Connection Science, 2004

Internal robotics

Domenico Parisi

Institute of Cognitive Sciences and Technologies

National Research Council, Rome

phone: 39-06-86090376

fax: 39-06-824737

d.parisi@istc.cnr.it

Abstract

Robotics can contribute significantly to our understanding of the behaviour of organisms because of its emphasis on the role of the body and its physical interactions with the external environment in determining the organism's behaviour. However, behaviour is the result of the interactions of an organism's nervous system with both the external environment and the internal environment, i.e. with what lies within the organism's body. While robotics has concentrated so far on the first type of interactions (external robotics), to more adequately understand the behaviour of organisms we also need to reproduce in robots the inside of the body of organisms and to study the interactions of the robot's control system with what is inside the body (internal robotics). The paper discusses seven differences between the two types of interactions and briefly describes some examples of internal robotics: robots which evolve a biological clock that allows them to look for food during the day and to rest at night, robots which evolve a pain signal associated with some damage in their body which allows the robot to stop moving when ill in order to recover more quickly from the bodily damage, and robots which can be hungry and/or thirsty and respond adaptively or maladaptively to conflicts between these two motivations.

1. Introduction

Robotics can contribute in very significant ways to our understanding of the behaviour of organisms because it makes entirely clear the role of an organism's body in determining the organism's behaviour. Organisms are, before anything else, physical entities which interact with a physical environment. The behaviour of organisms depends on the external morphology of their body (size, shape, etc.), on the spatial arrangement of the body's sensors, on how external physical/chemical events are encoded in the sensors, and on the body's effectors and their degrees of freedom. Organisms with different bodies tend to behave differently. The knowledge of the environment which allows an organism to interact with the environment in ways that tend to increase its survival and reproductive chances is derived from the physical interactions of the organism's body with the environment. By moving its effectors the organism modifies either the physical relation of its sensors to the external environment or the external environment itself, in both cases actively influencing the sensory input from the environment. The external environment is known and understood by the organism in that the environment reacts in more or less predictable ways to the organism's movements.

Robots are physical artefacts that interact with a physical environment using their sensors and effectors. Hence, they are especially appropriate as tools and models for understanding how the physical interactions of an organism with its environment shape the organism's behaviour. Evolutionary robotics is the design of robots and of the control systems that determine their behaviour through an evolutionary process in which a population of variants reproduce selectively and with the constant addition of new variants (Cliff et al. 1993, Pfeifer and Scheier 1999, Nolfi and Floreano 2000). The behaviour of the robots is not decided or programmed by the researcher but it spontaneously emerges and self-assembles in a succession of generations of individual robots that interact with a physical environment. The methods of evolutionary robotics have been applied with success to both the design of robots' bodies (Sims 1994, Lipson and Pollack 2000; Hornby and Pollack 2002) and to the design of the control systems that guide their behaviour (Nolfi and Floreano 2000). In the design of control systems both evolutionary methods at the population level and learning methods at the level of the individual have been applied, including the study of the interactions between evolution and learning (Parisi et al. 1990, Nolfi et al. 1994).

However, robotics so far has mostly been an external robotics. When one uses robots to study the behaviour of organisms, one tends to concentrate on the interactions of the robot's control system, say, a neural network, with the external environment. Given the physical structure of the external environment and given the sensory organs of the robot, the input units of the robot's neural network encode a particular pattern of activation caused by some event in the external environment, the neural network elaborates this input in its internal structure, and responds by producing some movement of the robot's effectors. These movements change either the physical relation of the robot's body to the external environment or the external environment itself, thereby influencing the next sensory input from the external environment.

But the external environment is only one of the two sources of inputs for the nervous system of an organism. An organism's nervous system interacts not only with the environment which lies outside the organism's body, i.e. the external environment, but also with the organs and systems which make up the rest of the organism's body beyond the nervous system, i.e. the internal environment. The behaviour of an organism results from both the interactions of the organism's nervous system with the external environment and its interactions with the internal environment. Therefore, what is needed is not only an external robotics but also an internal robotics. If we want to understand the behaviour of organisms what we need to reproduce in artificial physical organisms, i.e. robots, is not only the external morphology of an organism's body and the interactions of the organism's nervous system with the external environment but also the internal physical structure of the organism’s body and the interactions of the organism’s nervous system with what lies inside the body.

2. The two worlds in which organisms live

Organisms live in two worlds: an external world and an internal world. Robotics can help understand both the external world and the internal world of organisms.

A nervous system is a physical system that maps inputs into outputs in ways that allow an organism to survive and reproduce in its environment. Inputs are caused by physical or chemical events outside the nervous system and outputs cause physical or chemical events outside the nervous system. But “outside the nervous system” need not necessarily mean outside the organism’s body. Nervous systems are physical systems which are contained in a larger physical system, the organism’s body, which in turn is contained in a still larger physical system, the organism’s physical environment. Where do the inputs to an organism’s nervous system come from? Where do the outputs of the nervous system cause their effects?

