

Designing intelligent robots – on the implications of embodiment

ROLF PFEIFER, FUMIYA IIDA and GABRIEL GOMEZ

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

Abstract

Traditionally, in robotics, artificial intelligence, and neuroscience, there has been a focus on the study of the control or the neural system itself. Recently there has been an increasing interest into the notion of embodiment in all disciplines dealing with intelligent behavior, including psychology, philosophy, and linguistics. In this paper, we explore the far-reaching and often surprising implications of this concept. While embodiment has often been used in its trivial meaning, i.e. „intelligence requires a body“, there are deeper and more important consequences, concerned with connecting brain, body, and environment, or more generally with the relation between physical and information (neural, control) processes. It turns out that, for example, robots designed by exploiting embodiment are frequently simpler, more robust and adaptive than those based on the classical control paradigm. Often, morphology and materials can take over some of the functions normally attributed to control, a phenomenon called “morphological computation”. It can be shown that through the embodied interaction with the environment, in particular through sensory-motor coordination, information structure is induced in the sensory data, thus facilitating perception and learning. A number of case studies are presented to illustrate the concept of embodiment. We conclude with some speculations about potential lessons for robotics.

Keywords: Embodiment, morphological computation, information theoretic implications of embodiment, self-stabilization.

1. INTRODUCTION

While 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 attention the research community is devoting to using robots for studying intelligent systems and to the development of robots that share their ecological niche with humans. If we are to design these kinds of robots, embodiment must be taken into account.

There is an extensive literature in cognitive science, neuroscience, psychology, and philosophy on the conceptual and empirical groundwork for embodiment and embodied cognition (e.g. Anderson, 2003; Clark, 1999; Gallagher, 2005; Iida et al., 2004; Pfeifer et al., submitted; Lakoff and Johnson, 1999; Pfeifer and Scheier, 1999; Pfeifer and Bongard, 2006; Smith and Gasser, 2005; Sporns, 2003; Wilson, 2002; Ziemke, 2002). Inspired by this body of work, in this paper we address specifically how these ideas and insights can be fruitfully exploited for robotics research.

There is a trivial meaning of the term embodiment, namely that “intelligence requires a body”. It is obvious that if we are dealing with a physical agent, we have to take gravity, friction, torques, inertia, energy dissipation, etc. into account. However, there is a non-trivial meaning of embodiment, which relates to the interplay between brain, body, and environment, or more generally the physical and the information theoretic processes underlying an agent’s behavior. One simple but fundamental insight, for example, is that whenever an agent behaves in whatever way in the physical world, it will by its very nature of being a physical agent, affect the environment and in turn be influenced by it, and it will induce – generate – sensory stimulation. The sensory signals caused by the movement will, depending on the kind of behavior, have certain properties, and typically they will be correlated. For example, if you walk in the street optic flow will be induced in your visual sensors, and tactile and proprioceptive stimulation is generated in your feet and motor system. Thus, there is a continuous tight interaction between the motor system and the various sensory systems, i.e. there is a sensory-motor coordination. Typically, behavior in natural – and adaptive artificial – agents is sensory-motor coordinated. It turns out that through sensory-motor coordinated interaction with the environment, information structure (e.g. spatio-temporal correlations in a visual input stream, redundancies between different perceptual modalities) is induced in the various sensory channels which facilitates perception and learning. This insight which is among the most powerful and extensive implications of embodiment is an instance of what we call “information theoretic implications of embodiment.”

By this latter term we mean the effect of morphology, materials, and environment on neural processing, or better, the interplay of all these aspects. Materials, for example, can take over some of the processes normally attributed to control, a phenomenon called “morphological computation”. Although in an embodied agent, by the mere fact of its being physical, all aspects – sensors, actuators, limbs, the neural system – are always highly connected, 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. We start with sensor morphology which is followed by two locomotion examples. We then show how morphology and materials can be exploited for grasping using an artificial hand. We then turn to sensory-motor coordination and will try to integrate what has been said so far into a coherent picture. Finally, we will discuss what has been achieved, and what lessons there might be for robotics research.




