

New Robotics: Design Principles for Intelligent Systems

Rolf Pfeifer¹, Fumiya Iida¹, Josh Bongard²

¹ Artificial Intelligence Laboratory, Department of Information Technology, University of Zurich
Andreasstrasse 15, CH-8050 Zurich, Switzerland
pfeifer@ifi.unizh.ch, phone: +41 1 63 5 4320/31, fax: +41 1 635 68 09

² Sibley School of Mechanical and Aerospace Engineering, Cornell University
Ithaca, New York, USA

Paper prepared for the Special Issue of Artificial Life on
New Robotics, Evolution and Embodied Cognition
6 April 2003

Abstract

New Robotics designates an approach to robotics that, in contrast to traditional robotics, employs ideas and principles from biology. While in the traditional approach there are generally accepted methods (e.g. from control theory), designing agents in the New Robotics approach is still largely considered an art. In recent years, we have been developing a set of heuristics or design principles, that on the one hand capture theoretical insights about intelligent – adaptive – behavior, and on the other provide guidance in actually designing and building systems. In this paper we provide an overview of all the principles but focus on the principles of “ecological balance” which concerns the relation between environment, morphology, materials, and control, and “sensory-motor coordination” which concerns self-generated sensory stimulation as the agent interacts with the environment and which is a key to the development of high-level intelligence. As we will argue, artificial evolution together with morphogenesis is not only “nice to have” but is in fact a necessary tool for designing embodied agents.

1. Introduction

In the past the focus in the field of robotics has been on precision, speed, and controllability; more recently there has been an increasing interest in adaptivity, learning, and autonomy. The reasons for this are manifold, but an important one is the growing concern in the research community for using robots to study intelligent systems. This has led to a new set of goals for this new kind of robotics. Of course, one of the important goals remains the development of useful artifacts, but it is no longer the

major one. More important is an increased understanding of the principles underlying intelligent behavior. A typical starting point is the modeling of some aspects of biological systems, such as how ants find their way back to the nest after finding food, how dogs catch a Frisbee while running, how rats navigate in a maze, or how people recognize a face in a crowd. Experience has shown that building a robot (or a computer model) considerably enhances our understanding of the natural system. A third goal is the abstraction of general principles of intelligent behavior, which is the main topic of this paper.

Recently, a host of terms have been invented to characterize this new approach to the design of intelligent systems, one that, in contrast to the traditional ones of artificial intelligence and robotics, draw inspiration from biology. The field of adaptive systems, as loosely characterized by conferences such as SAB (Simulation of Adaptive Behavior) or AMAM (Adaptive Motion in Animals and Machines), or the journals of Artificial Life and Adaptive Behaviour, is very heterogeneous and there is a definite lack of consensus on the theoretical foundations. As a consequence, agent design is – typically – performed in an ad hoc and intuitive way. Although there have been some attempts at elaborating principles (e.g. Brooks, 1991; Bryson, 2002; Maes, 1993), general agreement is still lacking. In addition, much of the work on designing adaptive systems is focused on software, i.e. the programming of the robots. However, what we are really interested in is not just the programming aspects, but rather designing entire systems. The research conducted in our laboratory, but also by many others, has demonstrated that often, better, cheaper, more robust and adaptive agents can be developed if the entire agent is the design target rather than only its controller. This implies taking embodiment into account and going beyond the programming level proper. Below, we will characterize in detail what we mean by embodiment; for now we take it to

mean the physical setup of the agent which includes the body plan, sensors, actuators, etc. Because we are not only looking at programs, we prefer to use the term “engineering agents” rather than “programming agents”.

If this idea of engineering agents is the goal, the question arises as to what form the theory should have, i.e. how the experience gained so far can be captured in a concise scientific way. Philosophers have a preference for verbal and logic-based descriptions. However, verbal theories often gloss over important details and thus lack the required precision. Logic-based theories are typically very limited in their expressive power. Moreover, they suggest assumptions about the real world that simply do not hold. For example, there is an underlying assumption of discrete states that can be achieved by applying certain types of operators, an assumption that has misled many researchers in the past. Researchers from traditional artificial intelligence and cognitive science believe that the best way to make progress is to employ symbol processing concepts (e.g. Newell and Simon, 1976). This idea has led to interesting applications from a computer science perspective but has contributed little to our understanding of the nature of natural forms of intelligence.

Physicists and control engineers have a preference for differential equations, but they are better suited for analysis than for design. The mathematical theory of dynamical systems that capitalizes on this formalism is considered by many to ultimately be the best candidate. While this approach yields very interesting results for relatively simple robotic systems (e.g. Smithers, 1995; Steinhage and Schoener, 1997; Berthouze, 2000; Yamamoto and Kuniyoshi, 2001), it seems to be very difficult to apply it to complex systems beyond the metaphorical level. Moreover, applying this type of theory for design purposes does not seem to be straightforward. For the time being, it seems that progress over the last few years in the field has been slow, and we may be well-advised to search for an “intermediate” solution, between no theory at all (or purely verbal ones), and a rigorous mathematical one.

A set of design principles as a theoretical framework for understanding intelligence seems desirable for a number of reasons. First, at least at the moment, there do not seem to be any real alternatives. The information processing paradigm, another potential candidate, has proven ill-suited to come to grips with natural, adaptive forms of intelligence. Second, because of the current unfinished status of the field, a set of principles is flexible and can be dynamically changed and extended. Third, design principles present heuristics for actually building systems. In this sense, they instantiate the synthetic methodology (see below). And fourth, evolution can also be viewed as a designer, albeit a “blind one”, but an extremely powerful one nevertheless. We hope to convince the reader that a framework founded on a set of design principles is a good way to make progress, and that researchers will take it up, modify the principles, add new ones, and try to make the

entire set more comprehensive and coherent. The response so far has been highly encouraging and researchers as well as educated lay people seem to be able to relate to these principles very easily.