For all animals, including humans, there are two answers to both these questions. The inputs that arrive to the nervous system from outside the nervous system can come either from the physical environment surrounding the organism’s body or from within the organism’s body. A light can be reflected by an object which is present in the external environment and cause a pattern of activation in the nervous system’s visual receptors. Some mechanical energy can be generated in the contact between the organism’s body and an object in the external environment and cause a pattern of activation in the nervous system’s touch receptors. Some chemical molecules can come out of a bottle of perfume and cause a pattern of activation in the nervous system’s smell receptors. These are all cases in which the input to the nervous system comes from outside the organism’s body, from the external environment. But inputs to the nervous system can also come from within the organism's body. Movements and postures of the body cause somatosensory and proprioceptive inputs to the nervous system. Specific structures of the endocrine system produce molecules that travel in the blood circulatory system and reach the nervous system causing specific effects on the nervous system. Even the nervous system itself can produce such molecules that modulate the nervous system's own activity.

The same applies to the nervous system's outputs. All the outputs of the nervous system cause changes outside the nervous system. However, these changes can occur in the physical relation of the organism's body with the external environment and in the external environment itself, but they can also occur inside the organism's body. The organism's nervous system can respond to inputs with movements of the eyes, head, arms, hands, legs, phono-articulatory organs, but it can also respond by causing changes inside the body: movements of the stomach's muscles, contractions of the body vessels, the production of specific molecules that affect the body, the stimulation of the endocrine system for the production of such molecules.

Therefore, while all behaviour is the result of the interactions of the organism's nervous system with the environment outside the nervous system, it is important to distinguish between two components of the environment: (1) the physical environment which lies outside the organism's body (external environment), and (2) the rest of the organism's body outside of the nervous system (internal environment) (figure 1). The interactions of the nervous system with the two environments create the two worlds in which organisms live.

[Insert figure 1 about here]

The first kind of interactions take place, in one direction, when physical or chemical events outside the organism's body influence the activity of some special neurons, called sensory receptors, which are located on the surface of the organism's body and, in the opposite direction, when the nervous system influences the environment outside the organism's body because its motor neurons determine movements of the entire body or of some body part. In other circumstances the nervous system causes movements of the phono-articulatory organs which in turn cause acoustic energy to be present in the environment and this acoustic energy is perceived as sound by both the individual itself and other nearby individuals.

The second type of interactions, those between the nervous system and the rest of the organism's body, occur, in one direction, when the movements of the body affect the nervous system's somatosensory and proprioceptive receptors or when chemical substances that are produced inside the body influence the functioning of the nervous system and, in the opposite direction, when the nervous system moves various organs inside the body, e.g. heart or stomach, or when chemical substances produced by the nervous system influence various structures inside the body, including the nervous system itself.

In both cases the interactions of the nervous system with what lies outside itself tend to be circular in that the effects produced by the nervous system outside itself, either in the external environment or inside the rest of the body, become causes that affect the nervous system (cf. Piaget's circular reactions and von Foerster's "double closure"; Piaget 1952, Ziemke in press). Therefore, in both cases the nervous system is largely responsible for the causal influences that from outside affect its functioning. Furthermore, the two cause-effect circuits depicted in Figure 1 are not separated but they are interconnected. The nervous system may respond to causal influences from the external environment by producing effects inside the rest of the body and to causal influences from inside of the body with movements that affect the external environment.

While both types of interactions of the nervous system with the outside world are important to understand and explain behaviour, almost all research in robotics and, more generally, in Artificial Life has concentrated on the first type of interactions, those between the nervous system and the external environment. Little research has been dedicated to an internal robotics, that is, to the study of the interactions between the nervous system and the inside of the organism's body and of how these interactions can contribute to the overall behaviour of the organism.

In this paper we outline some of the properties that characterize the interactions between the nervous system and what lies within the body and contrast them with the properties of the interactions between the nervous system and the external environment, and we briefly describe some simple simulations that try to capture the first type of interactions and to contribute to the development of an internal robotics.

3. Seven differences between the two worlds

In this Section we describe seven differences that distinguish the interactions of the nervous system with what lies inside the organism's body from its interactions with the external environment.

(1) The nervous system's interactions with the external environment are predominantly physical whereas those with the internal environment are predominantly chemical

The effects of the nervous system on the external environment are mostly mediated by movements of the organism's body or body parts. The output units of a neural network encode movements which are physically realized by moving eyes, heads, arms, hands, legs, phono-articulatory organs, etc. Therefore, these effects can be studied by remaining within the boundaries of the discipline of physics. More rarely the nervous system has effects on the external environment which are mediated by the production of chemical molecules which diffuse in the external environment, such as the production of odours. Also, the inputs that arrive to the nervous system from the external environment are, in humans, predominantly caused by physical energies (light or sound) whereas more rarely the external environment has effects on the nervous system through chemical molecules (smell or taste).

On the contrary, the interactions between the nervous system and what is inside the body are mostly chemical and need to be studied with the help of chemistry. Many effects of the nervous system on the rest of the body take place through the direct production of chemical molecules (neurohormones) that affect various structures inside the body or through stimulation of portions of the endocrine system which produce chemical molecules (hormones). In its turn, the rest of the body influences the nervous system by producing chemical molecules and sending these molecules to the nervous system, although in some cases the body sends to the nervous system inputs which are similar to the inputs from the external environment such as the somatosensory and proprioceptive inputs caused by postures and movements of the body.

