

Some Armchair Worries about Wheeled Behavior

Fred A. Keijzer

Section of Experimental and Theoretical Psychology
Leiden University, P.O. Box 9555, 2300 RB Leiden
The Netherlands
E-mail: keijzer@rulfsw.leidenuniv.nl

Abstract

Autonomous agents research aims to understand adaptive behavior by building models that exhibit such behavior. In this paper, I describe some theoretical worries about the tendency to simplify the sensory-motor control of these behaving models as much as possible. Wheeled robots provide a good example of this tendency. The worries derive from the idea that a complex sensory-motor system is a necessary requirement for reliable functional behavior in a natural environment. The evidence on which I base this conviction is three-fold: (a) when animals, but not wheeled robots, move across a surface this results from variable spatio-temporal patterns across a complex neuro-musculo-skeletal system, (b) a theoretical analysis according to which the stability of distal behavior arises as a result of the variability of proximal behavior, and (c) the possibility to interpret the behavior of animals, but not that of wheeled robots, as a process of self-organization. The use of wheeled robots therefore seems to side step a core problem for understanding adaptive behavior.

1. Introduction

Research on autonomous agents aims to understand adaptive behavior by building models—robots or animats—that exhibit such behavior. This paper presents some theoretical worries about current attempts to understand adaptive behavior in this way.

These worries derive from the fact that organismal and robot behavior differ considerably in the ratio of sensory-motor complexity to the complexity of the functional characteristics of the behavior. Behaving multicellular organisms embody a nervous system and a complex array of sensors and muscles and their opposing rigid structures. A complex neuro-musculo-skeletal system is at work behind every form of adaptively relevant, functional behavior—even when the latter remains very basic. On the other hand, within autonomous agent research (AAR) most robots—of course not all (Brooks, 1989; Beer, 1995)—are built with an eye to achieving complex adaptive relations with an environment while relying on a minimally complex sensory-motor system.

As compared to classical cognitive science, AAR pays a great deal of attention to the mediating details of behavior. The recent interest in adaptive behavior is based on the conviction that the concrete, physical sensory-motor interactions between a specific, embodied behaving system and the characteristics of the environment together generate the phenomena of adaptive behavior. Nevertheless,

the actual use of detailed and complex sensory-motor interactions in modeling adaptive behavior is kept to a minimum. It is thought to be sufficient that they produce the ‘really interesting’ functional characteristics of behavior, the characteristics that make behavior intelligent and adaptive.

This state of affairs worries me. I am led to think that studying the wider, functionally adaptive aspects of behavior apart from their naturally complex mediation is insufficient, and even a mistake. In this paper, I will argue that the most important bottle-neck for building robots whose behavior is on a par with that of biological organisms lies in coming to grips with the biological capacity for complex sensory-motor coordinations. AAR should pay much more attention to the organizing principles—the mechanisms—of organismal behavior because a complex sensory-motor system is a prerequisite for the very occurrence, as well as the stability, of the functional characteristics of behavior. I claim that *the core problem for autonomous agent research is to model functionally relevant behavior by means of a complex neuro-musculo-skeletal system*. The major challenge at present is not to model many different kinds of functionally relevant behavior, it is to model the organizing principles on which such functional behavior depends.

To contain this discussion about the importance of sensory-motor complexity, I will focus on one particular example of behavior where this issue comes to the fore very markedly: moving about over a firm surface. In AAR this behavioral task is often accomplished by wheeled robots, whereas in nature very different solutions dominate. What I want to focus on here is, of course, not the use of wheels in particular. Wheeled behavior is merely a good example of the tendency to simplify the sensory-motor processes in order to attend to the ‘adaptive’ or ‘intelligent’ aspects of behavior.

2. Movement over hard surfaces

Moving about over a firm surface is a common accomplishment for both animals and robots. Most robots effect this behavior by driving about on wheels. Two commonly used robots, the *Nomad 200* (Sharkey & Heemskerk, 1997; Milán, 1994) and *Khepera* (Floreano & Mondada, 1994), provide good examples of the sensory-motor system that allows a wheeled robot to move. Despite their differences in size (they are respectively 90 cm and 3 cm high) and approximate weight (respectively 60 kilograms and 70 grams), both are statically stable artifacts that rest on wheels. Electric motors set the wheels in motion and so propel the robot forward, backward or let it turn about. Because these are good robots—in the traditional robotic sense of being precise and dependable—there is a reliable one to one relation between a small number of different signals sent to the motor and the movement of the robot itself.