Although most of the literature is still about programming, some of the research explicitly deals with complete agent design and includes aspects of morphology (e.g. Sims, 1994a, b; Pfeifer, 1996; Lipson and Pollack, 1999; Pfeifer and Scheier, 1999; Hara and Pfeifer, 2000a; Bongard and Pfeifer, 2001; Bongard, 2002, 2003; Pfeifer, 2003). Our own approach over the last six years or so has been to try and systematize the insights gained in the fields of adaptive behavior and intelligence in general by incorporating ideas from biology, psychology, neuroscience, engineering, and artificial intelligence into a set of design principles, as argued above; they form the main topic of this paper.

A first version of the design principles was published at the 1996 conference on Simulation of Adaptive Behavior (Pfeifer, 1996). A more elaborate version has been published in the book “Understanding Intelligence” (Pfeifer and Scheier, 1999). More recently, some principles have been extended to incorporate ideas on the relation between morphology, materials, and control (Ishiguro et al., 2003; Hara and Pfeifer, 2000a; Pfeifer, 2003). An updated detailed summary will be published in Pfeifer and Glatzeder (in preparation).

We start by giving a very short overview of the principles. We then discuss the information theoretic implications of embodiment, i.e. the relation between physics and information processing, using a number of case studies. In particular we illustrate the concept of “ecological balance”. We then show how the principle of “sensory-motor coordination” can be used to explain important aspects of the development of higher levels of intelligence. Subsequently it is demonstrated how artificial evolution together with morphogenesis can be employed to design ecologically balanced systems. We will speculate that, together with a sophisticated physics-based simulation, it might eventually lead to the design of systems with higher levels of intelligence. We will make clear that these considerations are only applicable to embodied systems.

This is not a technical paper but a conceptual one. The goal is to provide a framework within which technical research can be conducted that takes into account the most recent insights in the field. In our argumentation we will resort mostly to research conducted in our own laboratory but also to research performed in the community at large. In this sense, the paper has somewhat of a tutorial and review flavor and should be considered as such. Moreover, it is intended for an interdisciplinary audience and thus largely avoids technical jargon.

2. Design principles: Overview

There are, in essence, two different types of design principles: some are concerned with the general “philosophy” of the approach. We call them “design procedure principles”, as they do not directly pertain to the design of the agents themselves but rather to the way of proceeding, to the methodology. Another set of principles is concerned more with the actual design of the agent. We use the qualifier “more” to express the fact that we are often not designing the agent directly but rather the initial conditions and the learning and developmental processes or the evolutionary mechanisms as we will elaborate later. The current overview will, for reasons of space, be very brief; a more extended version is forthcoming (Pfeifer and Glatzeder, in preparation).

Table 1: Overview of the design principles.

A-Princ 5	Cheap design	Exploitation of niche and interaction; parsimony
A-Princ 6	Redundancy	Partial overlap of functionality based on different physical processes
A-Princ 7	Ecological balance	Balance in complexity of sensory, motor, and neural systems: task distribution between morphology, materials, and control
A-Princ 8	Value	Driving forces; developmental mechanisms; self-organization

Label	Name	Description
		<i>Design procedure principles</i>
P-Princ. 1	Synthetic methodology	Understanding by building
P-Princ 2	Emergence	Systems should be designed for emergence (for increased adaptivity)
P-Princ 3	Diversity-compliance	Tradeoff between exploiting the givens and generating diversity solved in interesting ways
P-Princ 4	Time perspectives	Three perspectives required: “Here and now”, ontogenetic, phylogenetic
P-Princ 5	Frame-of-reference	Three aspects must be distinguished: perspective, behavior vs. mechanisms, complexity
		<i>Agent design principles</i>
A-Princ 1	Three constituents	ecological niche (environment), tasks, and agent must always be taken into account
A-Princ 2	Complete agent	Embodied, autonomous, self-sufficient, situated agents are of interest
A-Princ 3	Parallel, loosely coupled processes	Parallel, asynchronous, partly autonomous processes, largely coupled through interaction with environment
A-Princ 4	Sensory-motor coordination	Behavior sensory-motor coordinated with respect to target; self-generated sensory stimulation

P-Princ 1: The synthetic methodology principle. The synthetic methodology, “understanding by building”, implies on the one hand constructing a model – computer simulation or robot – of some phenomenon of interest (e.g. how an insect walks, how a monkey grasps a banana, how babies learn to make distinctions in the real world, or how we recognize a face in a crowd). On the other we want to abstract general principles from the constructed model (some examples are given below). The term “synthetic methodology” was adopted from Braitenberg’s seminal book “Vehicles: Experiments in synthetic psychology” (Braitenberg, 1984).

P-Princ 2: The principle of emergence. If we are interested in designing adaptive systems we should aim for emergence. The term emergence is controversial, but we use it in a very pragmatic way, in the sense of not being preprogrammed. When designing for emergence, the final structure of the agent is the result of the history of its interaction with the – simulated or real world – environment. Strictly speaking, behavior is always emergent, as it cannot be reduced to internal mechanism only; it is always the result of a system-environment interaction. In this sense, emergence is not all-or-nothing phenomenon, but a matter of degree: the further removed from the actual behavior the designer commitments are made, the more we call the resulting behavior emergent. Systems designed for emergence tend to be more adaptive and robust. For example, a system specifying initial conditions and developmental mechanisms will automatically will automatically exploit the environment to shape the agent’s final structure. Another example from locomotion (see below) is the exploitation of the intrinsic material properties of an agent: if the spring-like properties of the muscles are exploited, the details of the trajectories of the joints are emergent and need not be controlled.

P-Princ 3: The diversity-compliance principle. Intelligent agents are characterized by the fact that they are on the one hand exploiting the specifics of their ecological niche and on the other by behavioral diversity. In a conversation I have to comply with the rules of grammar of the particular language, and then I have to react to what the

other individual says, and depending on that, I have to say something different. Always uttering one and the same sentence irrespective of what the other is saying would not demonstrate great behavioral diversity. This principle or trade-off comes in many variations in cognitive science, i.e. the plasticity-stability tradeoff in learning theory (Grossberg 1995), assimilation-accommodation in perception (Piaget 1970), and exploration-exploitation in evolutionary theory (Eiben and Schippers 1998).