This difference between the mostly physical interactions of the nervous system with the external environment and the mostly chemical interactions of the nervous system with the rest of the body has an important implication for research on neural networks and robotics. By addressing the interactions of the nervous system with the external environment, research on neural networks tends to consider <u>cellular</u> biology as the biological discipline of reference (after all, a neural network's units are cells and connections are extensions of cells) whereas studying the interactions of the nervous system with the rest of the body necessarily requires one to go one step down in the hierarchy of biological entities and to look at <u>molecular</u> biology. Therefore, internal robotics would bring research on neural networks more in line with recent developments in biology and medicine which also tend to go one step down and to analyse and explain biological phenomena at the level of molecular processes. For robotics, the implication is that physics (mostly mechanics) may be not enough to construct robots if robots must have an "inside" and not only an "outside" and the control system of robots must not only interact with the external environment but also with the "inside" of the robot itself.

(2) Two kinds of influences on the nervous system

The interactions of the nervous system with the external environment are mediated by the specific pattern of neuron-to-neuron connections and by the specific weights of the individual connections. The external environment influences the neural network in that it causes a specific pattern of activation in a specific subset of the network's units, the input units. This input activation pattern is transformed into other specific activation patterns in the successive layers of internal units until a specific activation pattern appears in another specific subset of the network's units, the output units. This activation pattern causes a specific movement of some part of the organism's body. As a consequence, the processes which take place inside the network as a result of the inputs from the external environment and the effects of these processes on the external environment have a character which can be defined as qualitative and specific, rather than quantitative and diffuse.

On the contrary, the interactions of the nervous system with the rest of the body tend to have a more aspecific and diffuse character, more quantitative than qualitative. The action of what lies inside the body on the nervous system is mediated by the arrival of chemical molecules which diffuse in large zones of the intercellular space, affect in a more or less similar way large number of neurons rather than targeting single neurons, and have largely quantitative effects such as increasing or decreasing the activation threshold of neurons. Neurons influence each other because they are connected by axons, not because they are close to each other in physical space. The manner in which the nervous system responds to the inputs from the external environment can be called neuro-transmissory whereas the manner in which it responds to influences from the rest of the body can be called neuro-modulatory. Neurotransmission is based on the particular topology or architecture of connections of the nervous system, which is at least in part independent of space. Neuromodulation is more dependent on space. (Some recent work has addressed neuromodulation in neural networks that control the behavior of robots. See, for example, the work of the GasNet Group (Husbands et al. 1998) and research on the role of neuromodulators such as dopamine, serotonin, noradrenaline, and acetylcholine in regulating reinforcement learning (Doya 2002)).

(3) The circuit "nervous system-rest of the body" is entirely evolved whereas the circuit "nervous system-external environment" is evolved only for the part of the circuit which is in the nervous system

The circuit "nervous system-rest of the body" (circuit (a) in figure 1) goes through a physical system, the rest of the body, which has evolved together with the nervous system and, therefore, it has characteristics which have been shaped by the interactions with the nervous system (Purves 1990). On the contrary, the circuit "nervous system-external environment" (circuit (b) in figure 1) goes through a physical system, the external environment, which has not co-evolved with the nervous system and which therefore has characteristics which have not been shaped by the interactions with the nervous system. Aside from the manner in which the external environment is encoded (if it is encoded) in the input units of the neural network of particular species of organisms, the external environment is the same for all species of organisms living in the same environment, whereas the inside of the body is different in different species.

A current direction of research in external robotics is the study of how the external morphology of the body (size, shape, etc.) evolves together with the control system of the robot (Sims 1994, Lipson and Pollack 2000; Hornby and Pollack 2002). Internal robotics requires the study of how not only the external morphology of the body but also the inside of the body, and the manner in which the inside of the body influences the control system, co-evolve together with the control system.

In human beings not only their internal environment but also their external environment co-evolves with their behaviour in that the external environment is shaped and created by their behaviour through technology. Although some non-human animals modify the external environment in adaptive ways (Kirsh 1996) and of course there is a lot of both intra- and interspecific co-evolution in the animal world, human beings are unique in that they almost entirely create their own external environment. Therefore, one can imagine that a future extension of external robotics will be the study of how robots can evolve a strategy of modifying the external environment to make the external environment more adapted to themselves rather than modifying themselves to become more adapted to a fixed external environment. Already existing examples of this type of technological robotics are simulations of insects' stigmergic communication (Bonabeau 1999, Camazine et al. 2001), creating obstacles to defend oneself from predators (Buason and Ziemke 1993), and enriching the environment with artificial signals that can make wayfinding easier (Ziemke et al. in press).

(4) For the nervous system the rest of the body is something which is always present and always more or less the same whereas the external environment can be present or absent and it can be very different at different times

The rest of the body is a constant source of causal influences for the nervous system, creating a world which is always present and which remains basically the same at all times. On the contrary, the external environment creates a world which can be present but also absent (e.g. when one is sleeping) and which can be wildly different at different times. As suggested by Damasio (1999, cf.

also Ledoux et al. 2003), the stability and constant availability of the rest of the body to the nervous system can be the basis for the emergence in the nervous system of a “self” or a “sense of self” which is opposed to the external environment as “nonself” or “rest of the world”. Hence, an internal robotics might be a necessity if one wants to construct robots that have a “self” or a “sense of self”.

