Embodied Artificial Intelligence
On the role of morphology and materials in the emergence of cognition*

Rolf Pfeifer
Artificial Intelligence Laboratory
Department of Information Technology, University of Zurich
Winterthurerstrasse 190, CH-8057 Zurich, Switzerland
pfeifer@ifi.unizh.ch, phone: +41 1 63 5 4320/31, fax: +41 1 635 68 09,
http://www.ifi.unizh.ch/~pfeifer

Abstract

In the cognitivistic paradigm, intelligence was viewed as algorithmic, an approach that has lead to many impressive results and applications. However, this view turned out to be too limited: If we are to understand natural forms of intelligence, embodiment must be taken into account. In this paper we explore the implications of embodiment by providing a number of case studies. One concept that we will investigate is "ecological balance", i.e. the interplay of morphology, materials, and neural processing. We develop a method of how this can be systematically investigated using artificial evolution and morphogenesis.

1. Introduction

The field of artificial intelligence has essentially two goals, understanding natural forms of intelligence and developing intelligent artifacts. For several decades, i.e. from the 50s until the mid-80s artificial intelligence was mostly concerned with algorithms, for example for playing chess, checkers (and other games), solving cryptarithmetic puzzles, logical inference, proof of mathematical theorems or natural language processing of written text. Viewing intelligent behavior at the level of algorithms is also called the cognitivistic paradigm. As is well-known, this approach has been highly successful: the victory of the chess playing computer "Deep Blue" over the world champion Garry Kasparov in 1997 in New York, is convincing testimony. Moreover, algorithms from the field of artificial intelligence have silently crept into many everyday applications:

---

* This is a slight edited version of a paper published in the proceedings of the workshop "The legacy of Grey Walter", Bristol, UK, 14-16 Aug., 2002.

whenever you turn on a computer or an appliance of sorts, there will be some pertinent algorithms involved. This includes search engines on the internet, text processing programs, speech-based directory information systems, sound systems, dishwashers, cars with their fuel injection systems, vacuum cleaners, and elevator control systems.
While successful in many ways, the approach has failed to contribute significantly to our understanding of natural forms of intelligence. Apparently, the latter differ in fundamental ways from the algorithmic kind. Over time, it became clear that intelligence was not so much a question of algorithms but of the interaction of an agent with the real world and the attention turned to embodiment (e.g. Brooks, 1991a, b). As a consequence, researchers started using robots as their workhorse. This change in orientation entails many new research issues that go beyond the level of algorithms.

Initially somewhat separate from artificial intelligence, the field of adaptive behavior has been growing rapidly during the last decade. Rather than trying to reproduce high-level cognition (thinking, reasoning, and abstract problem solving) directly, the focus is now on processes of adaptation and learning in the real world. This change in focus has lead to an interest in relatively simple biological types of intelligence. The argument was that if we are ever to understand higher levels of intelligence, for example, related to human natural language, we must first understand simple behaviors, as they are the precursors in evolutionary history. Adaptive behavior has attracted many researchers from artificial intelligence – the two disciplines are beginning to merge (this holds only for the biologically motivated branch of artificial intelligence, not the algorithmic one).

In this paper we investigate some of the consequences of the embodied approach to the study of intelligence. One of the important implications is that not only the neural system, but the entire agent, its morphology, and the materials from which it is constructed, are responsible for the agent's adaptive behavior. The relation between these aspects of an agent is called "ecological balance." In fact, the concept of "ecological balance" includes two different but related ideas. The first concerns the interplay of the sensory system, the motor system, and the neural substrate, the second the relation between morphology, materials, and control. In this paper we only deal with the second aspect (for references to both ideas, see, e.g. Hara and Pfeifer, 2000; Pfeifer, 1996; Pfeifer, 1999, 2000; Pfeifer and Scheier, 1999).

We proceed as follows. First, we introduce the notion of embodiment using a number of illustrations. We then discuss the notion of "ecological balance". This is followed by an argument of how this relation can be explored using artificial evolution and morphogenesis. Finally, we conclude with a discussion of what has been achieved and what the future prospects might be.

2. Embodiment – the interdependence of morphology and control

The goal of this section is to introduce the novel ideas that have been developed within the framework of embodied artificial intelligence. In particular we will show that embodiment means much more than simply "using a robot" or "having a body" – it requires an entirely new way of thinking, and it necessitates reflecting on the interaction with the real world; the latter is messy and not as neat as the world of the virtual machine. We start with a few comments on embodiment and then present a series of case studies.