P-Princ 4: The time perspectives principle. A comprehensive explanation of behavior of any system must incorporate at least three perspectives: (a) state-oriented, or the “here and now”, (b) learning and development, the ontogenetic view, and (c) evolutionary, the phylogenetic perspective. The fact that these perspectives are adopted by no means implies that they are separate. On the contrary, they are interdependent, but it is useful to tease them apart for the purpose of scientific investigation. Note the connection to the principle of emergence: if a time perspective can be explained as being emergent from another, we have a deeper kind of explanation. For example, if the “here and now” can be explained as being emergent from the ontogenetic one, this constitutes scientific progress.

P-Princ 5: The frame-of-reference principle. There are three aspects to distinguish between whenever designing an agent: (a) the perspective, i.e. are we talking about the world from the agent’s, observer’s, or designer’s perspective; (b) behavior is not reducible to internal mechanism; trying to do that would constitute a category error; and (c) apparently complex behavior of an agent does not imply complexity of the underlying mechanism. Although it seems obvious that the world “looks” very different to a robot than to a human because the robot has completely different sensory systems than a human, this fact is surprisingly often ignored. Second, behavior cannot be completely programmed, but is always the result of a system-environment interaction. Again, it is surprising how often this obvious fact is ignored even by roboticists. And third, the complexity of the environment plays an essential role in the behavior and thus in the ways in which this complexity is perceived by an observer. Thus, behavioral complexity cannot be attributed to the agent alone, but to the agent-environment interaction (see the discussion of “Simon’s Ant” in (Simon 1969); also (Seth 2002)).

A-Princ 1: The three-constituents principle. (a) The definition of the ecological niche (the environment), (b) the desired behaviors and tasks, and (c) the agent itself. The main point of this principle is that it would be a fundamental mistake to design the agent in isolation. This is particularly important because much can be gained by exploiting its physical and social environment.

A-Princ 2: The complete agent principle. The agents of interest are autonomous (a relative notion relating to the degree of independence of other agents), self-sufficient (i.e. they can sustain themselves over extended periods of

time), embodied (i.e. realized as physical systems), and situated (i.e. they can acquire information about the environment through their own sensory systems as a result of their interaction with the real world). This perspective, although extremely powerful and obvious, is not considered explicitly very often.

A-Princ 3: The principle of parallel, loosely coupled processes. Intelligence is emergent from an agent-environment interaction based on a large number of parallel, loosely coupled processes that run asynchronously and are connected to the agent’s sensory-motor apparatus. The term “loosely coupled” is used in contrast to hierarchically coupled processes in which there is a program calling a subroutine and the calling program has to wait for the subroutine to complete its task before it can continue. In that sense, this hierarchical control corresponds to very strong coupling. However, on a complete agent, there can be a very strong coupling of processes by the fact that the system is embodied: two joints coupled by a physical link (bones) are very strongly coupled as well. “Loosely coupled” also refers to the coupling through the interaction with the environment.¹

A-Princ 4: The principle of sensory-motor coordination. All intelligent behavior (e.g. perception, categorization, memory) is to be conceived as sensory-motor coordination. This sensory-motor coordination, in addition to enabling the agent to interact efficiently with the environment, serves the purpose of structuring its sensory input. One of the powerful implications is that the problem of categorization in the real world is greatly simplified through the interaction with the real world because the latter supports the generation of “good” patterns of sensory stimulation, “good” meaning correlated, and stationary (at least for a short period of time). This principle is essential in the development of higher-level cognition (see below).

A-Princ 5: The principle of cheap design. Designs must be parsimonious, and exploit the physics and the constraints of the ecological niche. This principle is related to the diversity compliance principle in that it implies, for example, compliance with the laws of physics. An example are robots with wheels that exploit the fact that the ground is mostly flat (which is given, for example, in office environments).

A-Princ 6: The redundancy principle. Agents should be designed such that there is an overlap of functionality in the different subsystems. Examples are sensory systems where, for example, the visual and the haptic systems both deliver spatial information, but they are based on different physical processes (electromagnetic waves vs. mechanical touch). Merely duplicating components does not lead to

¹ In this case “indirectly coupled” might be the better term. For example, the legs in insect walking are partly coordinated through the interaction with the real world: if one leg is lifted, the force on all the other legs changes instantaneously, which can be exploited for coordination (Cruse et al. 2002).

useful redundancy; the partial overlap of functionality and the different physical processes are essential. Note that redundancy is required for diversity of behavior and to make a system adaptive. If there is a haptic system in addition to the visual one, the agent can also function in complete darkness, whereas one with 10 cameras ceases to function if the light goes out.

A-Princ 7: The principle of ecological balance. This principle consists of two parts: the first one concerns the relation between the sensory system, the motor system, and the neural control. The “complexity” of the agent has to match the complexity of the task environment, in particular: given a certain task environment, there has to be a match in the complexity of the sensory, motor, and neural system. The second is about the relation between morphology, materials, and control: given a particular task environment, there is a certain balance or task distribution between morphology, materials, and control (for references to both ideas, see, e.g. Hara and Pfeifer, 2000a; Pfeifer, 1996; Pfeifer, 1999, 2000; Pfeifer and Scheier, 1999). Often, if the morphology and the materials are right, control will be much cheaper. Because we are dealing with embodied systems, there will be two dynamics, the physical one or body dynamics and the control or neural dynamics. There is the important question of how the two can be coupled in optimal ways. The research initiated by Ishiguro and his colleagues (e.g. Ishiguro et al., 2003) promises deep and important pertinent insights. We will be giving examples of this principle later in the paper.