Figure 1: Morphological computation through sensor morphology – the Eyebot. The specific non-homogeneous arrangement of the facets compensates for motion parallax, thereby facilitating neural processing. (a) Insect eye. (b) Picture of the Eyebot. (c) Front view: the eyebot consists of a chassis, an on-board controller, and sixteen independently-controllable facet units, which are all mounted on a common vertical axis. (d) Schema of Facets: Each facet unit consists of a motor, a potentiometer, two cog-wheels and a thin tube containing a sensor (a photo diode) at the inner end. These tubes are the primitive equivalent of the facets in the insect compound eye.

2. Sensor morphology: navigation in insects and robots

In previous papers we have investigated in detail the effect of changing sensor morphology on neural processing (e.g. Lichtensteiger, 2004; Pfeifer, 2000, 2003; Pfeifer and Scheier, 1999). Here we only summarize the main results. The morphology of the sensory system has a number of important implications. In many cases, when it is suited for the particular task environment, more efficient solutions can be found. For example, it has been shown that for many tasks (e.g. obstacle avoidance) motion detection is all that is required. Motion detection can often be simplified if the light-sensitive cells are not spaced evenly, but if there is a non-homogeneous arrangement. For instance, 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). This non-homogeneous arrangement, in a sense, compensates for the phenomenon of motion parallax, i.e. the fact that at constant speed, objects on the side travel faster across the visual field than objects towards the front: it performs the “morphological computation”, so to speak. 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, a principle that has also been applied to the construction of navigating robots (e.g. Hoshino et al., 2000). 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 without changing anything inside the controller (Lichtensteiger, 2004; Figure 1). Note that all this only works, if the agent is actually behaving in the real world and therefore is generating sensory stimulation.

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, which is why we use the term “morphological computation”.

3. Morphology of body and motor systems: locomotion

In this section we introduce two locomotion robots, the quadruped robot “Puppy”, which demonstrates the exploitation of materials and dynamics of the system-environment interaction, and the artificial fish “Wanda” which illustrates the exploitation of the environment to generate behavior. In the design of “Puppy”, a very simple kind of artificial “muscle” in the form of a normal spring is used. 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 employed because the response times are too slow. One of the fascinating aspects of “Puppy” is that not only fast but also robust locomotion can be achieved with no sensory feedback (Iida and Pfeifer, 2004).

The design of “Puppy” was inspired by biomechanical studies. Each leg has two standard servomotors and one springy passive joint (Figure 2). To demonstrate a running gait, we applied a synchronized oscillation based control to the four motors – two in the “hip” and two in the “shoulder” – where each motor oscillates through sinusoidal position control (i.e., amplitude and frequency of the motor oscillation). No sensory feedback is used for this controller except for the internal local feedback for the servomotors (Iida and Pfeifer, 2006).

Even though the legs are actuated by simple oscillations, in the interaction with the environment, through the interplay of the spring system, the flexible spine, and gravity, a natural running gait emerges. The controller of the robot is extremely simple and because it has no sensors, it cannot distinguish between the stance/flight phase, and it cannot measure acceleration, or inclination.

Nevertheless, the robot maintains a stable periodic gait, which is achieved by properly exploiting its intrinsic dynamics. The behavior is the result of the complex interplay of agent morphology, material properties (in particular the “muscles”, i.e. the springs), control (amplitude, frequency), and environment (friction and slippage, shape of the ground, gravity). Exploiting morphological properties and the intrinsic dynamics of materials makes “cheap” rapid locomotion possible because physical processes are fast – and they are for free, so to speak! (for further references on cheap locomotion, see e.g. Kubo and Full, 1999; Blickhan et al., 2003; Buehler, 2002; Iida, 2005). Because stable behavior is achieved without control – simply due to the intrinsic dynamics of the physical system – we use the term “self-stabilization”. Now, if sensors – e.g. pressure sensors on the feet, angle sensors in the joints, and vision sensors on the head – are mounted on the robot, structured – i.e. correlated – sensor stimulation will be induced that can potentially be exploited