I state the obvious here to stress the point that the use of wheels makes movement across flat, hard surfaces a trivially easy task. Wheels make it possible to investigate functional behavior (e.g. goal seeking, learning, imitation, social interactions) for which moving about is a prerequisite, without putting too much work in the necessary sensory-motor system.

How does the wheeled behavior of robots compare to the accomplishments of real animals? When considering movement over a hard surface the end result is the same in both cases, but the means by which it is accomplished differ considerably.

The first thing to notice is the wide variety of mechanisms used to produce forward motion. Legs may be the most familiar means by which animals move about, but there are many other options. In particular, soft-bodied animals exploit many different strategies.

Some animals, such as small, free-living flatworms and nemertines, use *cilia* to crawl (Barnes, Calow & Olive, 1993; Trueman, 1975). Cilia are relatively short, bristle-like cell organelles that project from cells at the underside of the animal. A constant propulsive force is generated by the coordinated movement of a large number of cilia. Each cilium normally has the same pattern and frequency of beating, but individual cilia are at different phases of this cycle—so as not to interfere with one another’s powerstrokes—and waves of activity can be seen to pass across the ciliated surface. There are

strong size-imposed limitations on ciliary locomotion and its use remains restricted to small animals of usually less than 1 mm long.

Looping is a form of movement used by slightly bigger animals. It derives from muscle power and occurs in elongated, 'tube-like' animals. Some nemertines use looping as a fast escape mechanism, but it has been perfected by leeches (Trueman, 1975). Looping is based on the combined use of circular and longitudinal muscles, the forces of which antagonize one another through a hydrostatic skeleton, a relatively fluid body mass of fixed volume. The longitudinal muscles make the animal short and thick, the circular muscles turn it thin and elongated. Leeches have suckers at their posterior and anterior extremities which attach the leech to the substrate. When the posterior sucker is attached, the leech uses its circular muscles to stretch its body forward, after which it places its anterior sucker on the substrate at almost a full body length in the direction of travel. Then it loosens its posterior sucker, flexes its longitudinal muscles to draw in its body, and places its posterior sucker just behind the anterior one, after which the cycle repeats itself.

A similar musculo-skeletal organization is used to move by *pedal locomotory waves*. These are waves of contraction that pass along the body. The cyclical contraction of circular muscles followed by that of longitudinal muscles, as in looping, is here replaced by the successive contraction of different *regions* of the longitudinal muscles (Trueman, 1975). "A muscular locomotory wave consists essentially of raising a part of the body from the substratum, moving it forwards and reattaching it." (ibid., p.22). The waves that move across the body can be both direct (in the direction of the organism's movement) or retrograde (in the direction the organism is moving from). In the retrograde case, the thinning and elongating parts of the body are lifted free from the substratum and move forward while the thick and shortened parts of the body form the so-called *points d'appui* which anchor the body to the substrate. When the waves are direct the fully elongated body parts form the *points d'appui*, while the thickened parts of the body are the ones that move forward. In this case, the thin portions of the body usually extend stiff, bristle-like appendages to prevent slippage (Barnes, *et al.*, 1993).

There are many more ways in which animals move over hard surfaces: snakes undulate, and many animals walk. The details of these various ways of moving are all different. Still there are commonalities which group all these forms of locomotion together and contrast them with robot movement provided by wheels. Forward movement as it is produced by biological organisms occurs as collective property of many small-scaled components, together making up a neuro-musculo-skeletal system. Forward movement is here not the result of a simple, singular event which sets a motor on or off, as is the case in many robots.

Of course, everybody doing adaptive behavior research will be perfectly aware that driving across a hard surface differs in this respect from ciliating, looping, crawling or walking over it. However, the dominant idea is that the details of such mediating processes can be ignored when one studies the functional aspects of behavior. The use of wheels allows one to study some highly interesting aspects of adaptive behavior (motivation, learning, social relations, imitation etc.) which would be extremely difficult if one had to implement a complex neuro-musculo-skeletal system first.

Given the difficulty of modeling both the full complexity of the mediating processes involved in animal behavior as well as the functional regularities which result from these, what is wrong with a division of labor and focus on one of these issues at a time? Why worry? My nagging doubts here are that an important aspect of natural behavior is left out of consideration. This doubt is reinforced by some old theoretical ideas concerning the relation between sensory-motor coordinations and functional organism-environment interactions.

3. The many faces of behavior

Behavior is a multi-faceted phenomenon. It is both difficult to define and to analyze. This "does not arise from any lack of ways to do it but from an embarrassment of riches. We can describe an action as a set of muscle twitches, or as a sequence of movements of limbs and other parts, or as a sequence of goal-directed actions, or in even larger units" (Miller, Galanter & Pribram, 1960, p.13).