A-Princ 8: The value principle. This principle is, in essence, about motivation. It is about why the agent does anything in the first place. Moreover, a value system tells the agent whether an action was good or bad, and depending on the result, the probability of repetition of an action will be increased or decreased. Because of the unknowns in the real world, learning must be based on mechanisms of self-organization. There is a frame-of-reference issue in that values can be implicit or explicit. If an agent is equipped, say, with neural networks for Hebbian learning, we can, as outside observers, say that this constitutes value to the agent because in this way it can learn correlations, certainly a useful thing. If as a result of a particular action, a particular internal signal – neural or hormonal – is generated that modulates learning, we talk about an explicit value system (we adopted the term from Edelman, 1987).

The topic of value systems is central to agent design and must be somehow resolved. However, it seems that to date no generally accepted solutions have been developed. As this is not the central topic of this paper, we will not further elaborate on this issue. We anticipate that there will be a subset of principles devoted to precisely these issues.

Although it does capture some of the essential characteristics of adaptive systems, the set of principles described above is by no means complete. A series of

principles for designing evolutionary systems and collective systems is currently under development.

As mentioned earlier, all these principles only hold for embodied systems; they could not possibly be applied in the context of traditional symbol processing systems. In this paper, we focus on the principles of ecological balance and sensory-motor coordination which lie at the heart of embodiment.

3. Information theoretic implications of embodiment

There is a trivial meaning of embodiment, namely that “intelligence requires a body”. In this sense, anyone using robots for his or her research is doing embodied artificial intelligence. It is also obvious that if we are dealing with a physical agent, we have to take into account gravity, friction, torques, inertia, energy dissipation, elasticity of materials, etc. However, there is a non-trivial meaning of embodiment, namely that there is a tight interplay between the physical and the information theoretic aspects of an agent. The design principles all directly or indirectly refer to this issue, but some focus specifically on this interplay, such as the principle of sensory-motor coordination (embodied interaction with the environment induces sensory-motor patterns); the principle of cheap design (proper embodiment leads to simpler and more robust control); the redundancy principle (proper choice and positioning of sensors leads to robust behavior); and the principle of ecological balance (capitalization of the relation between morphology, materials, and neural control). For the purpose of illustration we will focus on ecological balance and sensory-motor coordination in this paper. We proceed by presenting a number of case studies illustrating the application of these principles to designing adaptive behavior.

A short note on terminology is required here. We mentioned information theoretic implications of embodiment. What we mean is the effect of morphology, materials, and environment on neural processing, or better, the interplay between all of these aspects. The important point is that the implications are not only of a purely physical nature.

Whenever we have an embodied system, through the embodiment itself, all aspects of an agent – sensors, actuators, limbs, the neural system – are always highly connected: changes to one component will potentially affect every other component. From this perspective we should never treat, for example, sensory and motor systems separately. However, for the purpose of investigation and writing, we must isolate the components, but at the same time we must not forget to view everything in the context of the complete agent. Having said that, we now proceed with a few case studies, first focusing on the sensory side, then the motor side, and finally on their integration.

Sensory systems

In previous papers we have investigated in detail the effect of changing sensor morphology on neural processing (e.g. Lichtensteiger and Eggenberger, 1999; Maris and te Boekhorst, 1996; Pfeifer, 2000a, b; Pfeifer and Scheier, 1999). Here we only summarize the main results; for details, the reader is referred to the literature.

The morphology of sensory systems has a number of important implications. In many cases, when the morphology of the sensory systems is suited for the particular task environment, more efficient solutions can be found. For example, it has been shown that for many tasks motion detection is all that is required. Motion detection can often be simplified even more if the light-sensitive cells are not spaced evenly, but if there is a non-homogeneous arrangement. For example, Franceschini and his co-workers found that in the house fly the spacing of the facets in the compound eye is more dense toward the front of the animal (Franceschini et al., 1992). Allowing for some idealization, this implies that under the condition of straight flight, the same motion detection circuitry – the elementary motion detectors, or EMDs – can be employed for motion detection for the entire eye. Based on these ideas, Franceschini et al. built fully analog robots exploiting this non-homogeneous morphological arrangement. It can be shown that this arrangement, in a sense, compensates for the phenomenon of motion parallax. It has been shown in experiments with artificial evolution on real robots that certain tasks, e.g. keeping a constant lateral distance to an obstacle, can be solved by proper morphological arrangement of the ommatidia, i.e. frontally more dense than laterally (Lichtensteiger and Eggenberger, 1999).

There is an additional important implication of morphology. It can be shown, that – again for specific interactions – learning speed can be increased significantly by having the proper morphology (Lichtensteiger and Pfeifer, 2002). The reason this works is that through the particular arrangement of the facets and the specific interactions, sensory data with particular statistical distributions are generated that support learning. The interesting aspect of this experiment, in contrast to many other experiments on machine learning, is that no specific distributions of the input data are assumed, but that they are generated through the interaction with the environment, through a sensory-motor coupling.

Note that the sensor morphology alone does not tell us very much; it is only if we take the specific interaction with the environment into account (which includes the actions of the motor system as well), that we are able to understand the role of morphology in behavior.

Franceschini and Changeux also found that there is a vergence scanning mechanism that pulls the retina back and forth in the focal plane (Franceschini and Changeux,

1994). The functional role of this retinal movement is still open to investigation, but one purpose is clearly that it enhances the sensor's visual range. Shimoyama and his co-workers have designed and built a millimeter-scale model of a compound eye (Hoshino et al., 2000). Not only is there a non-homogeneous arrangement of the microlens array – ranging from 2 degrees between two “ommatidia” in the center, to 6 degrees at the periphery – but there is active scanning in that the retina can be moved back and forth.

Note that in this case, there is generation of sensory data through a sensory-motor coordination: as the retina is moved back and forth, data are generated that enable the organism to cover the area directly in front of the fly that would be otherwise inaccessible (for motion detection) for geometric reasons. Once again, we see the importance of the motor system for the generation of sensory signals, or more generally for perception. It should also be noted that these motor actions are physical processes, not computational ones, but they are computationally relevant, or put differently, relevant for neural processing. This is another demonstration that not all physical actions can be replaced by computation.