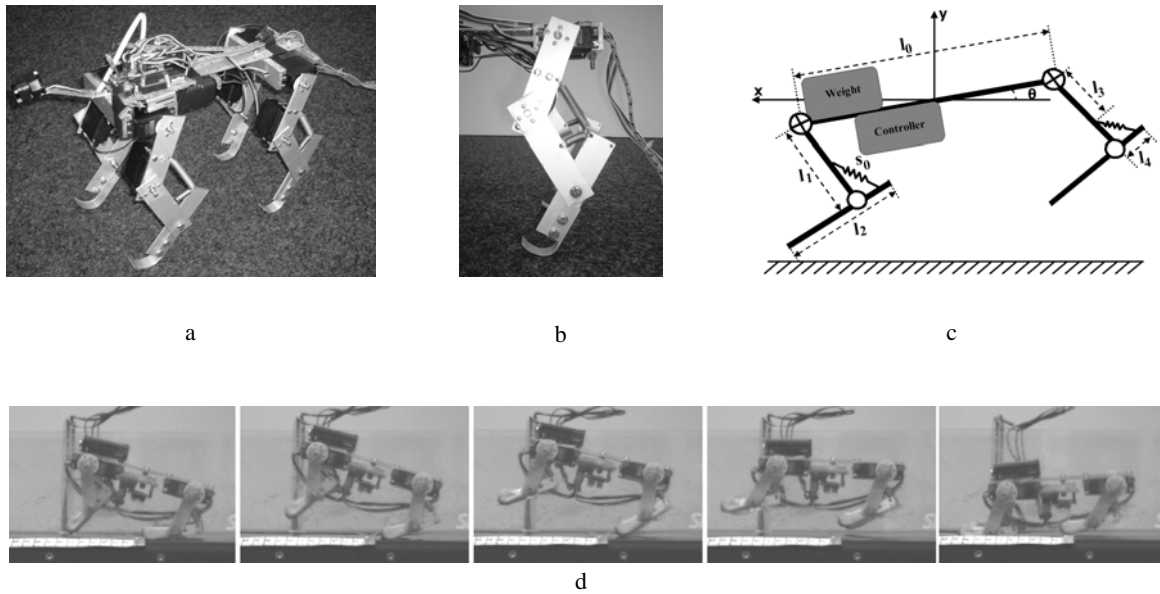


Figure 2: The quadruped “Puppy”. (a) Picture of the entire “Puppy”. (b) The spring system in the hind legs. (c) Schematic design of a slightly modified version of “Puppy”. (d) Running on a treadmill.

for learning about the environment and its own body dynamics (for more detail, see, e.g. Iida, 2005, or Pfeifer and Bongard, 2006, chapter 5).

The second locomotion case study concerns the artificial fish, “Wanda”, developed by Marc Ziegler and Fumiya Iida (Ziegler et al., 2005; Pfeifer and Iida, 2005; Figure 3). It shows how the interaction with the environment can be exploited in interesting ways to achieve a task: “Wanda” can reach any position in 3D space with only one degree-of-freedom (DOF) of actuation. One DOF means that it can basically wiggle its tail fin back and forth. The tail fin is built from elastic materials such that it will on average produce maximum forward thrust. It can move forward, left, right, up and down. Turning left and right is achieved by setting the zero-point of the wiggle movement either left or right at a certain angle. The buoyancy is such that if it moves forward slowly, it will sink, i.e. move down gradually. The speed is controlled by the wiggling frequency. If it moves fast and turns, its body will tilt slightly to one side which produces upthrust, so that it will move upwards. The fascinating point about this fish is the behavioral diversity that can be achieved in this way, through “morphological computation”. If material properties and the interaction with the environment are not properly exploited to achieve flexible movement in the water, one would need more complicated actuation, e.g. additional fins or a flexible spine and thus more complex control.