2.1 Implications of embodiment

Embodiment has two main types of implications, physical and information theoretic. The former are concerned with physical forces, inertia, friction, vibrations, and energy dissipation, i.e. anything concerned with the (physical) dynamics of the system, the latter with the relation between sensory signals, motor control, and neural substrate. Rather than focusing on the neural substrate only, the attention is now on the complete organism which includes morphology (shape, distribution and physical characteristics of sensors and actuators, limbs, etc.) and materials. Clearly, the neural processing required for a particular task depends on embodiment since the latter delivers, so to speak, the raw material, the signals for the neural system to process. Similarly, the motor system has a particular dynamics that depends on morphology (body shape, limbs) and materials (e.g. of the muscle-tendon system) and this dynamics needs to be controlled or modulated by the neural system. And last but not least, through the interaction with the real world, the agent actively generates sensory stimulation which is why we often talk about sensory-motor coupling. Note that there is no one-way path from sensors to internal representations.

One of the surprising consequences is that often, problems that seem very hard if viewed from a purely computational perspective, turn out to be easy if the embodiment and the interaction with the environment are appropriately taken into account. For example, as we will show in a number of examples, given a particular task environment, if the morphology is right, the amount of neural processing required may be significantly reduced. Because of this perspective on embodiment, entirely new issues are raised and need to be taken into account. An important one "ecological balance" which we introduced earlier and that we will further investigate below.

Before we look at the case studies, let us review an example that demonstrates the central role of embodiment in behavior. Simon (1969) has used the metaphor of an ant to illustrate some basic principles of behavior. For instance, even though the trajectory of an ant on the beach might look complex to an observer, the rules - implemented in the neural substrate of the ant - might be very simple, something like: obstacle on right, turn left, obstacle on left, turn right. Given a beach where we find pebbles, twigs, rocks, and puddles, the trajectory of the ant will be a zigzag line. Let us now make a thought experiment. Let us increase the size of the ant by a factor of 1000 (even though this is biologically not plausible) and we let the ant loose in exactly the same location, on the same beach with the same behavioral rules, i.e. the same "neural substrate". What happens is that the trajectory of the ant will be much more straight than the one of the small ant. We left the neural "program" and the environment unchanged, we only modified the size - the morphology - and the behavior turned out to be completely different. This illustrates on the one hand that behavior cannot be reduced to internal mechanism only, and on the other that in order to understand the behavior of the ant, we must know the characteristics of the body into which the brain is embedded and the properties of the environment.

2.2 Case studies illustrating embodiment

In previous papers we have investigated in detail the effect of changing sensor morphology (e.g. Lichtensteiger and Eggenberger, 1999; Maris and te Boekhorst, 1996; Pfeifer, 2000; Pfeifer and Scheier, 1999). In this paper we focus on the motor system and present mainly case studies on walking robots.

The passive dynamic walker

Let us start with an example illustrating the relation between morphology, materials, and control. The passive dynamic walker which goes back to McGeer (1990a, b), illustrated in figure 1a, is a robot (or, if you like, a mechanical device) capable of walking down an incline without any actuation and without control. In other words, there are no motors and there is no microprocessor on the robot; it is brainless, so to speak. In order to achieve this task the passive dynamics of the robot, its body and its limbs, must be exploited. This kind of walking is very energy efficient and there is an intrinsic naturalness to it. However, its "ecological niche" (i.e. the environment in which the robot is capable of operating) is extremely narrow: it only consists of inclines of certain angles. Energy-efficiency is achieved because in this approach the robot is – loosely speaking – operated near one of its Eigenfrequencies. To make this work, a lot of attention was devoted to morphology and materials. For example, the robot is equipped with wide feet of a particular shape to guide lateral motion, soft heels to reduce instability at heel strike, counter-swinging arms to negate yaw induced by leg swinging, and lateral-swinging arms to stabilize side-to-side lean (Collins et al., 2001).

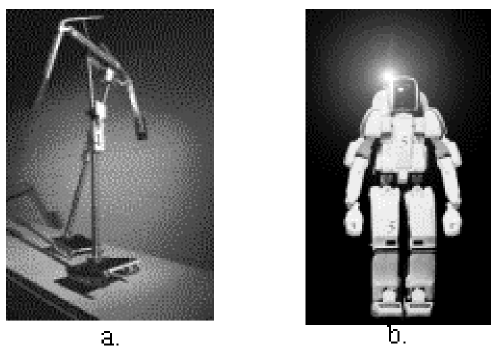

Figure 1: Two approaches to robot building. (a) The passive dynamic walker by Steve Collins (Collins et al., 2001) , (b) the Honda robot.