Not only do the retinas of insects have non-homogeneous morphology, but the retinas of mammals, including humans, are heterogeneous as well: the spacing in the center is more dense than on the periphery, which is unlike standard cameras in which the distribution of the light-sensitive cells is homogeneous. One of the reasons why animals can process visual signals so rapidly is that the retina already does a lot of pre-processing before the signals are sent on for further processing. This massively parallel peripheral processing ability is crucial to achieve real-time behavior. And in order to do this processing, morphology has to be exploited. This idea is also exploited in the design of artificial retinas and has a long history, as well as generally in the field of space-variant sensing (e.g. Ferrari et al., 1995).

We now turn to the motor system.

Motor systems

In this section we present three case studies, the “passive dynamic walker”, “Stumpy”, and the quadruped “Puppy” that can all be used to explain the concept of “ecological balance” as well as the principle of “cheap design”. While the passive dynamic walker has no actuation, “Stumpy” and the quadruped “Puppy” are equipped with simple artificial muscles.

The passive dynamic walker

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 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 consists only of inclines of certain angles. Energy-efficiency is achieved because the leg movements are entirely passive, driven only by gravity in a pendulum-like manner. 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 constrain 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 Asimo.

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 different types 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 highly complex, like the Honda robot, there is nothing in principle that prevents the exploitation of its passive dynamics. In human walking for example – and humans are certainly highly complex systems – 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 stairs, pushing a cart, opening a door, etc., implying that its ecological niche is considerably larger than that of the passive dynamic walker.

In terms of the design principles, this case study illustrates the principles of cheap design and ecological balance. The passive dynamic walker fully exploits the fact that it is always put on inclines that provide its energy source and

generates the proper dynamics for walking. Loosely speaking, we can also say that the control task, the neural processing, is taken over by having the proper morphology and the right materials. In fact, the neural processing reduces to zero. At the same time, energy efficiency is achieved. However, if anything is changed, e.g. the angle of the incline, the agent ceases to function. This is the trade-off of cheap design. In order to make it adaptive, we would have to add redundancy. There is no contradiction between cheap design and redundancy: even highly redundant systems such as humans exploit the givens of an ecological niche (e.g. gravity, friction, motion parallax).

Even though the passive dynamic walker is an artificial system (and a very simple one), it has a very natural feel to it. The term “natural” not only applies to biological systems, but artificial systems also have their intrinsic natural dynamics. Perhaps the natural feel comes from the exploitation of the dynamics, e.g. the passive swing of the leg.

In conclusion, as suggested by the principle of ecological balance, there is a kind of trade-off or balance: the better the exploitation of the dynamics, the simpler the control, and the less neural processing will be required.

Muscles – control from materials

The passive dynamic walker had no actuation. However, the energy efficiency of this approach can be preserved by incrementally adding actuation. This has been done by Martijn Wisse and his colleagues at Delft University in Holland (Wisse and Frankenhuyzen, 2003) in the construction of the almost passive dynamic walking robot “Mike”. Mike uses pneumatic actuators, which is a kind of artificial muscle: it consists of a rubber tube embedded in a fabric and contracts when air pressure is applied (more about artificial muscles below). We now present two case studies where very simple types of artificial “muscles” are used, which employ elastic materials (in “Stumpy” and in the springs of the quadruped “Puppy”).

Cheap diverse locomotion – “Stumpy”: Recently, there has been an increased interest in applying and further investigating these ideas through the construction of robots. An illustrative example is the walking and hopping robot Stumpy (Iida et al., 2002; 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 each. 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 Iida et al., 2002 and Paul et al., 2002).




Figure 2. The dancing, walking, and hopping robot Stumpy. (a) Photograph of the robot. (b) Schematic drawing (details, see text).

Let us briefly summarize the ideas concerning ecological balance, i.e. 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, control can often be significantly simplified while maintaining a certain level of behavioral diversity. Third, materials have intrinsic control properties. And fourth, because ecological balance is exploited, Stumpy displays a surprisingly diverse behavior (dancing, walking, and hopping in different ways). In this sense, Stumpy also illustrates the diversity-compliance principle: on the one hand, it exploits the physical dynamics in interesting ways and on the other it displays high diversity.

Cheap, rapid locomotion – the quadruped “Puppy”: Another case study that nicely illustrates the principle of ecological balance, is the quadruped “puppy” developed by Iida (Iida and Pfeifer, submitted).

One of the fundamental problems in rapid locomotion is that the feedback control loops, as they are normally used in walking robots, can no longer be used as the response times are too slow. One way to go about this problem is to try and eliminate as much as possible the need for sensory feedback. How this can be done is demonstrated in what follows. In addition, it is shown how rapid locomotion can be achieved through slow but powerful actuation. One of the fascinating aspects of the quadruped “Puppy” is that

not only fast but also robust locomotion can be achieved with no sensory feedback.

The design of “Puppy” was inspired by biomechanical studies. Each leg has two standard servomotors and one springy passive joint. It carries eight motors, batteries, and a micro-controller. To demonstrate a bounding gait using this robot, we applied a synchronized oscillation based control to four motors in the hip and shoulder (the motors in the elbows and knees are not used for this experiment), where each motor oscillates through sinusoidal position control. No sensory feedback is used for this controller except for the internal local feedback for the servomotors.




Figure 3. The quadruped “Puppy”. (a) Picture of entire “puppy”. (b) Diagram showing joints, servomotor actuated joints [circles with crosses], and flexible spine [dotted line]. (c) The spring system in the hind legs.

The legs exhibit simple oscillations, but in the interaction with the environment, through the interplay of the spring system, the flexible spine (note that the battery is attached to the elastic spine which provides precisely the proper weight distribution), and gravity, a natural quadrupedal gait occurs, which includes periods in which all four legs are off the ground: in other words, there is a clear distinction between a stance and a flight phase. The system has self-stabilizing characteristics: there are no sensors on the robot.

