantitatively and how a particular arrangement influences learning speed.

In general, there is a definite need for more quantitative methods in order to turn the field into a true scientific discipline. Gaussier et al., for example (Gaussier, et al., 2004) developed a formalism in the form of an algebra for cognitive processes based on the idea of perception-action coupling in autonomous agents. They apply the formalism to demonstrate how facial expressions can be learned and that there is no need to postulate innate mechanisms. Other examples of quantification will be discussed in the section on development.

While we must move towards more quantitative methods, there is a certain danger involved: Because of the limitations of formal description, there tends to be a focus on isolated, well-formalizable areas, as we know it from the field of classical robotics and control theory. For example, there is a lot of formal work on path planning and inverse kinematics which lends itself more readily to a formal treatment than, for example, locomotion of complex systems involving materials with different kinds of properties and many degrees of freedom. Formalizing the latter represents a big challenge.

From an engineering perspective, in addition to the materials of the motor system, there are challenges concerning the various sensory modalities: haptics for example, is a very fundamental and rich modality in natural organisms. But the technology is, compared to natural systems, very underdeveloped: low resolution, hard, non-bendable materials, pressure only. However, there are exciting developments towards overcoming these limitations, as illustrated by the soft robotic fingertip with randomly distributed sensors for measuring slip and texture by Hosoda (2004). The development of skin-sensors by which the entire body can be covered represents a big challenge, not so much for artificial intelligence, but for the material sciences, similar to the issue of artificial muscles. At the moment, this is a significant bottleneck: better materials would almost certainly entail a quantum leap in artificial intelligence.

Achieving higher level intelligence

The term “higher level” intelligence is used to designate behavior that is not purely sensory-motor, such as problem solving and reasoning, or generally thinking, natural language, emotion, and consciousness. Note that there is a frame-of-reference issue here: when we say “not purely sensory-motor” it is not really clear whether we are referring to behavior or mechanism. Inspection of the mechanisms underlying so-called non-sensory motor or cognitive behavior yields that almost universally the sensory and motor systems will be involved since in natural systems brains are intrinsically intertwined with embodiment and cannot clearly separated (e.g. Thelen and Smith, 1994). While it is possible in principle to “hand design” agents (see Table 1) endowed with higher level intelligence, all efforts to date have been met with only very limited success. One of the big unresolved issues to date is the one of symbol processing: How is it possible that humans have the capability for symbol processing? More precisely we would have to ask how it is possible that humans can behave in ways that it makes sense to describe their behavior as “symbolic”, irrespective of the underlying mechanisms, which might involve explicit symbol processing or not. The question is very broad and of general importance: it is about how organisms can acquire meaning, how they can learn about the real world, and how they can combine what they have learned to generate symbolic behavior, a problem known as the “symbol grounding problem.”. There is general agreement that learning will make substantial contributions towards a solution. However, learning alone will not suffice – embodiment must be taken into account as well.

Drawing inspiration from nature, a consensus has emerged that a productive approach might be to mimic at some level of abstraction a developmental process. Development, in contrast to learning, also incorporates growth and maturation of the organism. There is a vast literature on machine learning that might be potentially relevant here for solving the symbol grounding problem, but also for development in general. The book “Re-thinking innateness” has been viewed as a kind of landmark publication, employing a connectionist modeling approach (Elman et al., 1996). While a lot of ideas can be taken from this book, the approach does not deal with embodiment. This is the case with most of the machine learning literature.

As indicated earlier, the impact of taking embodiment into account can hardly be over-estimated. For example, there is the big challenge of general perception in the real world: How come we can recognize objects or faces under large variations of distance, orientation, partial occlusion, and lighting conditions? Again, many people seem to agree that a developmental approach might be useful. One of the basic issues is the fact that agents in the real world do not receive neatly structured input vectors – as is assumed in most simulation studies – but there is a continuously changing stream of sensory stimulation which strongly depends on the agent’s current behavior. One way to deal with this issue is by exploiting the embodied interaction with the real world: Through the – physical – interaction with the environment, the agent induces or generates sensory stimulation (e.g. Pfeifer and Scheier, 1999). The thus generated stimulation will typically be more structured, and will contain correlations within and between sensory channels that greatly facilitate the problem of focusing on the relevant stimulation and is in fact the enabler for learning (Lungarella and Pfeifer, 2001;

Sporns and Pegors, 2004). A very simple example is grasping and centering which stabilizes and normalizes the visual stimulation of an object on the retina, and at the same time produces correlated haptic and proprioceptive stimulation. This issue is covered in the principle of sensory-motor coordination which may be an important constituent in bootstrapping perception. Achieving general purpose, flexible and adaptive perception in the real world is certainly one of the very grand challenges. This is one of the big research topics in the field of “developmental robotics” or “cognitive robotics” that has recently picked up a lot of momentum. It has been suggested that the principle of sensory-motor coordination should be called more generally the principle of information self-structuring because the agent himself (or itself) interacts in particular ways with the environment to generate proper sensory stimulation.