How is this riches to be dealt with? Ever since Tolman in the 1930s, all psychologists and biologists are familiar with the distinction between molecular and molar behavior—molecular behavior referring to little units and molar behavior to large ones—but this distinction does not play an important role in practical psychological research. As Van der Heijden states:

it is not the complete behaviour that is observed and registered, but only some specific, quantifiable, aspects, such as the latency and/or correctness of the response. Other aspects, e.g. the duration and the detailed execution of the observable behaviour are simply neglected (Van der Heijden, 1992, p.14).

Confronted with the "embarrassment of riches" of actual behavior, researchers—both psychologists and biologists—tend to concentrate on specific characteristics and leave the rest to other researchers. The choice of these characteristics is constrained by the available experimental and modeling possibilities, but also by preconceptions about what are the most crucial aspects of adaptive behavior. One may agree that the topic of interest is the modeling of adaptive behavior, but this commitment leaves unspecified which aspects of adaptive behavior are chosen to be modeled.

3.1 Brunswik's Lens Model

To provide some handles on the many different aspects of behavior it will be useful to turn to some older ideas, first developed by Fritz Heider (Heider, 1959) and later elaborated by Egon Brunswik (Brunswik, 1952; Campbell, 1966). They provide an instance of the general distinction between molecular and molar behavior, which is particularly congenial for elucidating my worries about wheeled behavior.

Heider and Brunswik were both concerned with the structure of the environment as a relevant factor for the adaptive behavior of organisms. In the 1920s, Heider made a distinction between the 'near' and 'far' environments of an organism and addressed the implications of this distinction for adaptation (Heider, 1959). Brunswik, influenced by Heider, used the notions of proximal and distal to mark the distinction between distal objects or events (the far environment of the organism), and the proximal stimuli (the near environment) which actually reach the organism as a pattern of sensory stimulation (Brunswik, 1952).

The near environment of an organism consists of proximal stimulation which impinges on its body directly. This stimulation in itself tends to be neutral as far as adaptation is concerned. The presence or absence of light- or soundwaves, or of air-borne particles does not usually make a dramatic impact on the life-processes of an organism. The near, proximal environment provides a relatively neutral medium in which the organism can move about. The neutral proximal lights, sounds and smells do, however, convey cues about the existence and whereabouts of distal, nonlocal features of the environment.

Heider and Brunswik stated that the adaptively relevant elements in an environment consist primarily of the distal objects or 'things' therein. These are usually at some distance from the organism and provide the possible nourishment, threats or mates on which survival and reproduction depend. Of course, it is only when these distal events or things come close to the organism that they become a matter of life or death. However, such moments—when the tiger is already at your throat—are relatively rare, and when they occur these normally distal—but now proximal—entities also become irrelevant as far as behavior is concerned. Nothing further can be done at such a moment (except to eat

or being eaten).

The background assumption is that a certain time-frame is necessary in which an organism still has the opportunity to act. During this time, an organism uses the proximal stimuli as cues to guide it to or away from significant distal objects or events. Behavior then consists of bringing about (or avoiding) specific close encounters with such—most of the time—distal stimuli. As Brunswik says: "Organisms may in this sense be characterized as 'stabilizers' of [distal] events or of relationships" (1952, p.20).

The problem with using proximal stimuli as cues for distal stimuli is their unreliability as signs of distal states. The relation between the proximal stimuli and distal events is not one to one. Similar proximal stimulus patterns can correspond to different distal states of affairs, and, the other way round, the same distal state of affairs can give rise to very divergent proximal patterns. On the behavioral side, the same principle applies. The behavioral means available to the organism are also proximal: it can make movements. A series of 'properly' controlled movements can result in the achievement of an adaptively relevant event. However, the proximal set of movements that leads toward a similar distal event (feeding, escape from predators) will differ every time.

Heider and Brunswik state that the problem to be solved by an organism amounts to the following: given a state in which the organism is immersed in a highly variable, proximal environment—as well as the sensory as at the effector side—how does it achieve a consistent coordination with relatively invariant distal events? This basic question can be illustrated with Brunswik's expository "Lens Model" (Brunswik, 1952; Campbell, 1966).




Figure 1. Brunswik's Lens Model (adapted from Brunswik, 1952).

The Lens Model conveys the idea of proximal variability and distal stability when the organism is observed over several behavioral trials. This particular figure displays another distinction. It consists of two 'lenses'. The first one displays how a particular distal object is, on different occasions, connected to a particular central (mental) state by different proximal sensor states. The second lens portrays how the presence of a central state leads to the achievement of a distal, functional achievement through variable proximal movements on different occasions. The global picture is subsequently one in which a locally, or proximally, operating system—the s-r part can be seen as a single 'lens'—enforces stable distal S-R relations.