A different approach has been taken by the Honda design team. There the goal was to have a robot that could perform a large number of movements. The methodology was to record human movements and then to reproduce them on the robot which leads to a relatively natural behavior of the robot. On the other hand control – or the neural processing, if you like – is extremely complex and there is no exploitation of the intrinsic dynamics as in the case of the passive dynamic walker. The implication is also that the movement is not energy efficient. It should be noted that even if the agent is of high complexity as the Honda robot, there is nothing in principle that prevents the exploitation of its passive dynamics. In human walking, for example, the forward swing of the leg is largely passive as well. Of course, the Honda robot can do many things like walking up and down the stairs, pushing a cart, opening a door, etc., whereas the ecological niche of the passive dynamic walker is confined to inclines of a particular angle.

There are two main conclusions that can be drawn from these examples. First, it is important to exploit the dynamics in order to achieve energy-efficient and natural kinds of movements. The term "natural" not only applies to biological systems, but artificial systems also have their intrinsic natural dynamics. Second, there is a kind of trade-off or balance: the better the exploitation of the dynamics, the simpler the control, the less neural processing will be required.

Muscles – control from materials

Let us pursue this idea of exploiting the dynamics a little further and show how it can be taken into account to design actual robots. Most robot arms available today work with rigid materials and electrical motors. Natural arms, by contrast, are built of muscles, tendons, ligaments, and bones, materials that are non-rigid to varying degrees. All these materials have their own intrinsic properties like mass, stiffness, elasticity, viscosity, temporal characteristics, damping, and contraction ratio to mention but a few. These properties are all exploited in interesting ways in natural systems. For example, there is a natural position for a human arm which is determined by its anatomy and by these properties. Grasping an object like a cup with the right hand is normally done with the palm facing left, but could also be done – with considerable additional effort – the other way around. Assume now that the palm of your right hand is facing right and you let go. Your arm will immediately turn back into its natural position. This is not achieved by neural control but by the properties of the muscle-tendon system: On the one hand the system acts like a spring – the more you stretch it, the more force you have to apply and if you let go the spring moves back into its resting position. On the other hand there is intrinsic damping. Normally reaching equilibrium position and damping is conceived of in terms of electronic (or neural) control, whereas in this case, this is achieved (mostly) through the material properties. Or put differently, the morphology (the anatomy), and the materials provide physical constraints that make the control problem much easier – at least for the standard kinds of movements.

These ideas can be transferred to robots. Many researchers have started building artificial muscles (for reviews of the various technologies see, e.g., Kornbluh et al., 1998 and Shahinpoor, 2000) and used them on robots, as illustrated in figure 2. ISAC, a “feeding robot”, and the artificial hand by Lee and Shimoyama use pneumatic actuators, Cog the series elastic actuators, and the Face Robot shape memory alloys. Facial expressions also provide an interesting illustration for the point to be made here. If the facial tissue has the right sorts of material properties in terms of elasticity, deformability, stiffness, etc., the neural control for the facial expressions becomes much simpler. For example, for smiling, although it involves the entire face, the actuation is very simple: the “complexity” is added by the tissue properties. Another highly desirable property that one gets for free if using the right kinds of artificial muscles is passive compliance: if an arm, for example, encounters resistance it will yield elastically rather than pushing harder. In the case of the pneumatic actuators this is due to the elastic properties of the rubber tubes.

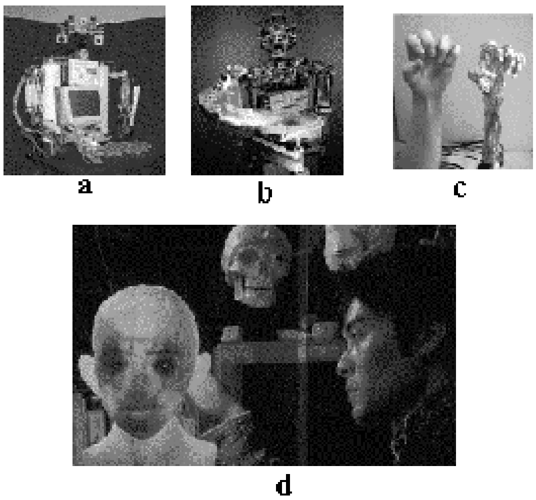

Figure 2: Robots with artificial muscles. The service robot ISAC by Peters (Vanderbilt University) driven by McKibben pneumatic actuators. (b) The humanoid robot Cog by Rodney Brooks (MIT AI Laboratory), driven by series-elastic actuators. (c) The artificial hand by Lee and Shimoyama (University of Tokyo), driven by pneumatic actuators. (d) The “Face Robot” by Kobayashi, Hara, and Iida (Science University of Tokyo), driven by shape-memory alloys.