Now the goal of this new field is not only perception, but development in general. An important direction is and has been imitation learning that seems to play a key role. This research has been inspired by the discovery of mirror neurons in the 1990s (e.g. Dipellegrino et al., 1992; Fadiga et al., 2000; Gallese et al., 1996) which demonstrated that motor and sensory systems are very closely intertwined in the brain. Designing and building a system capable of a wide range of imitation behaviors is certainly another one of the big challenges. Important first steps have demonstrated the in-principle feasibility of this approach (e.g. Kuniyoshi et al., 2004; Jansen et al., 2004; Yoshikawa et al., 2004). Robots will no longer have to be programmed, but the skills they should acquire can simply be demonstrated. While this ability will certainly improve the sensory-motor behavior of agents, the hope is that it will also contribute to the development of social behavior, and language and communication abilities. For a review of the research in developmental robotics, see Lungarella et al. (2004). One of the challenges for the research on imitation is that direct copying is not possible, because the caregiver has a morphology that considerably differs from the one of the baby, i.e. certain perceptual generalizations will have to be made by the baby in order to interpret the caregiver’s action. Over the last few years, there has been increasing consensus that joint attention plays a key role in learning and social development, a topic now being studied in developmental robotics (e.g. Nagai et al., 2003).

Let us briefly discuss a few additional grand challenges in development, acquisition of natural language, consciousness, emotion, and motivation. First steps toward acquisition of natural language, acquisition of a joint vocabulary, has been demonstrated in Luc Steels’s ingenious “Talking Heads” experiment. Steels also did some preliminary work on acquisition of syntax, but there is a long way to the final goal of complete natural language development.

Consciousness has always been considered as something like the ultimate criterion for true intelligence. An elusive and fascinating topic that has attracted quite a bit of attention in the field of embodied artificial intelligence. Owen Holland is also having a stab at the future of embodied artificial intelligence and asks the question of whether we will be able to achieve machine consciousness (Holland, 2004). A topic often discussed in investigating consciousness – and in building machine consciousness, are the so-called qualia. Qualia are the subjective sensory qualities like “the redness of red” that accompany our perception. Qualia symbolize the explanatory gap that exists between the subjective qualities of our perception and the physical brain-body system whose states can, in principle, be measured objectively. In our terminology, qualia are

closely related to embodiment, to the physical, material, and morphological structure of the sensory systems.

Emotions, another highly controversial topic, also relate to the issue of consciousness and the development of emotional machines is also a topic of interest (for a partial review of an embodied perspective, see e.g. Pfeifer, 2000) . Last but not least, a topic that anyone interested in intelligence and especially development will have to deal with is why an agent does anything in the first place? Why should it learn new things? This question is especially relevant if there are rich task environments with many behavioral possibilities. A chess computer only has one task, i.e. to make the next move, whereas in the real world there are always a host of possibilities – at least for those agents that we are potentially interested in (not for Braitenberg Type 1 vehicles). It is the entire issue of motivation, a topic with an enormous history. Luc Steels and Frederic Kaplan in this volume present two simple but powerful and highly plausible general solutions (Steels, 2004; Kaplan and Oudeyer, 2004). These are all fundamental questions of cognitive science.

In order to make development work, a number of engineering challenges must be resolved. From developmental studies it is known that sensory-motor coordination underlies much of concept development. This requires on the one hand the development of proper actuators: upper torso with head/neck, and arms with hands. Many researchers work with torsos only, but given the importance of locomotion for cognitive development, it would be desirable to have complete agents capable of walking freely in their environments. To date most robots are specialized, either for walking or other kinds of locomotion purposes, or for sensory-motor manipulation, but rarely are they skilled at performing a wide spectrum of tasks. This is due to conceptual and engineering limitations. Actuator technology is a major problem as today mostly electrical motors are employed, whereas – as argued earlier – artificial muscles would be more desirable. Skin sensors for the fingertips, but also for covering the entire body, would be essential for building up something like a body image, and ultimately to bootstrap cognition. Huge transdisciplinary efforts between engineering, biomechanics, and material science will be required to make progress here.