Brunswik's talk about proximal s-r laws and distal S-R laws may be less congenial nowadays than it was almost fifty years ago. However, the picture can be easily modified and made to concur with the current insight that responses feed back to the stimulus situation. In that case, the laws become couplings between the organism and its environment at two (actually a continuum of) different scales:

sensory-motor couplings for proximally guided behavior and perception-action couplings for distally oriented behavior.

To summarize, a central aspect of behavior consists of an ever varying proximal mediation which stabilizes distal events—and so turns them into distal regularities. To return to the example of forward movement, this distal task is accomplished in many different proximal ways—ciliating, looping, undulating, walking—the execution of which is itself always variable. The net result may in all cases be the same, this result is always accomplished in slightly different ways.

3.2 Behavior as results rather than means

When behavior is divided into relatively low-level proximal characteristics—what might be called the means—and higher-level, distal characteristics—the results, or the ends of behavior—an interesting question presents itself: Which of these two deserves our most particular attention? The question is particularly important for modeling adaptive behavior. Not every aspect of behavior can be replicated in a simulation or a robot, and choices will have to be made.

I think that within current studies of adaptive behavior there is a strong tendency to favor and focus on behavior's distal aspects: Behavior is interpreted in terms of adaptively relevant *results*, rather than in terms of proximally mediating mechanisms—the *means*—that bring these results about.

First, it should be noted that there is a general human tendency to describe behavior in distal terms whenever these apply. Vallacher and Wegner (1987) investigated in what manner people describe what they are doing. They state that "when both a lower and a higher level act identity are available," (e.g. 'lifting a mug to my mouth' and 'drinking coffee') "there is a tendency for the higher level to become prepotent." (p.5) Distal descriptions remain dominant as long as they apply, that is, as long as the distal result of the behavior occurs in a regular way. When the process leading up to the distal result is disturbed, more proximal descriptions come to the fore. For example, one experiment on which they based their conclusions was done with subjects who first had to drink a cup of coffee. One group drank from a normal cup, the other from a very unwieldy one, weighing approximately 500 grams. The first group applied more general descriptions to this activity (such as "getting energized" and "promoting my caffeine habit") while the group with the unwieldy cup identified it in terms like "drinking a liquid" and "lifting a cup to my lips". The message: when distal descriptions apply, then these will be the ones that are preferentially used by human observers.

The general tendency to prefer distal descriptions of behavior seems to have had a large impact on the study of behavior. The Lens Model provides an explanation. The most salient regularities of behavior appear at a distal scale. Distal behavior is easily identifiable because it is perturbation resistant and repeatable. Similar goals are reached time and again. At a proximal scale, there is a wide scattering of behavior. The mediating process is highly variable and there are many different proximal routes leading to the same distal behavior. Unless one is specifically interested in the specific sensory-motor details of specific animals or robots, it does not seem to matter much how exactly the distal behavior comes about. As long as it does.

This practical reason for focussing on behavior's functional regularities is reinforced by the often held conviction that these regularities are also the ones that a behavioral science most particularly needs to explain. Its functionality makes adaptive behavior different from mere physical movements. Brunswik already stated that the more important psychological generalizations were to be found at a distal level, and this attitude can still be witnessed in current work on modeling adaptive behavior. It leads to a serious underestimation of the sensory-motor problems involved in generating such distal regularities.

McFarland and Bösner (1993) offer an example of this tendency within AAR. In their book on intelligent behavior in animals and robots, they described intelligence as "an assessment of performance in terms of some functional criteria" (p.24). Even though they stress the importance of the

body which actuates adaptive behavior, their interest centers on ways to optimize choices between multiple—already existing—options for distal behavior. For example, they discuss how an imaginary housekeeping robot chooses between several "simple behavioral tasks" like "gotoN, gotoK, gotoF, gotoW, getmop, returnmop, getcloth, windowclean, floorclean, and sleep" (p.217). Creating a mechanism that could physically perform such tasks is not treated as a problem. According to McFarland and Bösser, 'a' mechanism will do. They even want to leave the kind of mechanism involved in generating behavior unspecified as "it is unsatisfactory to pin the label 'intelligence' on certain mechanisms and not others" (p.16). The problem how any mechanism can be constructed which generates real, biological-style behavior is passed over and treated as solved.

