

From Earwigs to Humans

Rodney A. Brooks
MIT Artificial Intelligence Laboratory
545 Technology Square
Cambridge, MA 02139, USA
brooks@ai.mit.edu

Abstract

Both direct, and evolved, behavior-based approaches to mobile robots have yielded a number of interesting demonstrations of robots that navigate, map, plan and operate in the real world. The work can best be described as attempts to emulate insect level locomotion and navigation, with very little work on behavior-based non-trivial manipulation of the world. There have been some behavior-based attempts at exploring social interactions, but these too have been modeled after the sorts of social interactions we see in insects. But thinking how to scale from all this insect level work to full human level intelligence and social interactions leads to a synthesis that is very different from that imagined in traditional Artificial Intelligence and Cognitive Science. We report on work towards that goal.

1 Introduction

There has long been a dichotomy in styles used in designing and implementing robots whose task is to navigate about in the real world.

Early on Walter (1950) developed simple robots that were based on reflex actions and simple associative learning Walter (1951). These robots used only a handful of switching elements, and their overall performance relied on interactions between their mechanical hardware, their electrical circuitry, and the properties of the environment. They were essentially a set of adaptable reflexes for what to do in particular perceptual situations.

Nilsson (1969)¹ describes an elaborate robot using a large mainframe computer of the day, that took perceptual inputs, built a world model, proved theorems about what must be true in that model, consulted goals given to it, and produced long plans of actions to achieve those goals.

The work of Walter is prototypical of what would now be called *behavior-based robotics*. That of Nilsson was

¹ See (Nilsson 1984) for a more complete collection of relevant papers from the late sixties and early seventies.

perhaps the ultimate defining moment of the grand days of *good old fashioned artificial intelligence (GOFAI)*.

This dichotomy extends through to today, although there was a long period where the Walter-style work disappeared completely. More recently, Brooks (1986b) introduced the *subsumption architecture* as an instance of a behavior-based approach to building robots that operate in the real world. There were three key ideas introduced in this work:

- Recapitulate evolution, or an approximation thereof, as a design methodology, in that improvements in performance come about by incrementally adding more situation specific circuitry² while leaving the old circuitry in place, able to operate when the new circuitry fails to operate (most probably because the perceptual conditions do not match its preconditions for operating). Each additional collection of circuitry is referred to as a new *layer*. Each new layer produces some observable new behavior in the system interacting with its environment.
- Keep each added layer as a short connection between perception and actuation.
- Minimize interaction between layers.

Brooks (1989) gives full details of the control system for a six-legged robot using the subsumption architecture—this is perhaps the purest illustration of the concepts. Brooks (1990) summarizes a large number of robots that use this exact approach.

The principles of behavior-based approaches to Artificial Intelligence, through building robots, are outlined in two manifestos, (Brooks 1991b)³ and (Brooks 1991a). The survey paper Brooks (1991c) gives at least one opinion on how these new ideas coalesced from a number of different groups (including Rosenschein & Kaelbling

²The *circuitry* might be implemented either directly as physical circuitry, or as software organized in a manner similar to circuitry.

³The original version of this was an MIT AI Laboratory memo (Brooks 1986a), which was widely circulated at the time. By the next year it was in its final published form and title (Brooks 1987), but it took a number of both years and rejections before it finally appeared in an archival journal.

(1986) and Agre & Chapman (1987)) working in similar directions. Maes (1990*b*) is a collection of papers on many of the early attempts within this framework, in both hardware and purely software implementations. Various conference proceedings include more up to date reports on hardware implementations of behavior-based mobile robots, e.g., (Meyer & Wilson 1991), (Meyer, Roitblat & Wilson 1993), (Cliff, Husbands, Meyer & Wilson 1994), (Brooks & Maes 1994), (Morán, Moreno, Merelo & Chacón 1995).

There is an interesting dualism between the two approaches to AI, where complementary aspects are designed-in explicitly to the systems, and converse aspects are emergent from the interactions of the parts that are designed in. E.g., in GOFAI behavioral responses are not designed in, rather they emerge from the interplay of the planner with the given goals, and the particular world model that has been constructed from sensory data. In behavior-based systems, by contrast, behavioral responses are explicitly designed into the system but there are not any explicitly represented goals. Nevertheless, it is natural for observers to attribute dynamically selected goals to the robots as they operate in an environment.