Note that although most people in developmental or cognitive robotics are interested in humanoids, this is by no means the only path. A developmental perspective can be beneficial for all kinds of animal studies.

High-level intelligence cannot only be achieved using a developmental approach, but also, at least theoretically, by means of evolutionary methods. We will discuss them in the subsequent paragraph, but given the state-of-the-art in artificial evolution, we will have to resort to more direct methods such as hand design or developmental approaches for the time being.

Automated design methods (artificial evolution and morphogenesis)

Using artificial evolution for design has a tradition in the field of evolutionary robotics (e.g. Nolfi and Floreano, 2001). The standard approach is to take a particular robot and use an evolutionary algorithm to evolve a controller for a particular task. However, if we want to explore morphological issues, and if we want to design entire agents

rather than controllers only, we have to devise powerful methods capable of handling these issues. Floreano et al. (2004) provide an excellent overview of the field with many illustrations and experiments.

Because of the many parameters and design considerations involved, automated methods must be employed because humans will no longer be able to “hand design” all aspects of such systems. There is the morphology of the body, the materials, the neural control, the interaction with the environment, and there is the possibility of having several agents, perhaps simpler ones, perform the task collectively. For individual organisms, there have been some initial successful attempts at designing systems by evolutionary means, the main approaches being the parameterization with recursive encoding (e.g. Sims, 1994; Lipson and Pollack, 2000), and those where ontogenetic development is based on abstract models of genetic regulatory networks using cell-to-cell signaling mechanisms (Eggenberger, 1997, 1999; Bongard, 2002, 2003; Bongard and Pfeifer, 2001; Banzhaf, 2004). The advantage of genetic regulatory networks is that they incorporate less of a designer bias and that they allow for incorporation of interaction with the environment during ontogenetic development, developmental plasticity (Bongard, 2003). Moreover, because they encode growth processes, they also, in some sense, contain the mechanisms for self-repair, an essential property of natural systems.

There are a number of challenges, here. First, it is the further development of models genetic regulatory networks to grow creatures of arbitrary complexity and to make the evolution open-ended in the sense that not only the parameters of the genetic regulatory networks can be manipulated, but that the mechanisms themselves are under evolutionary control. Moreover, understanding and controlling the highly involved complex dynamics of genetic regulatory networks will require a lot of research (see Bongard, 2003; Eggenberger, 1999; and Banzhaf, 2004, for some preliminary pertinent research). An important aspect will be the understanding of the emergence of hierarchical structures and modularity of the phenotypes (see also Floreano et al., 2004). Second, the physics-based simulation models need to be augmented to allow for more sophisticated agent-environment interactions. Also, deformable, flexible materials, additional sensors such as “skins” for covering the entire body, or olfaction, as well artificial muscles should be accounted for. Third, along these lines, the task environments must be made much more complex in order to put these design methods to a real test. In this way, we might be able to observe and better understand phenomena of centralization of neural substrate, i.e. the formation of brains. Eventually we might be able to see not only exploitation of physical interaction constraints, but also social ones. Whether the mechanisms of simulated genetic regulatory networks will in fact scale to very complex organisms capable of sophisticated social interaction, is an open question. The grand challenge remains to evolve truly complex creatures capable of communication, language, high-level cognition, and – perhaps – consciousness. Several orders of magnitude of scale will have to be bridged in the process, from molecules to macroscopic organisms. To what extent physically realistic simulations will be sufficient for this purpose, or whether evolution actually must happen in the real world with its indefinite richness, is a deep and currently unanswered issue.

This evolutionary level, designing the evolutionary mechanisms as well as the developmental processes based on genetic regulatory networks, might in fact provide a

proper level of formalization of ecological balance. While it is indeed hard to find a common currency for trading computation for materials and morphology, it might turn out to be much easier to formally specify the developmental processes as encoded in the genome. This is because, at this stage, it is still undecided how the tasks will be distributed to control, materials, and morphology for a particular task-environment.

Moving into the real world

The last grand challenge that we would like to discuss here concerns very generally speaking the “move into the real world.” The first significant step in this direction has been the introduction of the notion of embodiment and the insight that true intelligence always requires the interaction with the real world. Embodied artificial intelligence is based on this idea. Building intelligent robots, i.e. robots capable of performing a wide range of tasks, is, as we have argued throughout this paper, hard enough, and the robots we currently are capable of building are not to our satisfaction, and so building robots per se remains a grand challenge in the field.