Stumpy – a synthesis

Recently, there has been an increased interest in applying and further investigating these ideas to the construction of robots. An illustrative example is the walking and hopping robot Stumpy (Paul et al. 2002) (figure 3). Stumpy’s lower body is made of an inverted “T” mounted on wide springy feet. The upper body is an upright “T” connected to the lower body by a rotary joint, the “waist” joint, providing one degree of freedom in the frontal plane. The horizontal beam on the top is weighted on the ends to increase its moment of inertia. It is connected to the vertical beam by a second rotary joint, providing one rotational degree of freedom, in the plane normal to the vertical beam, the “shoulder” joint. Stumpy’s vertical axis is made of aluminum, while both its horizontal axes and feet are made of oak wood.

Although Stumpy has no real legs or feet, it can locomote in many interesting ways: it can move forward in a straight or curved line, it has different gait patterns, it can move sideways, and it can turn on the spot. Interestingly, this can all be achieved by actuating only two joints with one degree of freedom. In other words, control is extremely simple – the robot is virtually “brainless”. The reason this works is because the dynamics, given by its morphology and its materials (elastic, spring-like materials, surface properties of the feet), is exploited in clever ways. There is a delicate interplay of momentum exerted on the feet by moving the two joints in particular ways (for more detail, see Paul et al., 2002a, b).

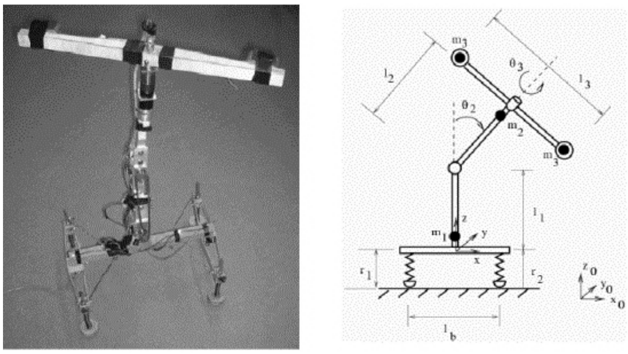

a b

Figure 3: The walking and hopping robot Stumpy. (a) Photograph of the robot. (b) Schematic drawing (details, see text).

Let us briefly summarize the ideas concerning the interplay between morphology, materials, and control. First, given a particular task environment, the (physical) dynamics of the agent can be exploited which leads not only to a natural behavior of the agent, but also to higher energy-efficiency. Second, by exploiting the dynamics of the agent, often control can be significantly simplified while maintaining a certain level behavioral diversity. And third, materials have intrinsic control properties.

We have now talked about ants on the beach, simple robots, and artificial muscles. How does this all fit together and how does it relate to intelligence? How can we deepen our understanding of these relationships, how can they be explored systematically, and how could these intuitions be made more quantitative? These are not questions that can be answered now; they constitute in fact major challenges. All we can do now is outline various approaches. This will be done in the remainder of the paper.

3. Exploring "ecological balance"

So far we have mostly argued that having the right morphology and materials may lead to simpler, and cheaper control. Initially, we briefly mentioned the concept of sensory-motor coupling which implies that an agent, through its interaction with the real world, can actively generate and structure its sensory data (e.g. Pfeifer and Scheier, 1997; Scheier et al., 1998). This structuring of the sensory data is now facilitated by having a proper morphology and proper materials: Because the palm of the hand and the finger tips are normally facing inwards there is a high probability that if the right arm, for example, moves from right to left, the palm and the finger tips will touch and grasp an object. This not only leads to rich haptic sensory stimulation (because of the high density of haptic sensors on the finger tips), but will also bring the object into the visual field. This way correlations in the sensory data within one channel and between channels can be induced (e.g. Lungarella and Pfeifer, 2001; Pfeifer and Scheier, 1997; Pfeifer and Scheier, 1999). Inducing correlations of this sort is essential for processes of categorization and concept development, as we have argued in detail elsewhere (e.g. Pfeifer, in press). These ideas have emerged from related studies that demonstrate the embodied nature of categorization and development (e.g. Edelman, 1987; Metta et al., 1998; Thelen and Smith, 1994). As categorization is one of the most basic cognitive abilities on top of which higher-level processes operate, this demonstrates the interdependence of embodiment and the development of cognition (For another line of argument making a similar point, see Lakoff and Nuñez, 2000). In this paper we focus on demonstrating how "ecological balance", the interplay of morphology, materials, and control, can be explored systematically using artificial evolution and morphogenesis.