The contrasts between the two styles of building robots and their programs can be summarized in the following table. Each line attempts to summarize loosely some aspect of the duality between the two approaches.

	GOFAI	Behavior-Based
<i>designed</i>	sensor fusion	matched filters
<i>in</i>	models	action selection
	goals	action schemas
	plans	behavior
<i>problems</i>	search	binding problem
<i>output</i>	choose next	concurrent
	action	actions
<i>emergent</i>	behavioral	(apparent) goals
	responses	and plans

In GOFAI there is never a sequence of actions explicitly represented *a priori* in the robot's program. Rather that sequence is produced as a result of reasoning about the world model and the goals assigned to the system. In the behavior-based approach there are often specific actions represented, coupled to other actions or perceptual conditions. But, the specific goals of the robot are never explicitly represented, nor are there any plans—the goals are implicit in the coupling of actions to perceptual conditions, and apparent execution of plans unroll in real-time as one behavior alters the robot's configuration in the world in such a way that new perceptual conditions trigger the next step in a sequence of actions.

2 Two Shifts in Viewpoint

The introduction of the behavior-based approach to building robots (see the collection Maes (1990*b*)) was based on one set of shifts in viewpoint. The work on Cog is inspired by a second set of shifts in viewpoint.

2.1 Traditional to behavior-based: first shift

In moving from traditional AI to behavior-based robotics, there were both distinct sudden changes in viewpoint and a more gradual background change in viewpoint. The first and rather sudden change was to drop the requirement that systems which were to act intelligently should have a central world model.

Such world models were typically symbolic in nature. For the purposes of this discussion we will define a *symbol* as a structure that is independent of its location in memory, and can be referred to by arbitrarily many other remote structures, giving separate computational processes working on those remote structures instant access to the symbol. Although not necessarily a requirement, in almost all traditional systems there was a correspondence between a symbol and some aspect of the world that was designed in by the human building the system, and this correspondence, at least in principle, was transparently obvious to another skilled person looking at the system.

The sudden change away from symbolic models was evident in all of Rosenschein & Kaelbling (1986), Agre & Chapman (1987), and Brooks (1986*b*).

The second change in moving from traditional to behavior-based approaches was more gradual, and is not as fully accepted in the behavior-based community today as is the first. The change was to get away from the idea that intelligence is a computational process that takes an input and produces an output⁴. In the new approaches, agents became situated in a world, so one had to consider the dynamics of their interactions with the world, along with their internal processes in order to understand what it is that they were capable of and were doing. This view was first evident in Agre & Chapman (1987), and became more explicit over time, e.g., Horswill & Brooks (1988) and Maes (1990*a*). See Brooks (1991*c*) for more discussions of these issues.

Such changes in perspective lead over time to the adoption of dynamical systems theory (see Beer (1995*a*), Smithers (1994), and Steinhage & Schoener (1996)) as a tool, framework and language, to describe the coupled agent and its environment, harking back to the early

⁴An even worse formalization confuses the state of the world with both the output of the perceptual apparatus and the effects of the motor system so that intelligence is then reduced to being a function I , mapping the current state s to the next state $I(s)$.

days of cybernetics, and in fact recapitulating that early work (Ashby 1952). The jury is still out as to whether these analysis tools will lead to great conceptual or practical advances.

With these two changes, the earlier vocabulary of classical AI becomes largely irrelevant. Gone are explicit representations of beliefs, desires and intentions. The very notion of a speech act no longer makes sense. And gone were the grand difficulties that AI artificially introduced by mistakenly trying to reduce an essentially external descriptive language to an internal calculus of reason at runtime.

2.2 Current behavior-based research

In order to make an argument about how we need to develop the ideas of the behavior-based approach further in order to build something as complex as a humanoid robot exhibiting a full range of human behavior, we first need to establish the level of complexity of systems built using current approaches. A full survey of the current state would be a massive undertaking. As a crude approximation we could consider the contents of the journal *Adaptive Behavior* from MIT Press. This journal is perhaps the journal most closely aligned with the behavior based approach. Its first three volumes (twelve issues) were published from 1992 to 1995, so are all rather current. We have carried out an analysis of the 42 papers which appeared in those volumes.