In designing higher-level intelligence we identified developmental approaches as a potentially suitable method. Development requires growth processes that we can currently only simulate. But there are some tricks that can be applied to make development somewhat more realistic vis-à-vis the real world. One possibility is to start with high-resolution, high-precision systems with many degrees of freedom. Growth, at least in some respects, can then be “simulated” by constraining the systems initially, freezing degrees of freedom, and simulating low resolution, for example, of the vision sensor in software by applying certain kinds of filters. These constraints can successively be released which in some sense reflects an organism’s maturational processes (Gómez et al., 2004).

However, biological organisms actually do grow in the real world by means of cell division and cell differentiation, a process that may in fact be essential for the emergence of cognition. Developing growing structures in the real world is one of the great engineering challenges that will require the cooperation of material scientists, engineers, molecular and developmental biologists, and nanotechnology experts. These are, by the way, all disciplines that are not normally associated with artificial intelligence.

If artificial evolutionary processes are not only to be simulated in a computer but performed in the real world, we will need growth processes as well. As mentioned earlier, it is not clear to what extent physics-based simulations will be sufficient for scalable artificial evolution, and to what extent evolution has to rely on processes in the real world. First steps in performing artificial evolution in the real world have been taken already in the 1960s by Ingo Rechenberg who evolved optimal shapes of fuel pipes by actually configuring the physical system “designed” by the evolutionary algorithm (an evolution strategy) and measuring the performance on the real fuel pipe system (Rechenberg, 1973). Another example is the work by Adrian Thompson at the University of Sussex who used FPGAs to test the circuits evolved using a genetic algorithm (Thompson, 1996). FPGAs, in contrast to microprocessors, rather than making a digital simulation of a circuit, actually configure a physical circuit. The

results achieved are truly amazing and provides a glimpse at the power of evolution in the real world.

A major step is taken by researchers in the EU-funded PACE (Programmable Artificial Cell Evolution) project by John McCaskill of the Ruhr University Bochum, in Germany, where the goal is to evolve an artificial cell in a chemical laboratory. Using micro-fluidic arrays, carefully controlled chemical reactions can be induced so that cells can be formed and their metabolisms influenced in precise ways. Part of the evolution will be performed in simulation and part in the real world. The goal is to evolve self-replicating cells in the laboratory, an enormous challenge. If successful, this would enable us to perform artificial evolution in the real world and thus we could generate any kind of structure required for performing a particular task. Because the cells can divide we would have actual growth processes in the real world. Some people like Ray Kurzweil believe that nanotechnology will be the key to engineer growth in the real world. Whether this will materialize we will only know in the future.

Cyborgs could also be viewed as a way to “move into the real world”: rather than constraining the neural substrate to function in a dish in isolation, it is connected to either a simulation or to a robot that behaves in the real world and sends its sensory signals back to the neural tissue in the dish (Bakkum et al., 2004). Coupling biological neural tissue to a real world artifact opens up entirely new avenues in man-machine interaction. This research in itself bears many great challenges, the general issue of coupling biological and technical substrate. On the one hand, we can expect to learn something about neural functioning, and on the other we might, in the future, be able to better understand how to control robots by observing the natural neurons. Medical applications in prosthetics (e.g. Yokoi et al., 2004), are of course obvious candidates for practical applications.

Finally, coming back to the research on self-repair, self-assembly, and self-reconfiguration discussed in the “Landscape” section, a big challenge, conceptually and from an engineering perspective, is the development of such systems in the real world. Again, while simulation of processes of self-repair, for example, represents a challenge and is far from being straight-forward, the ultimate challenge will be the transfer to the real world. Murata and his collaborators (2004) have demonstrated first ideas using modular robotic systems.

4. Conclusions, the future, and applications

The challenges outlined are big challenges and we must not expect to reach them in the near future. However, it is important to keep the long-term visions in mind when thinking about the next steps. The difficulty of research in any field, but in particular in artificial intelligence, is to map the big visions and challenges onto concrete, doable steps. We have also tried to outline what researchers in the field are currently attempting to do and what they are planning for the near future. And the papers presented in this volume provide an excellent starting point.