We found that this successful demonstration of bounding behavior relies strongly on the underlying mechanism of self-stabilization. The control of the robot is extremely simple: the controller does not distinguish between the stance/flight phase, acceleration, or inclination. Nevertheless, the robot maintains a stable periodic gait by properly exploiting its intrinsic dynamics. It is interesting to note that the foot-ground contact must exhibit little friction in order for this self-stabilization to work.

The self-stabilization property can now be used to control forward velocity by simply varying a single phase

parameter of the oscillation, i.e. a temporal delay between the fore and hind leg motors (the speed can vary between 20 and 50cm/sec). This is a nice illustration of the necessity of ecological balance in an adaptive agent. There are a number of other parameters which could control the forward velocity apart from the phase such as oscillation frequency, amplitude, and the spring constants (if possible). However, an important aspect here is that an indirect parameter, e.g. the phase, could be used for the control of the forward velocity.

This case study demonstrates that complicated and sophisticated designs typically employed in mechatronics are not always required in order to achieve behavioral diversity. Rather, the interaction of body dynamics (as determined by materials and mass distribution), environment (friction, shape of the ground), and control (amplitude, frequency) can be exploited for control purposes.

Reaching and grasping – sensory-motor coordination:

Let us pursue the idea of exploiting the dynamics a little further. Most robot arms available today work with rigid materials and electric motors. Natural arms, by contrast, are built of muscles, tendons, ligaments, and bones, materials that are non-rigid to varying degrees. All of these materials have their own intrinsic properties such as 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. Reaching for and 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: in a sense 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 returns to its resting position. Also, the human arm exhibits 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. The main task of the brain, if you like, is to set the material properties of the muscles, the spring constants. Once these constraints are given, the control task is much simpler.

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.

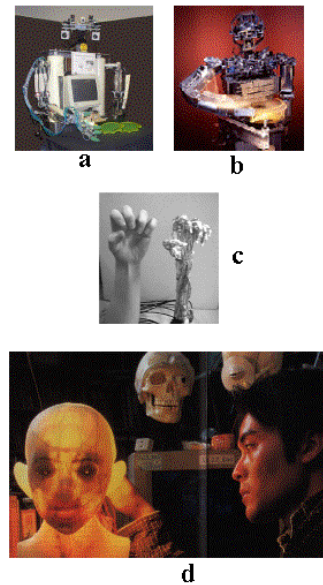


Figure 4. Robots with artificial muscles. (a) The service robot ISAC by Peters (Vanderbilt University) driven by McKibben pneumatic actuators. (b) The artificial hand by Lee and Shimoyama (University of Tokyo), driven by pneumatic actuators. (c) The humanoid robot Cog by Rodney Brooks (MIT AI Laboratory), driven by series-elastic actuators. (d) The “Face Robot” by Kobayashi, Hara, and Iida (Science University of Tokyo), driven by shape-memory alloys.

Facial expressions also provide an interesting illustration for our point about material properties. 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.

The important point here is that we are not simply replacing one type of actuator – an electrical motor – by a different type. This would not be very interesting. The point is that the new type of actuator – e.g. a pneumatic one – has intrinsic physical properties such as elasticity and damping that can be exploited by the neural control.

In section 2 we postulated a set of design principles for adaptive motion. The principle of ecological balance, for example, tells us that given a particular task environment, there is an optimal task distribution between morphology, materials, and control. The principle of emergence asks the question of how a particular “balance” has emerged: how it has come about. In the study of biological systems, we can speculate about this question. However, there is a possibility of systematically investigating this balance, by

using artificial evolution and morphogenesis. Pertinent experiments promise a deeper understanding of these relationships. We also postulated the principle of “sensory-motor coordination.” A fascinating question would be whether evolved agents would employ this principle. There is some evidence that this is indeed the case (e.g. Beer, 1996; Nolfi and Floreano, 2001; Pfeifer and Scheier, 1999). The remainder of this paper will be devoted to exploring the design principles in an evolutionary context.

4. Designing embodied agents using artificial evolution and morphogenesis

Using artificial evolution for design has a tradition in the field of evolutionary robotics (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 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 increases exponentially. 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 on a good solution.

This issue can be approached in various ways; here we just mention two. The first which we will not further discuss in this paper is to parameterize the shapes, thus introducing designer biases as to the types of shapes that are possible. An example that has stirred a lot of commotion in the media is provided by Lipson and 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. 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, such as genetic regulatory networks.

The mechanics of artificial genetic regulatory networks

Here we provide only a non-technical introduction to genetic regulatory networks (for details, see, e.g. (Bongard and Pfeifer 2001), or (Bongard 2002, 2003)). It should be stressed that although this computational system is biologically inspired, it does not constitute a biological model; rather, it is a 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.

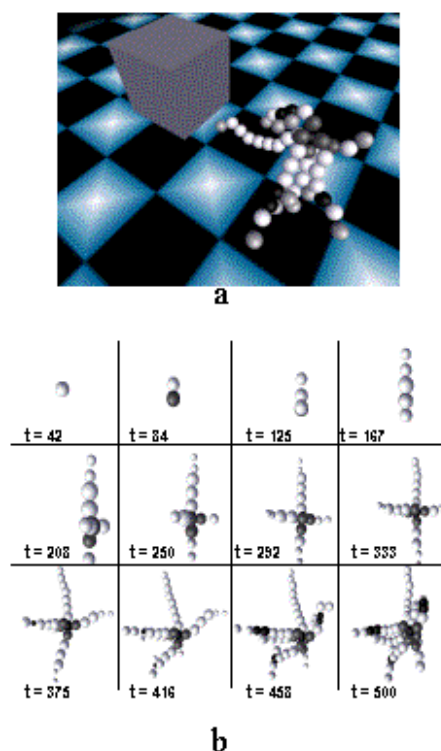


Figure 5. Examples of Bongard’s “block pushers”. (a) An evolved agent in its physically realistic virtual environment. (b) growth phase starting from a single cell, showing various intermediate stages (last agent after 500 time steps). Spheres with different shading indicate differentiation.