Figure 3: The artificial fish “Wanda”. (a) With one degree-of-freedom for wiggling the tail fin. (b) The forces acting on its body are illustrated by arrows. (c) A typical sequence of snapshots of an upward movement.

4. Complex motor systems: Grasping

In this case study, we discuss how morphology, materials, and control interact to achieve grasping behavior. The 13 degrees-of-freedom “Yokoi hand” (Yokoi et al., 2004; Figure 4) which can be used as a robotic and a prosthetic hand, is partly built from elastic, flexible, and deformable materials (the “Yokoi hand” comes in many versions with different materials, morphologies, sensors, etc.; here we only describe one of them). For example, the tendons are elastic, the finger tips are deformable and between the fingers there is also deformable material. When the hand is closed, the fingers will, because of its anthropomorphic morphology, automatically come together. For grasping an object, a simple control scheme, a “close” is applied. Because of the morphology of the hand, the elastic tendons, and the deformable finger tips, the hand will automatically self-adapt to the object it is grasping. Thus, there is no need for the agent to “know” beforehand what the shape of the to-be-grasped object will be. The shape adaptation is taken over by the morphology of the hand, the elasticity of the tendons, and the deformability of the finger tips, as the hand interacts with the shape of the object. In this setup, control of grasping is very simple, or in other words, very little “brain power” is required for grasping.

A certain degree of self-regulation of the hand can also be achieved by having relatively non-elastic tendons, and by putting pressure sensors on the fingertips and stopping to apply torques to the tendons whenever a certain pressure threshold is exceeded. Note that this kind of control requires more computation compared to the case with the elastic tendons. On the other hand, compared to elastic tendons, higher torques can be obtained. Of course, by placing additional sensors on the hand – e.g. for angle or torque – the abilities of the hand can be improved and feedback signals can be provided to the agent (the robot and the human) which can then be exploited by the neural system for learning purposes.

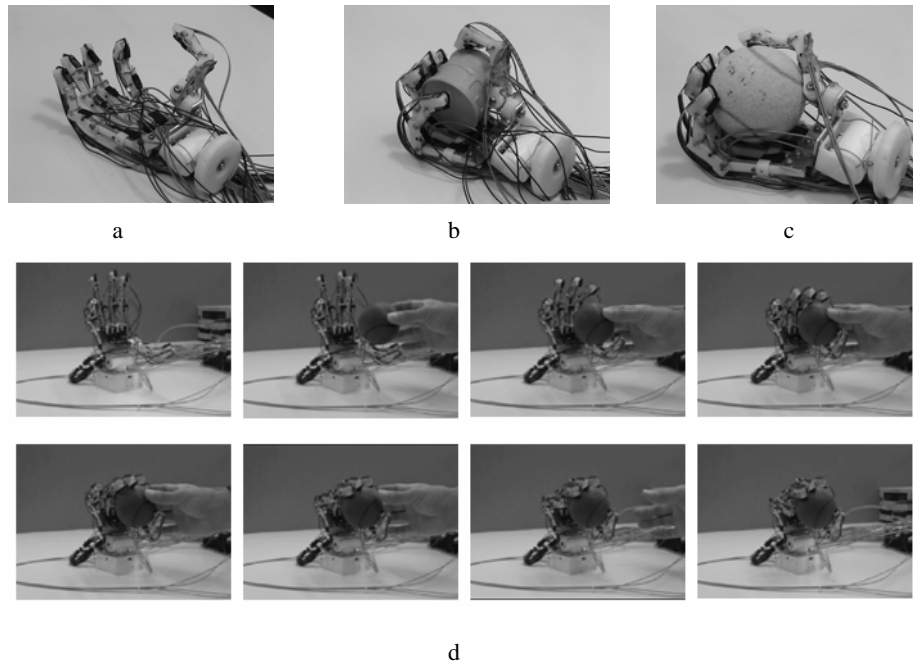


Figure 4: “Cheap” grasping: exploiting system-environment interaction. (a) The Yokoi hand exploits deformable and flexible materials to achieve self-adaptation through the interaction between environment and materials. (b)-(c) Final grasp of different objects. The control is the same, but the behavior is very different. (d) Sequence of a typical grasping experiment.