(5) The causal influences originating from within the body can result in the emergence of a private world whereas those originating in the external environment define a public world

For purely physical reasons the events and processes that take place inside an organism’s body can produce effects only in the nervous system of a single individual, the individual which owns that body, whereas the events and processes that take place in the external environment can produce the same effects, or similar effects, in the sensory organs of many individuals, provided that these individuals are all sufficiently spatially close to the physical origin of these effects. This is the basis for the emergence of a distinction between a private and a public world. However, although for all kinds of organisms, even very simple ones, events inside the body produce effects only in one individual and events outside the body produce effects in many individuals at the same time, the distinction of a private and a public world may not be an automatic consequence of this fact. The distinction between a private and a public world is likely to emerge only for organisms which have a sophisticated sociality and a complex nervous system, such as humans. For example, only a human being may be able to predict that a certain input (private, that is, originating within his or her body) will produce effects on his or her own behaviour but not on the behaviour of other nearby individuals whereas another kind of input (public, that is, originating in the external environment) will produce effects both on himself/herself and on other nearby individuals. Therefore, it is possible that a clarification of the distinction between “private” and “public” will only emerge from the joint contribution of internal robotics and collective, or social, robotics.

(6) The cognitive components of behaviour emerge from the interactions of the nervous system with the external environment whereas its emotional or affective components emerge from the interactions of the nervous system with the rest of the body

The cognitive components of behaviour can be explained, mostly, by considering only the interactions of the nervous system with the external environment. It is in these interactions that an individual exhibits its abilities, its problem solving capacities, its intelligence. On the contrary, the emotional or affective components of the behaviour of an organism are the result of the interactions of the organism's nervous system with what is inside the organism's body and one cannot even begin to explain these components if one ignores these interactions and the inside of the organism's body. Only the more superficial aspects of emotional or affective phenomena, such as the production of expressive movements in robots or in icons on a computer screen or the "recognition" of the emotions of the user, can be reproduced in artefacts if one ignores what is inside an organism's body and the interactions of the nervous system with what is inside the body. To construct robots that actually "feel" and not only "pretend to feel" one has to develop an internal robotics.

Not only behaviour has both cognitive and affective components but the two components cannot be separated and neither of the two components can really be understood without considering the other. For example, an organism can be able to do many different things (cognitive component) but what the organism actually does at any particular time depends on its motivational state (affective component). The current motivational state of an organism is a result of the interactions of the organism's nervous system with the inside of the body and it controls aspects of behaviour which seems to be purely cognitive such as selective attention.

(7) The interactions of the organism's nervous system with the external environment tend to produce effects that the organism is able to predict and are voluntary whereas the interactions of the nervous system with what is inside the body give rise to effects which the organism is unable to predict and are involuntary

The interactions of the nervous system with the external environment tend to produce voluntary behaviours and voluntary effects in the external environment. On the contrary, the effects of the nervous system on the rest of the body tend to be involuntary. While it may be difficult to clearly distinguish between voluntary and involuntary behaviours and effects, the voluntary character of behaviour appears to be linked to the predictability of its effects on the part of the individual which exhibits the behaviour, and the ability to predict can be simulated using neural networks that control the behaviour of robots (Parisi et al. 1990, Nolfi and Tani 1999). Although the effects of one's behaviour on the external environment can be within limits predicted, this is rarely the case for the effects that one's nervous system has on what is inside one's body (although of course there are some predictable regularities such as ceasing to be hungry when one eats).

Human beings adopt strategies that can make the effects of one's nervous system on the rest of the body and the effects of the rest of one's body on one's nervous system more predictable and controllable. One strategy is to acquire a greater knowledge/control of the internal states of one's body through various forms of body control techniques. Another strategy is exposing oneself to inputs from the external environment which one knows can have effects on one's nervous system which in turn have effects on the rest of one's body. Exposure to such inputs can take the form of talking with others - relatives, friends or psychotherapists - about emotional and affective topics and the creation of artistic artefacts or exposure to the artistic artefacts made by others.

4. Some examples of internal robotics

In this Section we briefly describe some very simple examples of internal robotics.

(A) Robots that sleep

Imagine a robot which has to move in the environment looking for food. The robot’s behaviour is controlled by a neural network with input units encoding the position of the nearest food element and output units encoding the robot’s displacements in the environment (figure 2a). Moving has a cost in terms of energy which is compensated by the energy extracted from the food that the robot finds in the environment. Imagine that the environment has a day/night cycle and that at night food is very difficult to see so that the cost of moving is not compensated by the food found. Therefore, for the robot it would be appropriate to look for food during daytime and to rest at night.