Understanding and exploring "ecological balance" - artificial evolution and morphogenesis

Using artificial evolution for design has a tradition in the field of evolutionary robotics. The standard approach there is to take a particular robot and use a genetic algorithm to evolve a control architecture for a particular task. However, if we want to explore ecological balance we must include morphology and materials into our evolutionary algorithms.

The problem with including morphology and materials is that the search space which is already very large for control architectures only, literally explodes. Moreover, if sophisticated shapes and sensors are to be evolved, the length of the genome which is required for encoding these shapes will grow very large and there is no hope that anything will ever converge.

This issue can be approached in various ways, we just mention two. The first which we will not further discuss is to parameterize the shapes, thus bringing in biases from the designer on the types of shapes that are possible. An example that has stirred a lot of commotion in the media recently is provided by Hod Lipson and Jordan Pollack's robots that were automatically produced (Lipson and Pollack, 2000). They decided that the morphology would consist of rods to which different types of joints could be attached. Rods can, for example, be parameterized as length, diameter, and material constants etc., thus limiting the space of possible shapes, or in other words, the types of morphologies, dramatically, but then the search space, even though it is still large, becomes manageable. While this example is impressive, it still implies a strong designer bias. If we want to explore different types of morphologies, we want to introduce as little designer bias as possible. This can be done using ideas from biology, i.e. genetic regulatory networks.

The mechanics of artificial genetic regulatory networks

We provide a non-technical introduction, for details, see, e.g. Bongard and Pfeifer (2001; in press). It should be stressed, that although this computational system is biologically inspired, it does not constitute a biological model. Rather, it is system in its own right. Also, when we use biological terminology, e.g. when we say that "concentrations of transcription factors regulate gene expression", this is meant metaphorically.

The basic idea is the following. A genetic algorithm is extended to include ontogenetic development by growing agents from genetic regulatory networks. In the example presented here, agents are tested for how far they can push a large block (which is why they are called "block pushers"). Figure 4a shows the physically realistic virtual environment. The fitness determination is a two-stage process: the agent is first grown and then evaluated in its virtual environment. Figure 4b illustrates how an agent grows from a single cell into a multicellular organism.

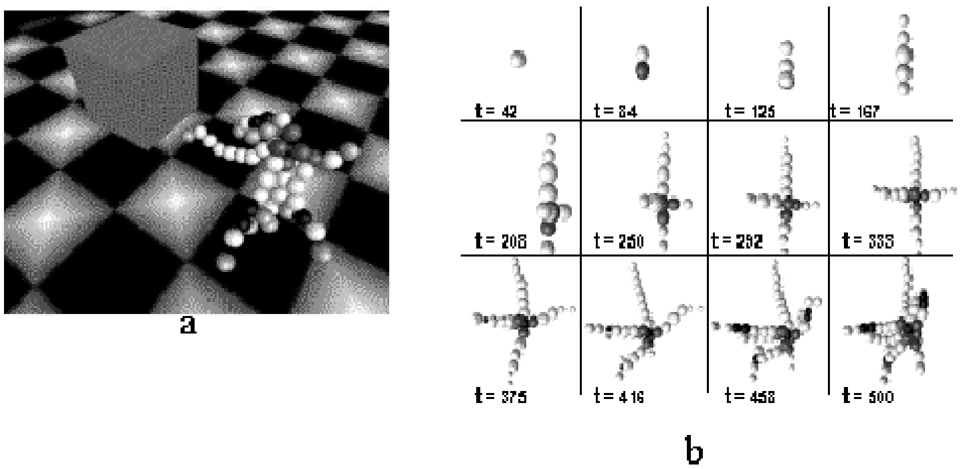

Figure 4: Examples of Bongard's "block pushers". An evolved agent in its physically realistic virtual environment. (b) growth phase starting from a single cell, showing various intermediate stages (last agent after 500 time steps).