For prosthetics, there is an interesting implication. EMG signals can be used to interface the robot hand non-invasively to a patient: even though the hand has been amputated, he or she can still intentionally produce muscle innervations which can be picked up on the surface of the skin by EMG electrodes. If EMG signals, which are known to be very noisy, are used to steer the movement of the hand, control cannot be very precise and sophisticated. But by exploiting the self-regulatory properties of the hand, there is no need for very precise control, at least for some kinds of grasping: the relatively poor EMG signals are sufficient for the basic movements.

5. Bringing it all together: Sensory-motor coordination

Grasping is a particularly important instance of the general class of sensory-motor coordinated actions. Of course, the hand is only useful, if it forms part of an arm of an agent which (or who) is also equipped with other sensory modalities such as vision and proprioception: the agent can then engage in sensory-motor coordination. Because of its almost universal presence in behaving organisms, sensory-motor coordination has been widely studied in psychology, neuroscience, and robotics (e.g. Dewey, 1896; Piaget, 1953; Edelman, 1987; Pfeifer and Scheier, 1999; Lungarella et al., 2005;

Harnad, 2005; Poirier et al., 2005). In robotics, an early demonstration of the idea of exploiting coordinated interaction with the environment, is a study by Pfeifer and Scheier (1997) in which it is shown that mobile robots can reliably categorize big and small wooden cylinders only if their behavior is sensory-motor coordinated. The artificial evolution experiments by Nolfi (2002) and Beer (2003) illustrate a similar point: the fittest agents, i.e. those that could most reliably determine the proper category of an object, were those engaging in sensory-motor coordinated behavior. Intuitively, in these examples, the interaction with the environment (a physical process) creates additional (i.e. previously absent) sensory stimulation which is also highly structured, thus facilitating subsequent information (or neural) processing. Metaphorically speaking, computational economy and temporal efficiency are purchased at the cost of behavioral interaction. As mentioned earlier, this sensory stimulation which is generated by the agent itself, has information structure, which is why this is called the principle of information self-structuring (for a quantitative perspective on information structure, see Lungarella et al., 2005).

Let us take this idea a step further. Grasping, because of the particular shape of the hand, is much easier than bending the fingers backwards, and owing to other aspects of morphology (e.g. the high density of touch sensors on the finger tips) rich structured sensory stimulation is induced. The natural movements of the arm and hand are – as a result of their intrinsic dynamics – directed towards the front center of the body. This in turn implies that normally a grasped object is moved towards the center of the visual field thereby inducing correlations in the visual and haptic channels, and these correlations simplify information processing, in particular categorization and learning. So we see that an interesting relationship exists between morphology, intrinsic body dynamics, generation of information structure, and learning. These ideas can be exploited for robot design.

6. Discussion

So far we have demonstrated that by exploiting embodiment, many functions can be achieved more “cheaply”, i.e. with much less computation, or better with less neural control: motion detection, running, swimming, and grasping are all facilitated through “morphological computation”. Moreover, movements corresponding to the natural intrinsic dynamics of an agent not only require less control but they are also more energy-efficient (compared to the case where morphological properties are not exploited). We also argued that through embodied interaction, information structure is induced in the sensory signals from various modalities. The question that immediately comes up is whether these kinds of considerations are confined to “low-level” sensory-motor processes and what their relation is to higher levels of cognition. This is a hard research question and needs much more exploration. But just very briefly, categorization, the ability to make distinctions in the real world, is one of the most fundamental cognitive abilities: a robot, is not going to be worth much if it can’t distinguish a power