Let us now return to the initial question of what thinking has to do with walking – the symbol grounding problem – and reflect on how the challenges outlined in the

paper will contribute to this question which metaphorically summarizes the goals of embodied artificial intelligence. In the early phases of embodied artificial intelligence, many people were working on navigation and orientation out of a conviction that locomotion and orientation are somehow the underlying driving forces in the development of cognition, in the evolution of the brain. This is corroborated by the question asked by the famous Oxford neuroscientist Daniel Wolpert “Why don’t plants have brains?”. And he suggested that the answer might actually be quite simple: “Plants don’t have to move!” Because of the “embodied turn”, researchers started working with robots, and because they were readily available and easy to use, wheeled robots were the tools of choice. Navigation in the real world is a challenging problem and there has been much exciting research in robotics in general (e.g. Bellot et al., 2004, who introduce the new method of Bayesian Programming) and in biologically inspired approaches in particular (e.g. Hafner, 2004). While there was a lot of progress – researchers were forced to deal with the intricacies of the interaction with the real world, such as noise, imprecisions, change, unpredictability – there were also some intrinsic problems with the approach. Remember that one of the aspects of the principle of ecological balance is the match in complexity of sensory, motor, and neural systems. Because it is easy to put a high-resolution camera on a robot, and because wheeled robots only have few degrees of freedom of actuation, many experimental designs were “unbalanced”: complex sensory systems, very simple motor systems. As a result of these unbalanced designs, these systems had a relatively uninteresting physical dynamics. One implication is that the algorithms used for control were largely arbitrary: Even though they were mostly biologically inspired, they were arbitrary with respect to the robot’s own dynamics; one algorithm can be exchanged by another, achieving essentially the same behavior. Something was missing and many suspected that this is a complex sensory-motor level with an interesting and rich dynamics.

As a consequence a number of researchers started working on complex body dynamics (e.g. Kuniyoshi et al, 2004; Iida and Pfeifer, 2004a; Proc. of the Int. Workshop on Adaptive Motion in Animals and Machines, AMAM-2003). This shift was interpreted by critics but also by people sympathetic to these developments, as a move away from the goal of understanding and building cognitive systems. However, and this is one of the big insights from embodied artificial intelligence, the exact opposite was the case: It turned out that a rich complex body dynamics is the foundation, the prerequisite for something like symbol processing to develop (see, e.g. Okada et al., 2003; Iida and Pfeifer, 2004b; Kuniyoshi et al., 2004). So what happened is that what seemed like a deviation from the road to cognition, turned out to be necessary. This view is also compatible with Núñez (2004) who argues that even very abstract mathematical concepts have their origins, are grounded, in our embodiment which provides the basis for metaphors. Because these metaphors must be sufficiently rich for bootstrapping interesting concepts, the embodiment must reflect this richness. Of course, at the moment, this is all speculation that must be corroborated by many experiments. But at the risk of being entirely wrong, let us speculate a little further.

There is another, unexpected idea that emerges from this research. The question of symbol grounding always entails the question of how it is possible that something like discrete symbol processing can emerge from a completely continuous dynamical system, such as a human. Rich, complex dynamics also implies many attractor states and

transitions between them. Attractor states are, within the continuous dynamics, objectively identifiable, discrete states, that can, of course, also be identified by the agent itself (or himself), given the proper neural system. Once identified, the agent can start using them, for example, for planning purposes (e.g. Okada et al., 2003; Kuniyoshi et al., 2004). It is interesting to note that a complex intrinsic sensory-motor dynamics implies that the neural control is no longer arbitrary, but has to be “in tune” with the physical substrate, quite in contrast to wheeled robots. Ishiguro and his colleagues (2004) have provided a beautiful demonstration, theoretically and in a robot case study, of how control and body dynamics in a complex agent have to be coupled. If coupled properly, control is not only simpler, but the entire system tends to be more energy-efficient. Lungarella and Berthouze (2004) in a robotics case study convincingly demonstrate that a judicious – non-arbitrary – choice of parameters coupling the neural and body dynamics facilitates the acquisition of motor skills in a developing organism. Whether these ideas on dynamics will ultimately lead to high-level cognition or to conscious agents, whether in this way we can achieve the goals set out by Holland (2004), is an entirely open question.

Tom Ziemke in his contribution (2004) quotes from Gerald Edelman “It is not enough to say that the mind is embodied: one has to say how.” (Edelman, 1992). Bootstrapping it from complex body dynamics might be part of the answer.

In their current state, evolutionary studies are, for the time being, restricted to providing ideas on the distribution of morphology, materials, control, and interaction with the environment. More varied and taxing task environments will be necessary to investigate agents with more complex sensory-motor dynamics on top of which cognition can bootstrap. But some of recent approaches demonstrate definite progress in this direction (e.g. (Bongard, 2003)). However, as alluded to in the previous section, in order to achieve truly complex organisms, it may be necessary to couple the artificial evolutionary process to the real world.