The algorithm starts with a string of randomly selected floating point numbers between 0 and 1. A scanning mechanism determines the location of the genes. Each gene consists of 6 floating point numbers which are the parameters that evolution can play with. They are explained in figure 5. There are transcription factors that only regulate the activity of other genes, there are transcription factors for morphology, and for neuronal growth. Whenever a gene is "expressed", it will diffuse a transcription factor into the cell from a certain diffusion site. The activity of this genetic regulatory network leads to particular concentrations of the transcription factors to which the cell is sensitive: whenever a concentration threshold is exceeded, an action is taken. For example, the cell may increase or decrease in size, if it gets too large, it will split, the joint angles can be varied, neurons can be inserted, connections added or deleted, structures can be duplicated, etc. The growth process begins with a single unit into which "transcription factors" are injected (which determines the primary body axis). Then it is left to the dynamics of the genetic regulatory network. The resulting phenotype is subsequently tested in the virtual environment. Over time, agents evolve that are good at pushing the block.

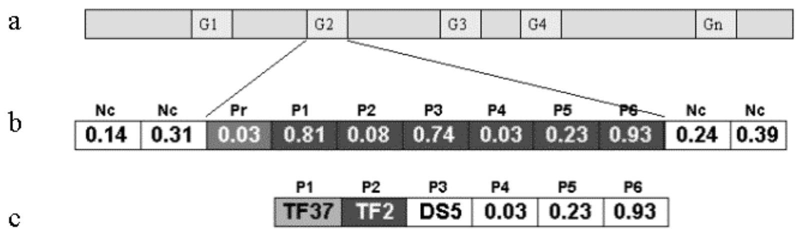

Figure 5: The mechanisms underlying the genetic regulatory networks. (a) Genes on the genome. Which regions are considered to be genes is determined by an initial scanning mechanism (values below 0.1 are taken as starting positions). (b) and (c) An example of a particular gene. Nc means "non-coding" region, Pr is a promoter site (start of gene), P1 through P6 are the parameters of the gene. P1: the transcription factor (TF) that regulates the expression of this gene [0,19]. P2: the TF the gene emits if expressed [0,42]. P3: the diffusion site, i.e. the location in the cell from which the TF is diffused. P4: the quantity of TF emitted by this gene, if expressed. P5, P6: lower and upper bounds of the concentrations within which the gene is expressed.

Emergence – the achievements of artificial evolution and morphogenesis

Although simple in their basic form, these mechanism lead to an interesting dynamics and produce fascinating results. Here are some observations: (1) Organisms early on in evolution are typically smaller than those of later generations: evolution discovers that in order to push a block of large size, it is necessary to have a large body. In other words, evolution had to manipulate morphology in order to achieve the task. (2) Evolution comes up with means of locomotion. In small creatures, these are very local reflex-like mechanisms distributed through the entire organism. Larger creatures tend to have additional tentacles that can be used to push against the block, which requires a different kind of control. (3) There is no direct relation between genotype length and phenotypic fitness – the two are largely dissociated. (4) There is functional specialization, i.e. cells differentiate into units containing both sensors and actuators (the white colored cells in figure 4), cells that only contain sensors but no actuators (gray coloring), and cells not containing anything, only providing structural support (black coloring). (5) There is repeated structure, i.e. some combination of cells occur in slightly modified form in various places on the agent. An example from biology are fingers that are similar but differ individually. (6) Some genes specialize to become "master regulatory genes", i.e. they regulate the activity of other genes. Thus, to an outside observer, it looks as if a hierarchical structure were evolving in the regulatory network. Note that this hierarchy is emergent and results from a "flat" dynamical system. Thus, it can change at a later point in time, unlike "structural" hierarchies. Again, metaphorically speaking, artificial evolution has discovered how to manage complexity, i.e. by evolving a hierarchical organization. It is important to mention that this has all been "discovered" by simulated evolution and has not been programmed into the system. Or stated differently, it is emergent from the mechanisms of simulated evolution and genetic regulatory networks.

The work of Eggenberger (1997, 1999) is among the first to employ genetic regulatory networks to model growth processes in computational systems. He succeeded in evolving three-dimensional shapes. As in the case of Bongard's system, the resulting shape (or organism) is emergent from a complex dynamical system.

.

4. Discussion and conclusions

We have argued that the cognitivistic paradigm has not succeeded in explaining natural forms of intelligence. The paradigm of embodied intelligence has been proposed instead which states that intelligence is emergent from an agent's interaction with a real physical world. We have demonstrated some of the implications of embodiment in various case studies. For example, given the right morphology and materials, it is often astonishing how simple control can produce sophisticated behaviors. This holds for the passive dynamic walker of Steve Collins, and for Stumpy. We have introduced the concept of "ecological balance", which states that given a task environment, there is an intimate interaction between morphology, materials, and control, morphology and materials taking over some of the control functions. As an example of this we have discussed the properties of the muscle tendon system. In order to understand "ecological balance" we can study natural systems where evolution has found appropriate distributions of tasks among these parts or aspects of an agent. However, it is important not only to explore what nature has discovered, but also "life as it could be", to use Chris Langton's phrase.