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 modify. 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 6. 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 (there are 20 different regulatory TSs, so this value is rounded into the range [0,19]. P2: the TF the gene emits if expressed (there are 43 TFs [20 regulatory and 23 that cause phenotypic change], so this value is rounded into the range [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 regulatory TF concentrations for which the gene is expressed.

Emergence – the achievements of artificial evolution and morphogenesis

Although simple in their basic form, these mechanisms 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 similar kind of control. This is in fact an instance of the equivalent of what is termed “exaptation” in biology: the exploitation of structures for a new function (in this case pushing with a tentacle) that were originally evolved for a different purpose (in this case locomotion). Because they have been created by artificial evolution and morphogenesis, they are, in some sense, ecologically balanced (for this particular task environment). (3) There is no direct relation between genotype length and phenotypic fitness – the two are largely dissociated. But of course, very short genomes can not produce highly complex phenotypes. (4) There is

functional specialization, i.e. cells differentiate into units containing both sensors and actuators (the white colored cells in figure 5), cells that only contain sensors but no actuators (gray coloring), or cells not containing anything, thus only providing passive 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. As can be seen in Figure 5, there is always a white sphere followed by a light gray one at the end of the tentacles. Inspection of the specific neural patterning in these pairs of spheres shows that they differ to some extent from one another. An example from biology is fingers that are similar but exhibit graded differences. (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. 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. Artificial evolution not only shapes the agent’s morphology and neural controller, but the interaction between its morphology, controller, and its environment, i.e. its ecological balance.

4. Discussion and conclusions

We have argued that there is still a lack of consensus in the field of adaptive behavior on its theoretical foundations. By employing design principles we have attempted to take a first step in the direction of providing a coherent framework for design. In the present form we have proposed the principles and have argued why they are plausible. The passive dynamic walker, the quadruped “Puppy”, and “Stumpy” provide illustrations of the principles of cheap design and ecological balance.

While this is acceptable and interesting, the design principles would be much more compelling and powerful if they could be demonstrated to emerge from an evolutionary process (which is one of the messages of the principle of emergence). Using the principles of genetic regulatory networks, we have worked out methods by which entire agents can be evolved, including their

morphology, their material properties, and their control systems.

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” borrowed from Gruau’s cellular encoding algorithm (Gruau, 1994) 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 such as energy supply, blood sugar level, operating temperature, etc. 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. As we have seen earlier, material properties are essential as they can be exploited and often lead to more simple and robust control (see, for example, the case study on the quadruped “Puppy”). And last but not least, we need to be able to increase the complexity of our task environments which requires much more computational power. The incorporation of realistic and sophisticated physics will be an essential criterion: only if the environment is sufficiently “responsive” to an agent’s interaction will there be a need for evolution to increase an agent’s complexity.

At the moment we are confined to simulation; the experiments with artificial systems that can grow physically in the real world are only in their very initial stages (for example in the field of nano-technology, growth mechanisms are investigated). 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 the environment, of the evolved individual and its interactions with the real world (e.g. gravity, impact, friction), and 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, using the evolved creatures as blueprint. 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 raising an issue that is always in the air when working with relatively simple systems (such as block pushers), the one of scalability. By scalability we mean in this context whether the methods proposed (simulated genetic regulatory networks) will be sufficiently powerful to evolve much more complex creatures capable of many behaviors in very different types of environments. We believe this question is still open as it is not clear to what extent the real world plays

an essential role in evolution, or whether simulated environments can be made sufficiently complex.

One hope is, for example, that as the environments and agents become more complex, performing 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.

Earlier, we pointed out the importance of sensory-motor coordination for the development of higher levels of intelligence. This idea is inspired by developmental studies on humans in which the ability for sensory-motor coordination is considered essential: through sensory-motor coordination correlations within and between different sensory modalities are induced, thus facilitating the learning process. Because of its obvious advantages we would expect our fittest agents to engage in sensory-motor coordination. To display its real power, the task environment would have to be more demanding than mere locomotion. We started with a series of experiments where agents have to grasp a large object, a task that obviously requires sensory-motor coordination skills (Bongard, 2003). However, these experiments are too preliminary to be conclusive, except that it seems possible, in principle to evolve grasping in complex agents.

Related to the issue of scalability is the fact that we are building and evolving agents for relatively simple or “low-level” sensory-motor tasks like locomotion and pushing. What is the relevance of these studies for higher forms of cognition? We speculate that by looking at motion, locomotion, and orientation, we can approach the grounding problem, i.e. how categorization – the ability to make distinctions in the real world – and abstract concepts emerge within the growing organism. Metaphorically speaking, a “body image” provides the foundation even for abstract mathematical concepts, as, for example, argued by Lakoff and Nunez (2000). And, we suspect, that in the development of the body image, basic sensory-motor skills play an essential role. While we can expect significant contributions from developmental robotics, in which the system to be modeled is human babies that grow into adults, these studies will always start from the assumption of a given, human-like morphology. What we propose here is more fundamental because we want to explore what sorts of morphologies evolve and how they relate to task environment and neural control. The fact that we can evolve other kinds of morphologies, we can study not only “life as it is”, but “life as it could be”, and experience has taught that we can learn a lot by doing things differently from nature.

Case studies such as the ones presented, and experiments with artificial evolution and morphogenesis will hopefully lead to a better understanding of how biological intelligence evolved and to a better understanding of intelligence in general. In this sense, we feel that “new robotics” can make a significant contribution.

Acknowledgments

The first author (Pfeifer) would like to thank the members of the Artificial Intelligence Laboratory for many discussions, in particular Gabriel Gómez for discussing all aspects of the manuscript. Credit also goes to the Swiss National Science Foundation for supporting the research presented in this paper, grant # 20-61372.00.