To conclude, just few words about applications. While the classical approach has created many applications in terms of clever algorithms that are now widely used, the embodied approach seems to be more limited. The major applications have been in the entertainment and educational areas. As this paper demonstrates, the field is just beginning to develop a basic understanding and there are many big challenges lying ahead. We could also add a challenge, namely to exploit these technologies for practical applications in industry, the environment, and services for the benefit of society.

Research on humanoid robots has an interesting side-effect, so to speak. Humanoids require the development of sophisticated body parts, legs, arms, hands, etc., that can potentially be used, at least to some extent, as prosthetic devices. The fascinating research by Yokoi et al. (2004) and by Boblan et al. (2004) points in this direction. The ground breaking research by Potter and his co-workers (Bakkum et al., 2004) might eventually be employed for interfacing these devices smoothly with humans – an additional intriguing perspective.

As outlined in the section of ubiquitous computing, a better understanding of embodied intelligence will lead to many applications in terms of so-called embedded systems, i.e. systems that autonomously interact with the real world, not only through sensing, but also by influencing the world without human intervention. These systems are not robots in the restricted sense of the word (they are very different from human-

oid robots, for example), but they have many of their characteristics in terms of intelligent, autonomous interaction with the environment. These kind of systems, also called “robotic devices” are already present in many technical applications (cars, airplanes, household appliances, elevators, etc.), but by augmenting their “intelligence”, so to speak, many more applications will become possible. This way, the ideas that embodied artificial intelligence has spurred will spread to numerous scientific and technological areas for the benefit of society.

Acknowledgments

We would like to thank the scientific director of the International Conference and Research Center for Computer Science, Prof. Reinhard Wilhelm, for suggesting this conference, and the Swiss National Science Foundation for supporting the research presented in this paper, grant # 20-68198.02 (“Embodied Artificial Intelligence”). We would also like to thank the members of the Artificial Intelligence Laboratory of the University of Zurich for numerous stimulating discussions on this topic. Credit also goes to Max Lungarella for his many thoughtful comments on this paper.

References

- Bakkum, D.J., Shkolnik, A.C., Ben-Ary, G., Gamblen, P., DeMarse, T.B., and Potter, S.M. (2004). Removing some ‘A’ from AI: Embodied cultured networks (this volume)
- Ballard, D. (1991). Animate vision. *Artificial Intelligence*, 48, 57-86.
- Banzhaf, W. (2004). On evolutionary design, embodiment, and artificial regulatory networks (this volume).
- Boblan, I., Bannasch, R., Schwenk, H., Miertsch, L., and Schulz, A. (2004). A human like robot hand and arm with fluidic muscles: Biologically inspired construction and functionality. (this volume)
- Bellot, D., Siegwart, R., Bessière, P., Tapus, A., Coué, C., and Diard, J. (2004). Bayesian modeling and reasoning for real-world robotics: Basics and examples (this volume).
- Blickhan, R., Wagner, H., and Seyfarth, A. (2003). Brain or muscles?, *Rec. Res. Devel. Biomechanics*, 1, 215-245.
- Bonabeau, E., Dorigo, M., and Theraulaz, G. (1999). *Swarm intelligence: from natural to artificial systems*. New York, N.Y.: Oxford University Press.
- Bongard, J.C. (2003). Incremental approaches to the combined evolution of a robot’s body and brain. Unpublished PhD thesis. Faculty of Mathematics and Science, University of Zurich.
- Bongard, J.C. (2002). Evolving modular genetic regulatory networks. In *Proc. IEEE 2002 Congress on Evolutionary Computation (CEC2002)*. MIT Press, 305-311.
- Bongard, J.C., and Pfeifer, R. (2001). Repeated structure and dissociation of genotypic and phenotypic complexity in artificial ontogeny. In L. Spector et al. (eds.). *Proc. of the Sixth European Conference on Artificial Life*, 401-412.
- Brooks, R. A. (1991). *Intelligence Without Reason*. Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91), pp. 569–595.
- Brooks, R.A., and Stein, L.A. (1993). Building brains for bodies. Memo 1439, Artificial Intelligence Lab, MIT, Cambridge, Mass.