outlet from a human baby, or a newspaper from an apple (and analogously for a biological agent). So it seems reasonable to assume that categorization is one of the very basic and essential cognitive abilities – “To cognize is to categorize: cognition is categorization” (Harnad, 2005). We made a case that since embodied interaction leads to induction of information structure, subsequent processing, in particular categorization, is facilitated. This information structure emerges while the interaction takes place (it doesn’t exist before the interaction), however, once it has been induced, learning can pick up on it so that next time around, the responsible sensory-motor information structure is more easily generated. It follows that embodied interaction lies at the root of a powerful learning mechanism as it enables the creation of time-locked correlations and the potential discovery of higher-order regularities that transcend the individual sensory modalities. Also, in a developmental context, the emergent information structure – the resulting intra- and intermodally correlated sensory stimulation – seems to be essential for concept formation (e.g. Gallese and Lakoff, 2005; Bahrick et al., 2004). In this way, concepts can be formed that are naturally grounded in the agent’s embodied interaction with the environment. Of course, with these considerations we have not explained high-level cognition, but we have explored the very basis on top of which the organism – or the robot for that matter – can overlay additional processes.

And a final point: We have been discussing a few implications of embodiment revolving around induction of information structure through embodied interaction with the real world, or more generally the relation between physical and information processes in agent behavior. However, a host of fascinating research where embodiment plays a crucial role, has been – and currently is – conducted in various areas and we cannot possibly do justice to all, we just point out some of the projects here: Yokoi’s “amoeba robot” (Yokoi et al., 1998), and Ishiguro’s “Slimebot” (Ishiguro et al., 2006), where locomotion is emergent from non-linear local interactions between the physical modules; Hosoda and his colleagues’ experiments in developmental robotics, where joint attention behavior emerges from the embodied interaction between human and robot (Nagai et al., 2003); Bovet’s “Artificial Mouse” that develops seemingly goal-directed behavior through embodied interaction with the environment, based on a simple Hebbian-style neural network where all sensors and motors are connected and no goals are programmed into the system (Bovet, 2006); or Kuniyoshi’s “Roll-and-Rise” robot capable of dynamically – using the inertia from a “rolling” movement – standing up from a lying position (e.g. Kuniyoshi et al., 2004). A detailed discussion is beyond the scope of this paper.

7. Conclusions: lessons for robotics

To conclude the paper, let us speculate a bit what potential lessons there might be for robotics. While some of these points are obvious, it is interesting to note that in practical everyday research, they are often not, or not sufficiently taken into account.

First, when applying the ideas outlined in this paper to robot design, we have to be aware of the fact that behavior is always emergent, i.e. it cannot be predicted (and designed) by focusing on control only. The behavior of the fingers in the “Yokoi hand” differs significantly depending on the particular shape of the object, even though the control is the same. In fact, there may be a considerable “conceptual gap” between the control signal and the actual behavior. In Ishiguro’s “Slimebot”, for instance, the motor signals move rods by which the individual modules attach to their neighbors and manipulate their own friction, but the overall behavior of the entire “Slimebot” is a forward movement by means of a global propagating wave. In short, we need to “design for emergence” (Steels, 1991; Pfeifer and Bongard, 2006). A related point is that the control signal may or may not relate to the actual behavior of the robot in a direct way. Thus, we have to be careful when interpreting the actual “meaning” of efference copies (copies of the control system).

Second, a point that for reasons of space, we only mentioned in passing, the environment can take over a considerable amount of the work, if exploited properly. In running behavior, obviously gravity and friction are exploited, and in walking behavior where the forward swing of the leg is largely passive, the fact is exploited that the leg acts like a pendulum and gravity does the job of moving the leg forward, so to speak. The robot fish “Wanda”, although it can only wiggle its tailfin left and right, can move upwards by making a turn which has the effect that the robot slightly tilts to one side so that it will get the required upthrust. But this principle not only applies to very simple systems like our robots: in our – human – everyday behavior we almost routinely exploit the environment – we just may not be aware it.