McFarland and Bösser's treatment of intelligent behavior is just one example how behavioral research is skewed toward the distal, functional aspects of behavior, at the cost of more proximal aspects. Using wheeled robots to investigate adaptive behavior is important in this context. The primacy awarded to distal behavior and the use of wheeled robots mutually reinforce one another. The focus on distal behavior makes the exact execution of distal behavior less relevant so that wheeled robots—although biologically implausible—become a plausible way to model adaptive behavior. In return, wheeled robots make it relatively easy to model distally oriented behavior. With a minimally complex sensory-motor system, a wheeled robot is capable of performing intuitively interesting distal tasks, for example, making it possible to envision a housekeeping robot. If such a robot had to walk, crawl or loop, just moving about would become a serious achievement, leaving little place for housekeeping tasks.

At this place my basic worry about wheeled behavior comes to the fore. Wheels allow one to study distal behavior by simplifying the proximal mediating process: What if the proximal mediating mechanism is a much more central aspect of adaptive behavior than many tend to think? More specifically, what if a fairly complex proximal mediating process is an essential ingredient for bringing about natural adaptive behavior?

3.3 The importance of sensory-motor complexity

Brunswik's Lens Model points not only to behavior's distal aspects. The Lens Model also visualizes the proximal mediation of these distal aspects. Rather than putting one of the two in a privileged position, the Lens Model suggests an interpretation of behavior in which the relation between proximal and distal behavior is a key issue for understanding and modeling adaptive behavior.

The Lens Model visualizes distal behavioral regularities as a result of a 'focussing process' that is somehow mediated by proximal behavior. The more easily identifiable distal behavior occurs only because there is a proximal focussing process which is highly variable and context dependent, and therefore difficult to pin down.

The focussing metaphor implies that a proximal behavioral 'lens' is required to acquire stable distal behavior. To interpret this metaphor it is good to remember Ashby's law of requisite variety: "only variety can destroy variety" (Ashby, 1956, p.207). The behaving system needs complex behavioral capacities at a proximal scale—call it room for manoeuvring—to counteract the many proximal disturbances that will always occur in a natural environment. Only by bringing into play many different options for proximal behavior—that is when the sensory-motor system has many degrees of freedom—can stable distal regularities be ensured under natural circumstances.

When the 'diameter' of the proximal lens is too small—that is when the behaving system has too few sensory-motor degrees of freedom available to counteract disturbances—the proximal countermeasures will be insufficient and the distal regularities will break down or not occur at all.

The use of wheeled robots produces an extreme simplification of the proximal sensory-motor couplings related to moving about. The capacity for maintaining intuitively convincing distal

perception-action couplings is therefore severely limited. In terms of the Lens Model, the researcher collapses the proximal lens by selecting a single proximal behavior (driving) which tends to accomplish the distal task of moving about. The resulting distal behavior is brittle. Wheeled robots are restricted to uninterrupted flat surfaces and their capacity to move about breaks down as soon as it encounters steep slopes, gaps, muck, or merely any rough surface with stones and treeroots.

When a wheeled robot encounters a natural terrain, it can no longer be expected to achieve forward movement. In comparison, nature's solutions to the problem of forward motion are not so easily compromised. These solutions are all relatively complex from a proximal point of view. They are also capable of negotiating naturally occurring terrain. To name just one case, looping leeches are not restricted to flat hard surfaces. They can move vertically, negotiate gaps and outperform anything on wheels many times over when it comes to moving about under natural conditions.

A complex mediating mechanism is thus highly important to establish and maintain adaptive behavior that deals effectively with a distal environment. It seems to me therefore that *a key issue of adaptive behavior is not merely to achieve distal goals, but to achieve these goals under always varying proximal circumstances*. Proximal sensory-motor variability is necessary to give rise to distal stability.

3.4 A challenge for modeling adaptive behavior

Researchers working on the simulation and modeling of adaptive behavior will certainly agree that good models of adaptive behavior need to incorporate both proximal and distal behavior. But the claim that it will be necessary to make the proximal behavior sufficiently complex goes beyond this general agreement. This claim has implications that may be less readily shared by all.

One implication for modeling behavior is that wheeled robots do not capture the characteristic signature of natural adaptive behavior: proximal variability leading to distal stability. Wheeled robots achieve the distal effect of moving about with only minimal proximal variability. As a model this does not suffice because a key factor of adaptive behavior is edited out of the model.

When modeling adaptive behavior, the mediating mechanism is not an arbitrary interchangeable producer of movements. It has to fulfil more stringent demands relating to flexibility and adaptivity by being capable of coordinating a large number of degrees of freedom. This implies that adaptive behavior research ought to pay much more attention to *behavioral mechanisms* than is often the case. It is not sufficient that a robot achieves a distal behavioral task. The robot should perform this task in a way which resembles an animal's way of performing this distal task: The robot's performance should be based on the same organizing principles as animal behavior.