[Insert figure 2 about here]

The problem can be solved by adding to the input units that allow the robot to perceive the location of food some other input units that encode the quantity of light present in the environment (light sensor) (figure 2b). Using a genetic algorithm one can then evolve a population of robots which respond to the input that specifies the location of food with movements that allow the robots to approach and eat the food when there is enough light in the environment (day) and to respond to the same information by stopping moving when there is little or no light (night). When the light sensor senses no light, the connection weights linking the light sensor to the motor output units inhibit the robot’s movements. When light returns, this inhibition ceases and the robot moves in response to information telling it where is the food.

Now imagine that after they have lived for some generations in an environment which contains only food the robots move to another environment which in addition to food elements contains some caves where light does not penetrate even during daytime. A robot which is used to stop moving when it is dark will respond by stopping moving when it enters one of these caves - which means that it will also stop eating for the rest of its life and will not reproduce. How is this problem to be solved?

One can solve this problem by replacing or complementing the light sensor with a biological clock (figure 2c). The biological clock consists of a set of special input units of the neural network controlling the robot's behaviour which encode inputs originating not in the external environment but inside the robot's body. The robot's body contains a structure which activates the biological clock input units in accordance with the day/night cycle. A genetic algorithm can be used for evolving this structure which controls both the cyclic activation pattern of the biological clock input units and its entrainment with the actual day/night cycle. Notice that the biological clock requires both the evolution - in fact, the co-evolution - of the structure that determines the appropriate cyclic activation pattern in the biological clock input units (this structure is part of the "rest of the body") and the evolution of the appropriate connection weights that link the biological clock input units to the motor output units (which are part of the nervous system).

Using the biological clock the evolved robots move in the environment looking for food during daytime and they rest at night but, what is crucial, they do not stop moving if they happen to enter a cave during the day even if there is no light in the cave. The biological clock is a more complex mechanism than the light sensor and it is somewhat slower and more difficult to evolve than the light sensor. If we allow the robots to evolve both the light sensor and the biological clock, the worst results are obtained if the light sensor and the biological clock input units independently influence the neural network's motor output units (figure 2d), intermediate results are obtained if the light sensor both directly controls the motor output units and sends its connections to the biological clock units (figure 2e), while we find the best results if the light sensor is not directly connected with the motor output units but it only regulates the biological clock (figure 2f) (Mirolli and Parisi 2003). This last architecture seems to be in agreement with what happens in the real brain where level of ambient light is directly communicated by the eyes to neurons in the suprachiasmatic nucleus that control the sleep/wake cycle.

(B) Robots that feel pain

Imagine that robots similar to those of the preceding simulation have to look for food to survive and reproduce but they now have a different kind of problem. There is no day/night cycle in their environment so they do not have to learn to move during the day and to rest at night like the preceding robots. The problem that the new robots have to solve is that for some accidental reason their body can incur some physical damage or illness which tends to recover spontaneously but which has a cost in terms of reproductive chances as long as it lasts. However, this cost is greater if the robot continues to move in the environment looking for food when its body happens to be damaged while the cost is smaller if the robot stops moving until it recovers from its damage/illness. Hence, for these robots it would be appropriate to stop moving when they happen to be ill and to move in the environment looking for food only when they are healthy.

This is where (physical) pain seems to play a role. Imagine that the neural network that controls the robot's behaviour has an additional set of input units (pain units) which are connected with the network's motor output units. If these units encode some particular activation pattern when the robot's body is damaged and a different activation pattern when the body is healthy, these two different activation patterns can regulate the robot's movements appropriately. When the first activation pattern appears in the pain units, signaling pain, the units should inhibit the robot's movements. When the second activation pattern appears in the pain units, signaling no pain, the inhibition should disappear and the robot should move around looking for food.

This more complex behaviour can be evolved using a genetic algorithm. In one simulation we evolved only the connection weights from the pain units to the motor output units while hardwiring the two distinct activation patterns in the pain units in correspondence with the two possible states of the body: damaged vs healthy. In a second simulation we evolved together both components of the pain mechanism. The tendency of the body to activate one activation pattern in the pain units when it was damaged and a different activation pattern when it was healthy evolved together with the appropriate connection weights of the neural network which caused the robot to stop moving when the first activation pattern appeared in the pain units and to resume moving when the second activation pattern appeared in the pain units. In other words, both the ability of the body to send pain signals to the nervous system and how the nervous system has to respond to these signals evolved together (Acerbi and Parisi in preparation).

This simulation, like the preceding one, illustrates the co-evolution of the control system (neural network) and of the internal environment with which the control system interacts which is typical of internal robotics. An extension of this work is the evolution of facial expression of pain. Pain signals from the body can cause an individual’s nervous system to produce facial expressions which function as inputs for other individuals, inducing these individuals to help the individual with the damaged or ill body (Williams 2002). The simulation may also have practical applications. Ceasing to move is a minimal response to a pain signal from the body. More active and specific responses to pain signals are possible such as touching and protecting the body part that hurts (which requires spatially specific pain signals) or seeing a doctor. Robots with practical uses may be informed by their body about possible internal malfunctionings and may be able to take action to eliminate these malfunctionings or minimize their consequences.