And third, not everything needs to be controlled by the brain: morphological computation takes over, or distributes computational or control functions to the morphology, materials, and system-environment interaction (e.g. the self-stabilization in “Puppy’s” running behavior, the self-adaptive grasping in the “Yokoi hand”). Recent insights in biomechanics, suggest that in rapid locomotion in animals, an important role of the brain is to dynamically adapt the stiffness and elasticity of the muscles, rather than very precise control of the joint trajectories, because in this way, the muscles can take over some of the control function, e.g. the elastic movement on impact and adaptation to uneven ground (e.g. Blickhan et al., 2003). For robotics, the idea of embodiment provides new ways of looking at behavior generation, because in the past the focus has been very much on the control side.

Acknowledgments

The author would like to thank Akio Ishiguro for the invitation to contribute this review paper. This research was supported by the project “From locomotion to cognition” of the Swiss National Science Foundation (Grant No. 200021-109210/1) and the EU-Project ROBOTCUB: ROBotic Open-architecture Technology for Cognition, Understanding and Behavior (IST-004370).

REFERENCES

- Anderson, M.L. (2003). Embodied cognition: a field guide. *Artificial Intelligence*, 149:91—130.
- Bahrick, L.E., Lickliter, R. and Flom, R. (2004). Intersensory redundancy guides the development of selective attention, perception, and cognition in infancy. *Current Directions in Psychological Science*, 13(3): 99—102.
- Beer, R.D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior*, 11(4):209—243.
- Blickhan, R., Wagner, H., and Seyfarth, A. (2003). Brain or muscles?, *Rec. Res. Devel. Biomechanics*, 1, 215-245.
- Bovet, S and Pfeifer, R. “Emergence of coherent behaviors from homogenous sensorimotor coupling,” in *Proc. 12th Int. Conf. on Adv. Robotics (ICAR)*, Seattle, 2005, pp. 324—330.
- Buehler, M. (2002). Dynamic locomotion with one, four and six-legged robots, *J. of the Rob. Soc. of Japan*, 20(3): 15-20.
- Clark, A. (1999). An embodied cognitive science? *Trends in Cognitive Sciences*, 3(9):345—351.
- Dewey, J. (1896). The reflex arc concept in psychology. *Psychological Review*, 3:357—370.
- Edelman, S. (1999). *Representation and Recognition in Vision*. Cambridge, MA: MIT Press.
- Franceschini, N., Pichon, J.M., and Blanes, C. (1992). From insect vision to robot vision. *Philos. Trans. R. Soc. London B.*, 337, 283-294.
- Gallagher, S. (2005). *How the Body Shapes the Mind*. Oxford: Clarendon Press.
- Gallese, V. and Lakoff, G. (2005). The brain’s concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, 22(3/4):455—479.
- Hara, F., and Pfeifer, R. (2000). On the relation among morphology, material and control in morpho-functional machines. SAB-6. *Proc. of the 6th Int. Conf. on Simulation of Adaptive Behavior*, 33-40.
- Harnad, S. (2005). Cognition is categorization. In: H.Cohen and C.Lefebvre (eds.), *Handbook of Categorization in Cognitive Science*. Elsevier
- 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. and Pfeifer, R. (2006). Sensing through body dynamics, *Journal of Robotics and Autonomous Systems*, (in press).
- Iida, F. (2005). Cheap design and behavioral diversity for autonomous adaptive robots. PhD Dissertation, Faculty of Mathematics and Science, University of Zurich, Switzerland.
- Iida, F. and Pfeifer, R. (2004) “Cheap” rapid locomotion of a quadruped robot: Self-stabilization of bounding gait, *Intelligent Autonomous Systems 8*, F. Groen et al. (Eds.), IOS Press, 642-649.
- Ishiguro, A., Matsuba, H., Maegawa, T., and Shimizu, M. (2006). A modular robot that self-assembles. In *Proceedings of the 9th International Conference on Intelligent Autonomous Systems (IAS-9)* . T. Arai et al. (Eds.), Amsterdam, IOS Press, Tokyo, Japan, 585-594.
- Kuniyoshi, Y., Rougeaux, S., Stasse, O., Cheng, G., and Nagakubo, A. (2000). A humanoid vision system for versatile interaction. *Biological motivated computer vision. LNCS 1811*, 512-526.
- Kubow, T. M., and Full, R. J. (1999). The role of the mechanical system in control: a hypothesis of selfstabilization in hexapedal runners, *Phil. Trans. R. Soc. Lond. B*, 354, 849-861.
- Lakoff, G. and Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. New York: Basic Books.
- Lichtensteiger, L. (2004). On the interdependence of morphology and control for intelligent behavior. PhD Dissertation, University of Zurich.
- Lungarella, M., Pegors, T., Bulwinkle, D. and Sporns, O. (2005). Methods for quantifying the informational structure of sensory and motor data. *Neuroinformatics*, 3(3):243—262.
- Nagai, Y., Hosoda, K., Morita, A., and Asada, M. (2003). A constructive model for the development of joint attention. *Connection Science*, Vol. 15, No. 4, pp. 211-229.
- Nolfi, S. (2002). Power and limit of reactive agents. *Neurocomputing*, 49:119—145
- Pfeifer, R., Lungarella, M., Sporns, O., and Kuniyoshi, Y. (submitted). On the information theoretic implications of embodiment – principles and methods.
- Pfeifer, R., and Bongard, J. (2006). *How the body shapes the way we think: a new view of intelligence*. Cambridge, Mass.: MIT Press.
- Pfeifer, R. and Iida, F. (2005). Morphological computation: Connecting body, brain and environment. *Japanese Scientific Monthly*, Vol. 58, No. 2, 48-54.
- 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. (2000). On the role of morphology and materials in adaptive behavior. In *Proc. SAB (Simulation of Adaptive Behavior) 6*, 23-32.
- Pfeifer, R., and Scheier, C. (1999). *Understanding intelligence*. Cambridge, Mass.: MIT Press.
- Pfeifer, R., and Scheier, C. (1999). *Understanding intelligence*. Cambridge, Mass.: MIT Press.
- Pfeifer, R. and Scheier, C. (1997). Sensory-motor coordination: the metaphor and beyond. *Robotics and Autonomous Systems*, 20:157—178.
- Piaget, J. (1953). *The Origins of Intelligence*. New York: Routledge.