- Collins, S.H., Wisse, M., and Ruina, A. (2001). A three-dimensional passive-dynamic walking robot with two legs and knees. *The International Journal of Robotics Research*, 20, 607-615.
- Dipellegrino G, Fadiga L, Fogassi L, Gallese V, Rizzolatti, G (1992). Understanding motor events - a neuro-physiological study. *Exp Brain Res* 91: 176-180.
- Edelman, G.E. (1992). *Bright air, brilliant fire. On the matter of the mind*. New York: Basic Books.
- Eggenberger, P. (1997). Evolving morphologies of simulated 3d organisms based on differential gene expression. In: P. Husbands, and I. Harvey (eds.). *Proc. of the 4th European Conference on Artificial Life*. Cambridge, Mass.: MIT Press.
- Eggenberger, P. (1999). Evolution of three-dimensional, artificial organisms: simulations of developmental processes. Unpublished PhD Dissertation, Medical Faculty, University of Zurich, Switzerland.
- Elman, J.L., Bates, E.A., Johnson, H.A., Karmiloff-Smith, A., Parisi, D., and Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development*. Cambridge, Mass.: MIT Press.
- Epstein, J.M. and Axtell, R.L. (1996). *Growing artificial societies: social science from the bottom up*. Cambridge, Mass.: MIT Press.
- Fadiga L, Fogassi L, Gallese V, Rizzolatti G (2000) Visuomotor neurons: Ambiguity of the discharge or 'motor' perception? *Int J Psychophysiol* 35: 165-177.
- Ferber, J. (1999). *Multi-agent systems. Introduction to distributed artificial intelligence*. Addison-Wesley.
- Floreano, D., Mondada, F., Perez-Uribe, A., and Roggen, D. (2004). Evolution of embodied intelligence (this volume).
- Gallese, V., Fadiga, L., Fogassi, L., and Rizzolatti G. (1996). Action recognition in the premotor cortex. *Brain* 119: 593-60.
- Gaussier, P., Prepin, K., and Nadel, J. (2004). Toward a cognitive system algebra. Application to facial expression learning and imitation (this volume).
- Gómez, G., Lungarella, M., Eggenberger Hotz, P., Matsushita, K. and Pfeifer, R. (2004). Simulating development in a real robot: on the concurrent increase of sensory, motor, and neural complexity. The 4th annual workshop of Epigenetic Robotics (EPIROBOT04), (in press).
- Hafner, V. (2004). Agent-environment interaction in visual homing (this volume).
- Hara, and R. Pfeifer (eds.) (2003). *Morpho-functional machines: the new species – designing embodied intelligence*. Tokyo: Springer-Verlag.
- Hara, F., and Pfeifer, R. (2000). On the relation among morphology, material and control in morpho-functional machines. In Meyer, Berthoz, Floreano, Roitblat, and Wilson (eds.): *From Animals to Animats 6. Proceedings of the sixth International Conference on Simulation of Adaptive Behavior 2000*, 33-40.
- Holland, O. (2004). The future of embodied artificial intelligence: Machine consciousness? (this volume).
- Hosoda, K. (2004). Robot finger design for developmental tactile interaction. Anthropomorphic robotic soft fingertip with randomly distributed receptors (this volume).
- Iida, F. and Pfeifer, R. (2004a) "Cheap" Rapid locomotion of a quadruped robot: Self-stabilization of bounding gait. F. Groesn et al. (eds.). *Intelligent Autonomous Systems 8*. IOS Press, 642-649.
- Iida, F., and Pfeifer, R. (2004b). Self-stabilization and behavioral diversity of embodied adaptive locomotion (this volume).
- Ishiguro, A., and Kawakatsu, T. (2003). How should control and body systems be coupled? A robotic case study (this volume).
- Janssen, B., de Boer, B., and Belpaeme, T. (2004). You did it on purpose! Towards intentional embodied agents (this volume).