By building systems that differ in some respects from the natural ones we can learn about the natural systems, but we can also acquire an understanding of behavior in general, beyond the natural system. This approach has been called the synthetic methodology (e.g. Braitenberg, 1984; Pfeifer and Scheier, 1999; Reeke et al., 1989), "understanding by building." In applying this methodology we are using particular technologies, robots or simulations, which necessarily differ from their natural counterparts.

As an instrument that can be used to explore the relation or "task distribution" between morphology, materials and control in general, not constrained to biological systems, we have proposed artificial evolution and morphogenesis and have shown how this system which enables the simultaneous evolution of morphology and neural substrate, comes up with highly interesting solutions that may or may not be found in natural systems. Hopefully, this will eventually lead to an improved understanding of this "task distribution" and of embodiment in general.

There are a number of limitations of this approach that we will put on the research agenda for the coming years. One is the incorporation of interaction with the environment during ontogenetic development. Moreover, the "rewrite rules" for neuronal growth will be replaced by more biological mechanisms. Third, instead of defining a fitness function, we will turn to "open-ended evolution" where the survival of the individual is the sole criterion. This requires the definition of pertinent resources that need to be maintained. Fourth, we need to incorporate the variation of material properties into the evolutionary algorithm, so that this aspect can be studied as well. And last but not least, we need to be able to increase the complexity of our task environments which requires much higher computational power.

At the moment we are confined to simulation; the experiments with artificial systems that can grow physically are only in their very initial stages. One way to get around this problem, at least to some extent, is on the one hand to have a good simulator that models the physics of an evolved individual and its interactions with the real world (e.g. gravity, impact, friction), on the other to have rapid robot building kits that enable the researchers to quickly build a robot to test some individuals in the real world. But even if done in simulation, evolving an organism from scratch is a big challenge as well.

One of the problems with the examples and ideas presented in this paper is that they are mostly qualitative. Clearly, more quantitative statements will be required to make the story more compelling. But we do hope that researchers will take up the challenges posed by embodiment.

Let us conclude by coming back to the title of the paper. Of course, we have not explained cognition. But rather than simply programming cognitive processes into algorithms, we have tried to study complete agents that might eventually show cognitive behaviors as emergent phenomena. While one would hardly call the activities of the block pusher or Stumpy "cognitive", we believe that by increasing the complexity of the task environments, eventually more complex agents will evolve, agents that will display more sophisticated behaviors that we then might want to call cognitive. One hope is, for example, that as the environments and agents get more complex, involving not only one or a few tasks, but perhaps hundreds or thousands, we will begin to see a certain centralization of the neural substrate which in the very simple creatures is largely distributed through the entire agent.

Acknowledgments

I would like to thank the members of the Artificial Intelligence Laboratory for many discussions, in particular Josh Bongard for his patience in explaining evolution to me and Gabriel Gómez for discussing the manuscript. Credit also goes to the Swiss National Science Foundation for supporting the research presented in this paper, grant # 20-61372.00.

References

Braitenberg, V. (1984). Vehicles: Experiments in synthetic psychology. Cambridge, MA.: MIT Press.

Bongard, J., and Pfeifer, R. (2001). Repeated structure and dissociation of genetypic and phenotypic complexity in artificial ontogeny. Genetic and Evolutionary Computatio Confer-ence, GECCO-2001, 829-836.

Bongard, J., and Pfeiffer, R. (in press). Envolving complete agents using artificial ontogeny, to appear in: F. Hara, and R. Pfeiffer (eds.). Morphofunktional machines the new species. Designing embodied intelligence. Berlin. Springer-Verlag.

Brooks, R. A. (1991a). Intelligence without representation. Artificial Intelligence, 47, 139-160.

Brooks, R. A. (1991b). Intelligence without reason. Proceedings of the International Joint Conference on Artificial Intelligence-91, 569-595.

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.

Edelman, G.E. (1987). Neural Darwinism. *The theory of neuronal group selection*. 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.

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.