This message might be difficult to accept. At least two reasons can be brought against the idea. Firstly, it seems strange to uphold that robots who accomplish the same functional behavior as animals or humans are to be distinguished because of their supposedly 'different organizing principles' or underlying mechanisms. Why would it matter how a task was accomplished, as long as the job is done? Secondly, referring to different organizing principles seems to be rather vague and mystifying. Any reference to such principles should be given a more clear meaning before they can be taken seriously.

The first point can be dealt with rather quickly. Postulating a theoretical difference between the behavior of animals and that of wheeled robots looks only awkward under the assumption that the same tasks *are* accomplished in both cases. However, the argument is that robot and animal behavior have different characteristics. A robot may achieve a distal task under a narrow range of proximal circumstances while an animal can achieve the same distal goal under a much wider range of proximal circumstances. Robots and animals do therefore not accomplish the same distal tasks at all. The difference made between animal and (current) robot behavior on the basis of underlying mechanisms is

also reflected in the functional behavioral regularities of animals and robots.¹

The difference made between the behavior of animals and robots is based on the conviction that animals, but not robots—at least not as they have been built so far—rely on still unspecified organizing principles which form a necessary precondition for achieving the distal behavioral regularities which are so characteristic of animal behavior. Saying something about those principles requires more work.

4. Behavior as a self-organizing process

Stating that the proximal behavior of animals and robots derives from different organizing processes may sound either mysterious or trivial. Interpreted as the obvious difference between organics and electronics the case looks trivial. Of course there is a difference here, but this difference is of a kind which many would think should not matter. There do not seem to be any principled reasons for claiming that artificial creatures, build from electronic and mechanical components, are forever incapable of generating behavior in the same ways as biologicals do. I agree. The idea that animal and (current) robot behavior reflects different organizing principles refers to a more abstracted level of organization: the behavior of animals, but not that of robots, consists of a process of self-organization.

The notion of self-organization is used often and with different meanings. My usage derives mainly from the work of Haken (1987) and Kelso (1995), but see also Goodwin (1994) and Kauffman (1993). Self-organization occurs when a large number of dispersed, small-scaled processes (the microscopy) become coordinated and so generate a pattern at a much larger scale (the macroscopy). A standard example consists of the large-scale convection rolls that arise in a container with fluid when the heat gradient between the cold top and the warm bottom exceeds a certain threshold. Individual molecules, at first each with its own independent direction of movement, start to move together on a large scale. Another example consists of swarms of fireflies where the light flashes of the individual flies become synchronized so that the swarm flashes its light on and of as a whole.

The notion of self-organization is very intriguing, unfinished and one which could have wide repercussions for the study of adaptive behavior (see e.g. Kelso, 1995). As it cannot be adequately described in a short space, I will simply endorse the following points: (1) The world is organized along multiple spatio-temporal scales, each having its own, relatively independent, set of dynamical relations. For example, one can both look at colliding billiard balls or consider the molecules making up these billiard balls. (2) Generally, dynamical relations are studied at a particular scale. In contrast, self-organization is a process where two or more of these—usually independent—dynamics influence one another and become coordinated. (3) When self-organization occurs, the specific characteristics that arise at a macroscopic scale depend on the interactions between *many* small-scaled—microscopic—events. I take these three points to characterize what happens (among other things) when a process of self-organization occurs (see also Keijzer, in press).

These points tie up nicely with animal behavior in its form of forward motion across a firm surface. Forward movement involves at least two different scales of dynamical relations. There are a large number of small-scaled neuro-musculo-skeletal events and there is the large-scale movement of the organism. The biological cases of forward movement described above all consist of a large scale pattern (coordinated cilia, looping, pedal locomotory waves, coordinated legs) that arises as a result of many, smaller scaled interacting parts.

¹ Elsewhere (Keijzer, 1998), I signal the distinction between animal and robot behavior by making a difference between movements and behavior: Robots (at least those that have been build up to now) merely move, they do not behave. Only animals have so far been capable of true behavior. I admit that it is difficult to draw a crisp borderline between genuine behavior and mere movement, but the distinction is a good way to make the point that not everything that moves automatically behaves.

For example, inside worms a neural network is connected to a dispersed system of muscle fibers which are antagonized by a hydrostatic skeleton. Under the influence of the neural network the muscle fibers contract and extend again. At the much larger scale of the whole worm-body, the pedal locomotory waves can be observed that make the worm move over the substrate. This large scale set of dynamical relations can be studied independent from the small-scaled happenings inside the body. However, in forward movement the otherwise separate dynamics are closely related and together the many small-scaled neuro-musculo-skeletal events organize themselves as the large-scale locomotory waves of the worm as a whole. These locomotory waves can be interpreted as a macroscopic order—similar to the convection rolls in heated water—which arises as a result of a mutual influence with a microscopic level of organization.