- Poirier, P., Hardy-Vallee, B. and DePasquale, J.F. (2005). Embodied categorization. In: H. Cohen and C. Lefebvre (eds.), *Handbook of Categorization in Cognitive Science*. Elsevier.
- Smith, L. and Gasser, M. (2005). The development of embodied cognition: six lessons from babies, *Artificial Life*, 11(1/2):13—30.
- Sporns, O. (2003). Embodied cognition. In: M. Arbib (ed.), *Handbook of Brain Theory and Neural Networks*, Cambridge, MA: MIT Press.
- Steels, L. (1991). Toward a theory of emergent functionality. In J.A. Meyer and S.W. Wilson (Eds.). *From Animals to Animats*. Proc. SAB-1, 451-461.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin and Review*, 9(4):625—636.
- Yokoi, H., Hernandez, A., Katoh, R., Yu, W., Watanabe, I., and Maruishi, M. (2004). Mutual adaptation in a prosthetics application. In F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi (Eds.). *Embodied artificial intelligence*. Springer LNAI 3139, 146-159.
- Yokoi, H., Yu, W., Hakura, J., and Kakazu, Y. (1998). Morpho-functional machine: robotics approach of amoeba model based on vibrating potential method. In R. Pfeifer, B. Blumberg, J.-A. Meyer, and S.S. Wilson (Eds.). Proc. SAB-98, 297-302.
- Ziegler, M., Iida, F., and Pfeifer, R. (2005). “Cheap” underwater locomotion: Morphological properties and behavioral diversity, IROS05 Workshop on Morphology, Control and Passive Dynamics.
- Ziemke, T. (ed.) (2002). *Situated and embodied cognition*. *Cognitive Systems Research*, 3(3):271—554