(C) Robots that are hungry and thirsty

In this last simulation the robots live in an environment that contains both food elements and water elements. To survive and reproduce the robots must be able to eat as much a possible and to drink as much as possible (within limits) but they also have to maintain a balance between eating and drinking. Individuals that eat much more than they drink or drink much more than they eat, are not very likely to reproduce. The robots consume both energy (from food) and liquids (from water), so they need both food and water and cannot live either only on food or only on water.

The neural network that controls the robots' behaviour has two sets of input units that separately encode the position of the nearest food element and the position of the nearest water element and two sets of "motivational" units that respectively encode the quantity of energy and of liquids that are currently present in the robot's body. The first two sets of units encode information from the external environment while the two sets of motivational units encode information from within the body. Since the robot perceives both food and water at the same time, which are not normally located in the same place in the environment, it has to decide at any given time whether to approach food or to approach water. The information concerning its current motivational state - how much energy from food (hunger) and how much liquids from water (thirst) are currently present in its body - is used to adopt the appropriate decision.

In the simulation we do not simulate the evolution of the motivational state, that is, the manner in which the body "learns" to communicate to the nervous system how much energy and how much liquid it currently contains. We hardwire in the two sets of motivational units the activation patterns which reflect the current quantity of energy and of liquids which are present in the robot's body. What we are interested in is how the robots respond to this information. The robots use the motivational state to guide their behaviour. When they are hungry (little energy in their body) but not thirsty (sufficient liquids) they ignore water (even if they perceive water) and go toward to the food. When they are thirsty, they behave in the opposite way. Hence, their motivational state functions as a selective attention mechanism.

This is how the robots behave when there is a clear difference in the levels of their energy and of their liquids, i.e. when they are hungry but not thirsty or when they are thirsty and but not hungry. What is interesting is how the robots behave when there is no clear difference between the levels of energy and of liquids inside their body. This may create a situation of motivational conflict.

The results of the simulation show that when they are neither very hungry nor very thirsty (they have eaten and drunk enough recently), the robots will simply stop moving. Consider that in this simulation, as in the simulation with the day/night cycle, moving has a cost. Hence, the robots develop an adaptive response to a motivational state of neither hunger nor thirst: they avoid paying the costs of moving by stopping moving until they have consumed their energy and their liquids and they become hungry and/or thirsty again.

But the more interesting result concerns the robots' behaviour when they happen to be both very hungry and very thirsty. Since food and water are randomly distributed, this may happen for purely chance reasons, or an individual robot may not be very good at reaching food and water, so that it can become both very hungry and very thirsty.

In response to a motivational state of both strong hunger and strong thirst, the robots' behaviour appears to depend on the parameters of the simulation. If the costs in terms of reproductive chances of going below a given threshold of energy and liquids are not very great, a robot which is both very hungry and very thirsty will tend to go either toward food or toward water according to the preferences of each individual. This of course is an adaptive response. However, if the costs are very large, robots that are both very hungry and very thirsty simply tend to stop moving. They become paralyzed. This is a maladaptive response since it will completely destroy the individual's chances of surviving and reproducing. This behaviour may resemble some pathological behaviour in human beings which is caused by the existing conflict between two very strong motivations (Cecconi and Parisi 1993, Parisi 1996, Pelliccione and Parisi in preparation).

5. Conclusions

We need an internal robotics in addition to the current external robotics. External robotics is concerned with the external morphology of the body and with the interactions of the robot's control system (neural network) with the external environment. Internal robotics is the study of what lies inside a robot's body and of the interactions of the robot's control system with what lies inside the body. The behaviour of organisms results from both kinds of interactions, so we need to study both. In fact, the mind lives in two worlds, the external cognitive, public, sufficiently predictable and controllable world which results from the first type of interactions, and the more affective, private, much less predictable and controllable world which results from the second type of interactions. (In fact, the (human) mind lives also in a third world which is created by the rich structure of recurrent connections within the organism's nervous system. These recurrent connections self-generate inputs for the nervous system and underlie what is called mental life: visual images, rememberings, thoughts, reasonings, explicit predictions and plans.)

We have briefly described three simple simulations that illustrate some aspects of internal robotics: the role of the body in controlling the overall level of activity of the nervous system, how the body can inform the nervous system about bodily states that may require specific action on the part of the nervous system, the co-evolution of the body's ability to send the appropriate inputs to the nervous system and the nervous system's ability to respond appropriately to these inputs, the role of some bodily inputs (motivations) to guide the nervous system's interactions with the external environment (selective attention). Of course, these simulations merely scratch the surface of a huge world of interesting and diverse phenomena: different states and levels of activation and vigilance, affective phenomena, the complex interplay of different motivations, the role of motivations in directing behaviour toward particular goals and attention toward particular objects and aspects of the environment, the distinction between a public and private world, the emotional side of consciousness. These phenomena cannot be considered as separate from the more cognitive aspects of a robot's behaviour. They are not phenomena that can be ignored even if we are only interested in the abilities of robots, in what they can do for us. One reason why human behaviour is more complex and more sophisticated than a typical robot's behaviour is because it results from the interplay between the interactions of the nervous system with the external environment, on one side, and its interactions with the internal environment, on the other side. The simulations that we have described are simple demonstrations that these phenomena can and should be tackled by a more ambitious robotics of the future.