To summarize, at the heart of behavior lies a capacity to organize spatio-temporal pattern, both across a nervous system and across a musculo-skeletal system. Certain of those spatio-temporal patterns result in forward movement on a hard surface.

The congeniality of self-organization and ‘simple’ forms of behavior is of course not a proof that self-organization actuality takes place in behavior. At present there is only a suggestive, general correspondence which will need a lot of ‘future work’ to turn it into a succinct and factual statement. Still, the link looks promising. Self-organization might provide the mechanism that is necessary to maintain distally stable behavior on the basis of microscopic variability. This is exactly what is necessary for adaptive behavior in a natural environment.

In addition, there are intriguing options for ‘steering’ such processes of self-organization by means of ‘internal control parameters’ situated within the set of small-scaled dynamical relations. This possibility has interesting implications for the debate on representations (Keijzer, in press), but this goes beyond the aim of the present paper. The message that I want to confer right now is merely that a process of self-organization makes the claim more solid that the behavior of animals derives from specific organizing principles, principles which are not necessarily replicated by any artifact capable of moving about.

Well, do wheeled robots rely on a process of self-organization in order to move about? It seems unlikely. In self-organization there are no direct links between small-scaled and large-scaled events, many small-scaled events make up a similar large-scale event on different occasions, but never in exactly the same way. In wheeled robots, the macroscopic forward motion does derive directly from a few small-scaled signals which control the motors and so what happens on a larger scale. The small-scaled (neural) processing merely selects which signal is sent to the motors, it does not generate a pattern on a bodily scale. The principles involved in self-organization and in wheeled robots are dissimilar in this respect. In other words, if the behavior of animals consists of a process of self-organization, then the behavior of animals and wheeled robots do not reflect the same abstract organizing principles.

5. What are the implications if these worries are justified?

My worries about wheeled behavior hover around the idea that animal behavior requires complex sensory-motor interactions in order to maintain stable functional (distal) behavior in a natural environment. The use of wheels for robots is an example of sidestepping the issue of proximal behavior’s complexity in order to study distal behavioral regularities. Averting this issue leads to the neglect of, what I think is a central aspect of natural behavior: how do variable proximal (fast and short-term) sensory-motor couplings maintain stable distal (slower and longer-term) perception-action couplings.

Assuming for now that the present analysis is on the right track, what are the implications? Foremost, it should be stressed that adaptive behavior is not a clear and circumscribed research topic.

Rather, what adaptive behavior actually amounts to is an open research question. The tendency to emphasize the functional, distal aspects of behavior ought to be balanced by attention to the mediating mechanisms on which biological systems depend to generate these distal aspects. Behavior should not be treated as a singular (distal) output but rather as the generation of spatio-temporal patterns along a proximal-distal dimension.

Additionally, it seems that behavior requires much more elaborate theorizing which explicates what actually goes on in behavior. The casual observation of an organism or a robot performing some intelligent-looking task does not suffice, and neither does the painstaking and highly systematic observance of outward animal behavior such as is done in ethology. There should be a characterization of the many interlocking aspects of behavior, ranging from the neural dynamics, the sensory-musculo-skeletal system, the actually made movements (Jacobs, et al., 1988), to the distal functional regularities of behavior. The big theoretical question is, how does all of this fit together?

The notion of self-organization seems to offer a route to work on this question concerning the organizing principles involved in such a multiply-scaled system. In turn, self-organization is itself a concept that needs working on, and behavior might be a good testing-ground for developing this notion further. The combination of self-organization and behavior thus looks like an interesting cluster for theory development in autonomous agent research.

In this perspective, a theory of behavior explains how the spatio-temporal patterns of behavior come about. To gain insight in these processes, it is relatively unimportant whether the modelled behavior is very 'intelligent' in any traditional sense. What matters is whether a model offers insight into the underlying multiply-scaled organization. The most simple forms of behavior are the most promising in this respect. For example, the most simple nervous systems in existence—the neural nets of *Cnidarians* (such as jelly fish)—do not show any cephalization or ganglions, but form a mesh which somehow coordinates and synchronizes large-scale movement across the bell of the jelly fish, making it contract and expand as a whole in a rhythmic way (Golding, 1993). From the perspective advocated here, this would be a highly interesting system to work on.