- Kaplan, F., and Oudeyer, P.-Y. (2004). Maximizing learning progress: an internal reward system for development (this volume).
- Kubow, T. M., and Full, R. J. (1999). The role of the mechanical system in control: a hypothesis of self-stabilization in hexapedal runners, *Phil. Trans. R. Soc. Lond. B*, 354, 849-861.
- Kuniyoshi, Y., Yorozu, Y., Ohmura, Y., Terada, K., Otani, T., Nagakubo, A., and Yamamoto, T. (2004). From humanoid embodiment to theory of mind (this volume).
- Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, 30, 39-64.
- Lenat, D., Prakash, M., and Shepher, M. (1986). CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. *AI Magazine*, vol. 6, issue 4, 65-85.
- Langton, C. G. (1995). *Artificial life: an overview*. Cambridge, Mass.: MIT Press.
- Lipson, H., and Pollack J. B. (2000), Automatic design and manufacture of artificial life forms. *Nature*, 406, 974-978.
- Lichtensteiger, L. (2004). The need to adaptv and its implications for embodiment (this volume).
- Lungarella, M., and Berthouze, L. (2004). Robot bouncing: On the synergy between neural and body dynamics (this volume).
- Lungarella, M. and Pfeifer, R. (2001). Information-theoretic analysis of sensory-motor data. In *Proc. of the IEEE-RAS International Conference on Humanoid Robots*, 245-252.
- Lungarella, M., Metta, G., Pfeifer, R. and Sandini, G. (2003). Developmental robotics: a survey. *Connection Science*, 15 (4), 151-190.
- Mead, C.A. (1989). *Analog VLSI and neural systems*. Reading, Mass.: Addison-Wesley.
- Murata, S., Kamimura, A., Kurokawa, H., Yoshida, E., Tomita, K., and Kokaji, S. (2004). Self-reconfigurable robots: platforms for emerging functionality (this volume).
- Nagai, Y., Hosoda, K., and Asada, M. (2003). Joint attention emerges through bootstrap learning, *Proc. of the 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2003)*, 168-173.
- Nolfi, S. and Floreano, D. (2001). *Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- Núñez, R. (2004). Do real numbers really move? The embodied cognitive foundations of mathematics (this volume).
- Okada, M., Nakamura, D., and Nakamura, Y. (2003). On-line and hierarchical design methods of dynamics based information processing system. *Proc. of the 2003 IEEE/RJS Int. Conference on Intelligent Robots and Systems*, 954-959.
- Pfeifer, R. (2000). On the role of embodiment in the emergence of cognition and emotion. In H. Hatano, N. Okada, and H. Tanabe (eds.). *Affective minds*. Amsterdam: Elsevier, 43-57.
- Pfeifer, R., Iida, F., and Bongard, J. (2004). *New robotics: design principles for intelligent systems*. *Artificial Life* (in press).
- Pfeifer, R., and Scheier, C. (1999). *Understanding intelligence*. Cambridge, Mass.: MIT Press.
- Rechenberg, I. (1973). *Evolution strategies: optimization of technical systems with principles from biological evolution* (in German). Stuttgart, Germany: Frommann-Holzboog.
- Sims, K. (1994a). Evolving virtual creatures. *Computer Graphics*, 28, 15-34.
- Sporns, O., and Pegors, T.K. (2004). Information-theoretical aspects of embodied artificial intelligence (this volume).
- Steels, L. (2001). Language games for autonomous agents. *IEEE Intelligent Systems*, Sept/Oct issues.
- Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Sciences*, 7 (7), 308-312.
- Steels, L. (2004). The autotelic principle (this volume).

- te Boekhorst, R., Lungarella, M., and Pfeifer, R. (2003). Dimensionality reduction through sensory-motor coordination. Proc. of the 10th Int. Conf. on Neural Information Processing (ICONIP'03), p.496-503, LNCS 2174.
- Thelen, E., and Smith, L. (1994). A dynamic systems approach to the development of cognition and action. Cambridge, Mass.: MIT Press.
- Thompson, A. (1996). Silicon evolution. In J.R. Koza et al. (Eds.). Genetic Programming 1996: Proc. of the First Annual Conference, Cambridge, Mass.: MIT Press, 444-452.
- Webb B. and Consi R. C. (2000). Biorobotics -Methods & application-, Cambridge, Mass.: MIT Press.
- Weiser, M. (1993). Hot topics: Ubiquitous computing, IEEE Computer.
- Wisse, M and Frankenhuyzen, J.van, (2003) Design and Construction of MIKE; a 2D autonomous biped based on passive dynamic walking, Proceedings of the 2nd International Symposium on Adaptive Motion of Animals and Machines, Kyoto, March.4-8, 2003.
- Yamamoto, T. and Kuniyoshi, Y. (2001). Harnessing the robot's body dynamics: a global dynamics approach. Proc. of 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2001), pp. 518-525, Hawaii, USA.
- Yokoi, H. Arieta, A.H., Katoh, R., Yu, W., Watanabe, I., and Mruishi, M. (2004). Mutual adaptation in a prosthetic application (this volume).
- Yoshikawa, Y., Asada, M., and Hosoda, K. (2004). Towards imitation learning from a view point of an internal observer (this volume).
- Ziemke, T. (2004). Embodied AU as science: Models of embodied cognition, embodied models of cognition, or both? (this volume).