References

Acerbi, A. and Parisi, D., in preparation, *Pain in artificial organisms*.

Bonabeau, E. (ed.), 1999, Special issue on Stigmergy. **Artificial Life**, 5(2).

Buason, G. and Ziemke, T., 2003, Competitive co-evolution of predator and prey sensory-motor systems. In *RealWorld Applications of Evolutionary Computing - EvoWorkshops 2003*, edited by S. Cagnoni, J. Cardalda, D.W. Corne, J. Gottlieb, A. Guillot, E. Hart, E. Marchiori, J.A. Meyer, M. Middendorf, G.R. Raidl (Berlin: Springer), pp. 605-611.

Camazine, S., Deneubourg, J.L., Frank, N.R., Sneyd, J., Theraulaz, G., and Bonabeau, E., 2001, *Self-Organization in Biological Systems* (Princeton: Princeton University Press).

Cecconi, F., and Parisi, D., 1993, Neural networks with motivational units. In *From Animals to Animats 2*, edited by J-.A. Meyer, H.L. Roitblat, and S.W. Wilson (Cambridge, Mass.: MIT Press), pp. 346-355.

Cliff, D., Harvey, I., and Husbands, P., 1993, Explorations in evolutionary robotics. **Adaptive Behavior**, **2**, 73-110.

Damasio, A., 1999, *The Feeling of What Happens. Body and Emotion in the Making of Consciousness* (London: Harcourt).

Doya, K., 2002, Metalearning and neuromodulation. **Neural Networks**, 15, 495-506

Hornby, G.S., and Pollack, J.B., 2002, Creating high-level components with a generative representation for body-brain evolution. **Artificial Life, 8**, 223-246.

Husbands, P., Smith, T., Jakobi, N., and O'Shea, M., 1998, Better living through chemistry: evolving GasNets for robot control. **Connection Science**, 10, 185-210.

Kirsh, D., 1996, Adapting the environment instead of oneself. **Adaptive Behaviour**, 4, 415-452.

Ledoux, J.E., Debiek, J., and Moss, H. (eds.), 2003, *The Self: From Soul to Brain* (New York: New York Academy of Science).

Lipson, H., Pollack, J.B., 2000, Automatic design and manufacture of robotic lifeforms. **Nature, 406**, 974-978.

Mirolli, M. and Parisi, D., 2003, Artificial organisms that sleep. In *Advances in Artificial Life*, edited by W. Banzhaf, T. Christaller, P. Dittrich, J.T. Kim, and J. Ziegler (New York: Springer), pp. 377-386.

Nolfi, S., Elman, J.L., and Parisi, D., 1994, Learning and evolution in neural networks. **Adaptive Behavior, 3**, 5-28.

Nolfi, S. and Floreano, D., 2000, *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-Organizing Machines* (Cambridge, Mass.: MIT Press).

Nolfi, S., and Tani, J., 1999, Extracting regularities in space and time through a cascade of prediction networks. **Connection Science, 11**, 125-148.

Parisi, D., 1996, Motivation in artificial organisms. In *Machine Learning and Perception*, edited by G. Tascini, F. Esposito, V. Roberto, and P. Zingaretti (Singapore: World Scientific), pp. 3-19.

Parisi, D., Cecconi, F. and Nolfi, S., 1990, Econets: neural networks that learn in an environment. **Network,** 1, 149-168.

Pelliccione, A. and Parisi, D. in preparation, *Hunger and thirst in artificial organisms*.

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

Piaget, J., 1952, *The Origins of Intelligence in Children* (New York: International University Press).

Purves, D., 1990, *Body and Brain* (Cambridge, Mass.: Harvard University Press).

Sims, K., 1994, Evolving virtual creatures. In *Siggraph 94 Conference Proceedings. Annual Conference Series*, pp. 15-24.

Williams, A., 2002, Facial expression of pain: an evolutionary account. **Behavioral and Brain Sciences**, 25, 439-488.

Ziemke, T., in press, Cybernetics and embodied cognition: on the construction of reality in organisms and robots. **Kybernetes**.

Ziemke, T., Bergfeldt, N., Buason, G., and Svensson, H., in press, Evolving cognitive scaffolding and environment adaptation: a new research direction in evolutionary robotics? **Connection Science**.