None of this involves a break with the major ideas of AAR. It is rather a plea to proceed along a path that has already been chosen. A core idea of AAR was named *emergent functionality* by Steels (1991). The message of emergent functionality is that complex functional behavior can 'emerge' from relatively simple interactions between an organism and an environment. No complex internal knowledgeable system is needed. The present argument presses this approach a step further by making the claim that the 'simple interactions between an organism and an environment' are themselves actually quite complex and 'emerge' from the simultaneous activity of many small-scaled processes dispersed over an extended neuro-musculo-skeletal system.

AAR started out with the claim that it would be beneficial for understanding intelligence to study systems capable of solving simple (non-cognitive) behavioral tasks, examples being wall-following and goal-seeking. AAR always opposed the idea that such behavior is too simple to be of much interest. The argument here is that AAR will have to look at behavior which is much simpler still, if it wants to gain insight in how real animals behave.

Literature

Ashby, W.R. (1956). *An introduction to cybernetics*. London: Methuen.

Barnes, R.S.K., Calow, P. & Olive, P.J.W. (1993). *The invertebrates: A new synthesis* (2nd ed.). Oxford: Blackwell.

Beer, R.D. (1995). A dynamical systems perspective on agent-environment interaction. *Artificial Intelligence*, 72, 173-215.

Brooks, R.A. (1989). A robot that walks: Emergent behaviors from a carefully evolved network. *Neural*

- Computation*, 1, 253-262.
- Brunswik, E. (1952). *The conceptual framework of psychology*. Chicago: University of Chicago Press.
- Brusca, R.C. & Brusca, G.J. (1990). *Invertebrates*. Sunderland, MA: Sinauer.
- Campbell, D.T. (1966). Pattern matching as an essential in distal knowing. In K.R. Hammond (Ed.), *The psychology of Egon Brunswik* (pp. 81-106). New York: Holt, Rinehart and Winston.
- Floreano, D. & Mondada, F. (1994). Automatic creation of an autonomous agent: genetic evolution of a neural-network driven robot. In D. Cliff, P. Husbands, J-A Meyer & S.W. Wilson (Eds.), *From animals to animats 3* (pp. 421-430). Cambridge, MA: MIT Press.
- Golding, D.W. (1993). Control systems. In R.S.K. Barnes, P. Calow & P.J.W. Olive *The invertebrates: A new synthesis* (2nd ed.) (pp. 402-452). Oxford: Blackwell.
- Goodwin, B.C. (1994). *How the leopard changed its spots: The evolution of complexity*. New York: Scribner's Sons.
- Haken, H. (1987). Synergetics: an approach to self-organization. In F.E. Yates (Ed.), *Self-organizing systems: The emergence of order* (pp. 417-437). New York: Plenum Press.
- Heider, F. (1959). On Perception and Event Structure and the Psychological Environment. *Psychological Issues*, 1(4), 1-123.
- Jacobs, W.J. et al. (1988). Observations. *Psychobiology*, 16, 3-19.
- Kauffman, S.A. (1993). *The origins of order: Self-organization and selection in evolution*. New York: Oxford University Press.
- Keijzer, F.A. (1998). *The generation of behavior: On the function of representation in organism-environment dynamics*. Leiden University. unpublished doctoral thesis.
- Keijzer, F.A. (in press). Doing without representations that specify what to do. *Philosophical Psychology*.
- Kelso, J.A.S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: MIT Press.
- McFarland, D. & Bösner, T. (1993). *Intelligent behavior in animals and robots*. Cambridge, MA: MIT Press.
- Millán, José del R. (1994). Learning efficient reactive behavioral sequences from basic reflexes in a goal-directed autonomous robot. In D. Cliff, P. Husbands, J-A Meyer & S.W. Wilson (Eds.), *From animals to animats 3* (pp. 266-274). Cambridge, MA: MIT Press.
- Miller, G.A., Galanter, E. & Pribram, K.H. (1960). *Plans and the structure of behavior*. New York: Holt, Rinehart and Winston.
- Sharkey, N.E. & Heemskerk, J.N.H. (1997). The neural mind and the robot. In A.J. Brown (Ed.), *Neural network perspective on cognition and adaptive robotics* (pp. 169-194). London: Institute of Physics press.
- Steels, L. (1991). Towards a theory of emergent functionality. In J-A. Meyer & S.W. Wilson (Eds.), *From animals to animats* (pp. 451-461). Cambridge, MA: MIT Press.
- Trueman, E.R. (1975). *The locomotion of soft-bodied animals*. London: Arnold.
- Vallacher, R.R. & Wegner, D.M. (1987). What do people think they're doing? Action identification and human behavior. *Psychological Review*, 94, 3-15.
- Van der Heijden, A.H.C. (1992). *Selective attention in vision*. London: Routledge.