Morphological computation: connecting body, brain, and environment

Rolf Pfeifer and Fumiya Iida
Artificial Intelligence Laboratory, Department of Informatics, University of Zurich
Andreasstrasse 15, CH-8050 Zurich, Switzerland
[pfeifer,iida]@ifi.unizh.ch, phone: +41 1 63 5 4320/31, fax: +41 1 635 68 09

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 not only in robotics and artificial intelligenc, but also in the neurosciences, psychology, and philosophy. In this paper, we introduce the notion of morphological computation and demonstrate how it can be exploited on the one hand for designing intelligent, adaptive robotic systems, and on the other for understanding natural systems. While embodiment has often been used in its trivial meaning, i.e. „intelligence requires a body“, the concept has deeper and more important implications, concerned with the relation between physical and information (neural, control) processes. Morphological computation is about connecting body, brain and environment. A number of case studies are presented to illustrate the concept. We conclude with some speculations about potential lessons for neuroscience and robotics.

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. Now 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 work on embodiment. It is also 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, namely that there is a tight interplay between the physical and the information theoretic aspects of an agent, or generally, the information theoretic implications of embodiment. 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. A fish, for example, as it moves, will induce currents and turbulences in the water which then affect its own motion. 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 in your feet and motor system. Optic flow is about correlations in the visual sensors, proprioceptive and tactile stimulation about rhythmic patterns, also containing important correlations. An additional point is that because of the intrinsic physical dynamics there will be certain preferred walking patterns, corresponding to energy-efficient movement. Thus, there is a continuous tight interaction between the motor system and the various sensory systems, a sensory-motor coordination. Typically, behavior in natural – and artificial agents – are sensory-motor coordinated.

Before we continue, a short note on terminology is required. By information theoretic implications of embodiment we mean the effect of morphology, materials, and environment on neural processing, or better, the interplay of all these aspects. It turns out that materials, for example, can take over some of the processes normally attributed to control, a phenomenon that is called "morphological computation". There is no taxonomy of morphological computation yet, but we can roughly distinguish between sensor morphology taking over a certain amount of computation, similarly for shape and materials, and for the interaction with the environment.

In an embodied agent, by the mere fact of its being physical, all aspects – 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 system 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. We start with sensor morphology, followed by two locomotion examples, and we conclude with the example of grasping with an artificial hand. Finally, we will discuss what has been achieved, what lessons there might be for neuroscience and robotics research.

2. Sensor morphology

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 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 (e.g. obstacle avoidance) 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, a phenomenon that is studied in the field of space-variant vision, e.g. Kuniyoshi et al., 2000; Ferrari et al., 1995). 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). This non-homogeneous arrangement of the light-sensitive cells, the ommatidia in the case of the insects, 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 (Lichtensteiger, 2004).

Note that all this only works, if the agent is actually behaving in the real world and therefore 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. Locomotion

In this section two case studies, the quadruped "Puppy", and the artificial fish "Wanda" demonstrating the exploitation of materials and in particular dynamics of the system-environment interaction, will be introduced.

Muscles: control from materials – the running quadruped "Puppy": We now present a case study where 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 as the response times are too slow. 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 (Iida and Pfeifer, 2004).

The design of "Puppy" was inspired by biomechanical studies. Each leg has two standard servomotors and one springy passive joint. To demonstrate a running gait, we applied a synchronized oscillation based control to four motors in the "hip" and "shoulder", 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.

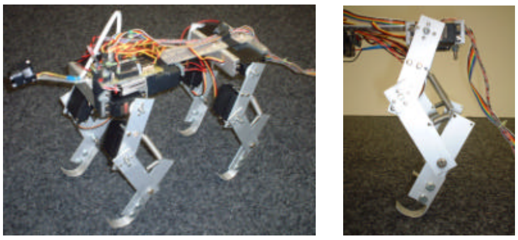

a **b**

Figure 2: The quadruped "Puppy". (a) Picture of the entire "Puppy". (b) The spring system in the hind legs.

Even tough 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 quadruped gait emerges, 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 controller of the robot is extremely simple: it 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. The morphological computation in this case 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, shape of the ground, gravity). Exploiting morphological computation makes cheap rapid locomotion possible because physical processes are fast and for free! (for further references on cheap locomotion, see e.g. Kubo and Full, 1999; Blickhan et al., 2003; Buehler, 2002). Now, if sensors – e.g. pressure sensors on the feet, angle sensors in the joints, and vision sensors on the head – are put on the robot, structured – correlated – sensor stimulation will be induced that can potentially be exploited.

Behavioral diversity from system-environment interaction – "Wanda": The artificial fish, "Wanda", built by Marc Ziegler and Fumiya Iida (Ziegler and Iida, in preparation), a very recent development in our laboratory, exploits the interaction with the environment in interesting ways. The fish has one single degree-of-freedom of actuation: 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. The speed is controlled by the wiggling frequency and amplitude. 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 through morphological computation: instead of having more complicated actuation, e.g. additional fins or a flexible spine and thus more complex control, the interaction with the environment can be exploited to achieve the task.