Kornbluh, R. D., Pelrine, R., Eckerle, J., and Joseph, J. (1998). Electrostrictive polymer artificial muscle actuators. *Proceedings of the IEEE International Conference on Robotics and Automation* 1998. New York, N.Y.: IEEE, 2147-2154.
Lakoff, G., and Núñez, R.E. (2000). *Where mathematics comes from. How the embodied mind brings mathematics into being.* New York, N.Y.: Basic Books.
Lipson, H., and Pollack J. B. (2000), Automatic design and manufacture of artificial life forms. Nature, 406, 974-978.
Lichtensteiger, L., and Eggenberger, P. (1999). Evolving the morphology of a compound eye on a robot. *Proceedings of the third European Workshop on Advanced Mobile Robots (Eurobot'99)*. IEEE, Piscataway, NJ, USA; 1999; 127-34 .
Lungarella, M., and Pfeifer, R. (2001). Robots as cognitive tools: information theoretic analysis of sensory-motor data. *Proc. of the IEEE-RAS £International Conference on Humanoid Robots*, 245-252.
Maris, M., and te Boekhorst, R. (1996). Exploiting physical constraints: heap formation through behavioral error in a group of robots. *Proceedings of the IROS'96, IEEE/RSJ International Conference on Intelligent Robots and Systems, 1655—1660.*
McGeer, T. (1990a). Passive dynamic walking. Int. *Journal of Robotics Research*, 9, 62-82.
McGeer, T. (1990b). Passive walking with knees. *Proc. of the IEEE Conference on Robotics and Automation*, 2, 1640-1645.
Metta, G., Sandini, G., and Konczak, J. (1998). A developmental approach to sensori-motor coordination in artificial systems. *Proceedings of IEEE Conference on System, Man and Cybernetics*, San Diego (USA), 11-14.
Paul, C., Dravid, R. and F. Iida (2002a) Control of lateral bounding for a pendulum driven hopping robot. to appear in *Proceedings of the International Conference of Climbing and Walking Robots* , Paris, France (to appear)..
Paul, C., Dravid, R. and F. Iida (2002b) Design and Control of a Pendulum Driven Hopping Robot. *Proc of the IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS-2002*, Lausanne, Switzerland (to appear).
Pfeifer, R. (1996). Building "Fungus Eaters": Design principles of autonomous agents. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S.W. Wilson (eds.): *From Animals to Animats 4. Proceedings of the fourth International. Conference on Simulation of Adaptive Behavior.* Cambridge, Mass.: A Bradford Book, MIT Press, 3-12.
Pfeifer, R. (1999). Dynamics, morphology, and materials in the emergence of cognition. In Burgard, W., Christaller, T., Cremers, A. B. (eds.): *KI-99 Advances in Artificial Intelligence. Proceedings of the 23rd Annual German Conference on Artificial Intelligence*, Bonn, Germany, 1999, Lecture Notes in Computer Science, Springer, 1701, 27-44.
Pfeifer, R. (2000). On the role of morphology and materials in adaptive behavior. 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*, 23-32.
Pfeifer, R. (2001). Embodied Artificial Intelligence: 10 years back, 10 years forward. In: R. Wilhelm (ed.). *Informatics – 10 years back, 10 years ahead. Lecture Notes in Computer Science.* Berlin: Springer, 294-310.

Pfeifer, R. (in press). Robots as cognitive tools. *Journal of Cognitive Technology* (to appear).

Pfeifer, R., and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. Robotics and Autonomous Systems, 20, 157-178.

Pfeifer, R., and Scheier, C. (1998). Representation in natural and artificial agents: an embodied cognitive science perspective. *Zeitschrift für Naturforschung,* 53c, 480-503.

Pfeifer, R., and Scheier, C. (1999). *Understanding intelligence.* Cambridge, Mass.: MIT Press.

Reeke, N.G., Finkel, L.H., Sporns, O., and Edelman, G.A. (1998). Synthetic neural modeling: a multilevel approach to the analysis of brain complexity. In G.M. Edelman, W.E. Gall, W.M. Cowan (eds.). *Signal and Sense: local and global order in perceptual maps.* New York, N.Y.: Wiley, 607-706.

Scheier, C., Pfeifer, R., and Kuniyoshi, Y. (1998). Embedded neural networks: exploiting constraints. *Neural Networks*, 11, 1551-1569.

Shahinpoor, M., Bar-Cohen, Y., Simpson, J.O., and Smith, J. (2000). Ionic Polymer-Metal Composites (IPMC) as biomimetic sensors, actuators & artificial muscles- A review. http://www.unm.edu/~amri/paper.html

Simon, H. A. (1969). *The sciences of the artificial* (2nd ed.). Cambridge, MA: MIT Press.

Thelen, E. and Smith, L. (1994). A dynamic systems approach to the development of cognition and action. Cambridge, Mass.: MIT Press, Bradford Books.