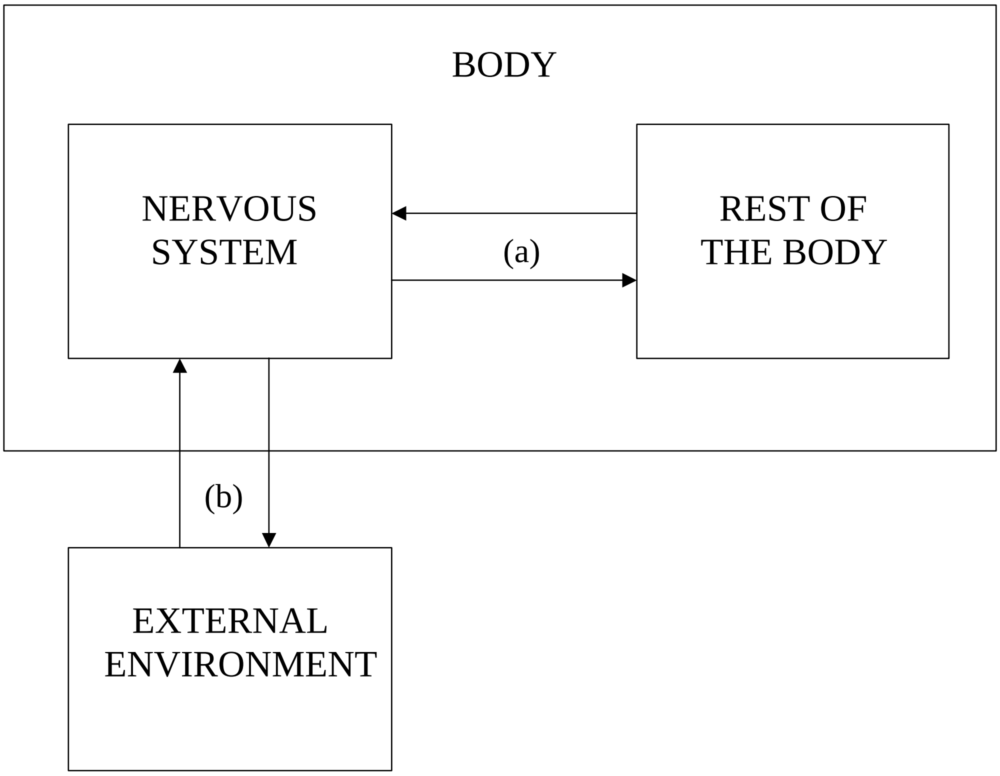
BODY
NERVOUS SYSTEM
REST OF THE BODY
(a)
(b)
EXTERNAL ENVIRONMENT

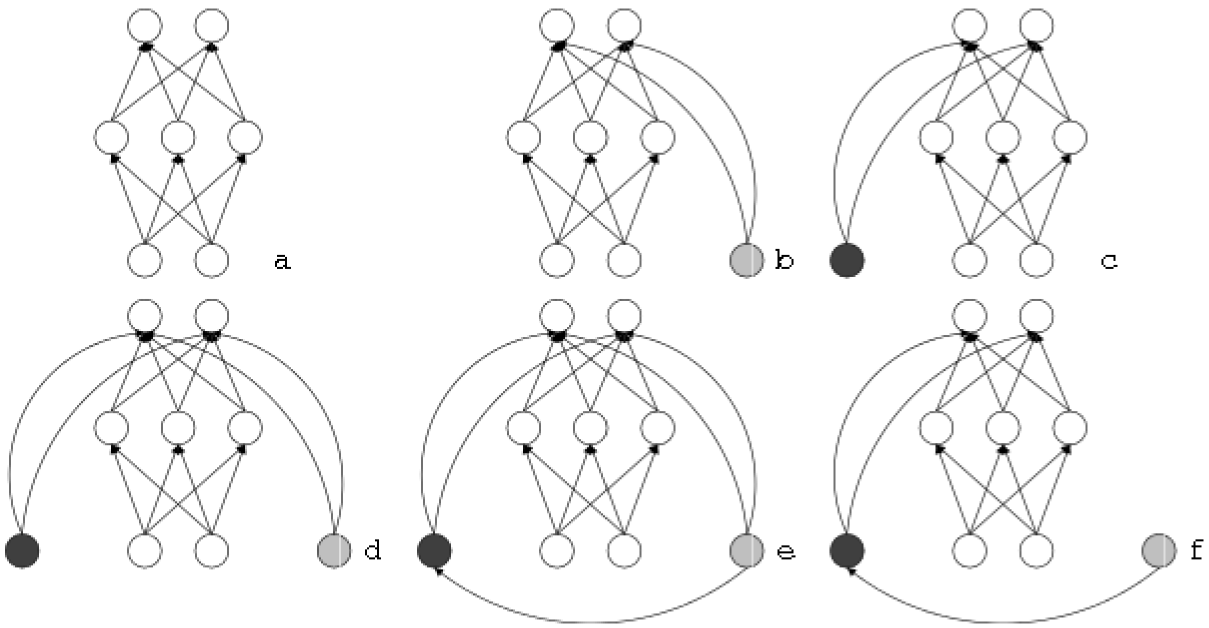
a
b
c
d
e
f

Figure 1. As a physical system the nervous system affects and is affected by two components of the environment which lies outside itself: the environment inside the organism's body, i.e. the rest of the organism's body (internal environment) (a), and the environment which lies outside the organism's body (external environment) (b).

Figure 2. Networks architectures for robots that move around looking for food during the day and rest at night: (a) simple architecture with input units encoding the position of the nearest food and output units encoding locomotion; (b) architecture with light sensor; (c) architecture with biological clock; (d) architecture with both light sensor and biological clock independently connected to the output units; (e) architecture with light sensor connected to both biological clock and output units; (f) architecture with light sensor only connected to biological clock.

Acknowledgment

Research partly funded by EU Project “Embodied and Communicating Agents”.