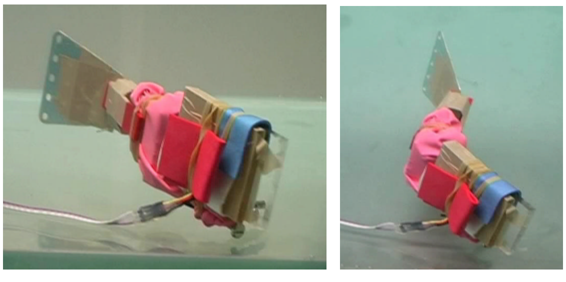

Figure 3: The artificial fish “Wanda” with one degree-of-freedom for wiggling the passive tail fin.

4. Grasping

“Cheap” grasping – the “Yokoi hand”: Because of the limited space we can only give a very simplified account. The 24 degrees-of-freedom “Yokoi hand” (Yokoi et al., 2004) that can be used as a robotic and a prosthetic hand, is partly built from elastic, flexible, and deformable materials. 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 morphological computation performed 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. Because of this morphological computation, control of grasping is very simple, or in other words, very little brain power is required for grasping. For prosthetics, there is an interesting implication. If EMG signals, which are known to be very noisy, are used, control cannot be very precise and sophisticated. But by exploiting morphological computation, there is no need for very precise control, at least for grasping. Of course, by adding touch, angle, and torque sensors, the abilities of the hand can be improved and feedback signals can be provided to the agent (the robot and the human) that can be exploited by the neural system for building object concepts. However, that is not the main point here.

a **b**

Figure 4: Cheap grasping: exploiting system-environment interaction. The Yokoi hand exploits deformable and flexible materials to achieve self-adaptation.

5. Conclusions: lessons for neuroscience and robotics

We introduced the concept of morphological computation which can be characterized as performing a kind of “task distribution” between the brain (neural system), or the controller in a robot, the morphology of the agent (shape, sensors, actuators, materials), and the environment. We showed that by exploiting morphology, materials, and system-environment interaction, hard tasks such as rapid locomotion or grasping can be achieved in a “cheap” manner. Let us speculate a bit what potential lessons there might be for neuroscience and for robotics. While some of these points are entirely obvious, it is interesting to note that in practical everyday research, they are normally not considered, or not considered sufficiently.

First, by looking at the neural system only the function of the neural system cannot be understood: we must take the way it is embedded into the agent and the specific types of interactions with the environment into account as well (e.g. the case study on motion parallax). Second, 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). Third, the interaction with the environment takes over essential aspects of the control task, which simplifies not only the control, but also the morphology of the agent (e.g. the artificial fish “Wanda”). If we are interested in brain function, i.e. the role the brain plays in subtending behavior, the entire agent and the interactions with the environment must be taken into account. Recent insights in biomechanics, for example, 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 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 morphological computation provides new ways of looking at behavior generation, because in the past the focus has been very much on the control side.

One problem with the concept of morphological computation is that while intuitively plausible, it has defied serious quantification efforts: We would like to be able to ask "How much computation is actually being done?" A first very crude approximation would be to compare a system that exploits morphological computation with one that doesn't, for example, how much computation is required for the controlled forward movement of a leg in biped walking compared to a passive swing. Another problem is that the notion of computation in the context of morphology or dynamics may in fact require fundamental reconceptualization, which is a challenging research topic.

In summary, we hope to have demonstrated the power of morphological computation as a new way of designing robots and of understanding biological systems, thus giving the term "embodiment" a non-trivial meaning that goes substantially beyond "intelligence requires a body."

Acknowledgments

The author would like to thank the members of the Artificial Intelligence Laboratory in Zurich and Hiroshi Yokoi, who is a frequent visitor at our laboratory, for many discussions and for building exciting robots, and to Yoshi Nakamura for inviting me to write this article.

References

Blickhan, R., Wagner, H., and Seyfarth, A. (2003). Brain or muscles?, Rec. Res. Devel. Biomechanics, 1, 215-245.

Buehler, M. (2002). Dynamic locomotion with one, four and six-legged robots, J. of the Rob. Soc. of Japan, 20(3): 15-20.

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.

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.

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. (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

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 self-stabilization in hexapedal runners, Phil. Trans. R. Soc. Lond. B, 354, 849-861.

Lichtensteiger, L. (2004). On the interdependence of morphology and control for intelligent behavior. PhD Dissertation, University of Zurich.

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. (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 Scheier, C. (1999). Understanding intelligence. Cambridge, Mass.: MIT Press.

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.

Ziegler, M., and Iida, F. (in preparation). Cheap underwater locomotion: Influence of material properties and behavioral diversity.