References

- Beer, R.D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In P. Maes, M. Mataric, J. Meyer, J. Pollack and S. Wilson (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. pp. 421-429. MIT Press.
- Berthouze, L., (2000). Bootstrapping the developmental process: the Filter Hypothesis. In J. Demiris and A. Birk (eds.) *Robot Learning - An Interdisciplinary Approach*, Series in Robotics and Intelligent Systems Vol. 24, World Scientific, 8- 30
- 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.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA.: MIT Press.
- Breazeal, C.L. (2002). *Designing Sociable Robots (Intelligent Robotics and Autonomous Agents)*. Cambridge, Mass.: MIT Press.
- Brooks, R. A. (1991). *Intelligence Without Reason*. Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91), pp. 569-595
- Bryson, J.J. (2002). *Intelligence by design: Principles of modularity and coordination for complex adaptive agents*. Unpublished PhD Dissertation. Massachusetts Institute of Technology.
- 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.
- Cruse, H., Dean, J., Dürr, V., Kindermann, T., Schmitz, J. and Schumm, M. (2002). A decentralized, biologically based network for autonomous control of (hexapod) walking. In J. Ayers et al. (eds.) *Neurotechnology for Biomimetic Robots*. Cambridge, MA: MIT Press, 383-400.
- 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.
- Eiben, A. E. and Schippers, C. A. (1998). On evolutionary exploration and exploitation. *Fundamenta Informaticae*, 35, 35-50.
- Ferrari, F., Nielsen, P.Q.J., and Sandini, G. (1995). Space variant imaging. *Sensor Review*, 15, 17-20.
- Franceschini, N., Pichon, J.M., and Blanes, C. (1992). From insect vision to robot vision. *Philos. Trans. R. Soc. London B.*, 337, 283-294.
- Grossberg, S. (1995). The attentive brain. *American Scientist*, 83, 438-449.
- Gruau, F. (1994). Automatic definition of modular neural networks. *Adaptive Behaviour*, 3, 151-183.
- 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.
- Hoshino, K., Mura, F., and Shimoyama, I. (2000). Design and performance of a micro-sized biomorphic compound eye with a scanning retina. *Journal of Microelectromechanical Systems*, 9, 32-37.
- Iida, F., Dravid, R. and Paul, C. (2002) 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 Proceedings of International Conference on Intelligent, pp.2141-2146.
- Iida, F. and Pfeifer, R. "Cheap" Rapid Locomotion of a Quadruped Robot -Self-Stabilization of Bounding Gait-, *Intelligent Autonomous Systems 8 (to appear)*
- Ishiguro, A., Ishimaru, K., Hayakawa, K., and Kawakatsu, T. (2003). Toward a "Well-Balanced" Design: A Robotic Case Study -How should Control and Body Dynamics be Coupled?-, in *Proc. of The 2nd International Symposium on Adaptive Motion of Animal and Machines*
- 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 Nunez, R. (2000) Where mathematics comes from. *New York: 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.
- Lichtensteiger, L., and Pfeifer, R. (2002). An Optimal Sensor Morphology Improves Adaptability of Neural Network Controllers. In: Jose R. Dorronsoro (Ed.): *Artificial Neural Networks - ICANN 2002*, International Conference, Lecture Notes in Computer Science, LNCS 2415, Springer, Heidelberg. pp.850-855
- Liu, J. (1999). *Autonomous agents and multi-agent systems. Explorations in learning, self-organization and adaptive computation*. Singapore: World Scientific.
- Maes, P. (1993). Behavior-based artificial intelligence. In J.A. Meyer, H.L. Roitblat, and S.W. Wilson (Eds.). *From Animals to Animats. Procs. 2nd Int. Conf. on the Simulation of Adaptive Behavior*,. Cambridge, MA: MIT Press, 2-10.
- Manzotti, R. (2000). Intentional robots. The design of a goal-seeking environment-driven agent. Unpublished Doctoral Dissertation, University of Genova.
- 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.
- Newell, A., and Simon, H.A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM (Association for Computing Machinery)*, 19, 113-126.
- Nolfi, S. and Floreano, D. (2001). *Evolutionary robotics: the biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- Paul, C., Dravid, R. and F. Iida (2002) Control of lateral bounding for a pendulum driven hopping robot. *Proceedings of the International Conference of Climbing and Walking Robots (CLAWAR 2002)*, pp. 333-340.
- 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. (2000a). 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. (2000b). 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. (2002). Robots as cognitive tools. *Int. Journal of Cognition and Technology*, 1, 125-143.
- Pfeifer, R. (2003). Morpho-functional machines: basics and research issues. In F. Hara, and R. Pfeifer (eds.). *Morpho-functional machines: the new species*. Tokyo: Springer, 2003.
- Pfeifer, R., and Glatzeder, B. (in preparation). *How the body shapes the way we think: the embodied revolution in artificial intelligence*. Cambridge, Mass.: MIT Press.
- Pfeifer, R., and Scheier, C. (1999). *Understanding intelligence*. Cambridge, Mass.: MIT Press.
- Piaget, J. (1970). Piaget's theory. In P.H. Mussen (ed.) Carmichael's manual of child psychology. New York, N.Y.: Wiley, 703-732.
- Picard, R.W. (1997). *Affective computing*. 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.
- Seth, A. K. (2002). Agent-based modeling and the environmental complexity thesis. In J. Hallam et al (eds.). *Proceedings of the Seventh International Conference on the Simulation of Adaptive Behaviour*. Cambridge, MA: MIT Press, 13-24.
- 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.
- Sims, K. (1994a). Evolving virtual creatures. *Computer Graphics*, 28, 15-34.
- Sims, K., (1994b). Evolving 3D morphology and behavior by competition. In R. Brooks, and P. Maes (eds.). *Artificial Life IV Proceedings*, 28-39.
- Smithers, T. (1995). On quantitative performance measure of robot behaviour. *Robotics and Autonomous Systems*, 15, 7-133.
- Steinhage, A. and Schoener, G. (1997). Self-Calibration Based on Invariant View Recognition: Dynamic approach to Navigation, Robotics and Autonomous Systems 20, 133-156
- 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, October 29 - November 3