There has been very little work in the way of evaluating the *complexity* or *level of performance* of robotic systems, behavior-based or otherwise. The closest perhaps to this is the practice, such as in (Wilson 1985), of categorizing the environments in which the creatures, animats, or robots operate. Following that practice we defined six categories of the type and complexity of coupled environment-robot systems and then placed each of the 42 papers into one of these categories. We then analyzed the sorts of activities the robots were to perform.

2.2.1 Categories of domain

Our six categories are as follows:

There was no attempt to model a robot in any sense; all these papers either reported on animal or human studies, or looked at the mathematics of some particular technique in isolation.

R0 The ‘robots’ are simulated but have no spatial extent, either in size or position, e.g., the paper might concern experiments with mating and use a probability to control whether two individuals mate, rather than spatial adjacency.

R1 The robots were simulated and there was a direct correspondence between the world model and what the sensors delivered, and how the actuators worked in the world.

R2 As with category **R1**, but with stochastic noise added.

R3 Again simulated robots, but now with either, or both, a simulated physics of the sensors and actuators.

R4 An physical robot with sensors and actuators operating in a physical world.

category	number of papers
#	7
R0	5
R1	17
R2	7
R3	2
R4	4
	42

Very little work is done with either physical robots or even realistic simulations of robots. In fact the vast majority of papers can best be described as computational experiments. The table above was derived from the following raw classifications.

#. The majority of papers here describe experiments with or observations of particular animals, sometimes with neural network implementations trying to match the learning observed in animals. The particular animals used are fish (Halperin & Dunham 1992), toads Ewert, Beneke, Buxbaum-Conradi, Dinges, Fingerling, Glasgow, Schürg-Pfeiffer & Schwippert (1992), worms (Staddon 1993), toads (Wang 1993), human infants (Rutkowska 1994), and rabbits (Kehoe, Horne & Macrae 1995). Beer (1995*b*) tries to classify the possible classes of dynamics of recurrent neural networks for very small network sizes.

R0. These dimensionless robots are variously concerned with responding to visual stimuli (Fagg & Arbib 1992), modelling mating behavior (Todd & Miller 1992), communication acts (MacLennan & Burghardt 1993), sequence generation (Yamauchi & Beer 1994), and abstract group decision making (Numaoka 1994).

The papers from the remaining four classes get analyzed in other ways below, so we simply list them by category here.

R1. Simulated robots with exact correspondences between the world and sensors and actuators were described in Beer & Gallagher (1992), Müller-Wilm, Dean,

Cruse, Weidemann, Eltze & Pfeiffer (1992), Koza, Rice & Roughgarden (1992), Klopff, Morgan & Weaver (1993), Baird III & Klopff (1993), Schmajuk & Blair (1993), Tyrrell (1993), Peng & Williams (1993), Cliff & Bullock (1993), Colombetti & Dorigo (1994), Nolfi, Elman & Parisi (1994), Ling & Buchal (1994), Cliff & Ross (1994), Gruau (1994), Wheeler & de Bourcier (1995), Schölkopf & Mallot (1995), and Corbacho & Arbib (1995).

R2. Simulated robots with noisy correspondences between the world and sensors and actuators were described in Grefenstette (1992), Arkin (1992), Liaw & Arbib (1993), Szepesvári & Lörincz (1993), Ram, Arkin, Boone & Pearce (1994), Tyrrell (1994), and Cruse, Brunn, Bartling, Dean, Dreifert, Kindermann & Schmitz (1995).

R3. Simulated robots with simulated sensors and actuators were described in Cliff, Harvey & Husbands (1993) and Ekeberg, Lansner & Grillner (1995).

R4. Experiments with physical robot systems were described in Espenschied, Quinn, Chiel & Beer (1993), Kube & Zhang (1993), Ferrell (1994), and Hallam, Halperin & Hallam (1994).

Some further classification of the domains is also possible.

Six of the **R1** papers ((Tyrrell 1993), (Peng & Williams 1993), (Cliff & Bullock 1993), (Nolfi et al. 1994), (Cliff & Ross 1994), (Corbacho & Arbib 1995)) concerned environments which were discrete grids rather than a continuous model of space. Two of the **R2** papers ((Szepesvári & Lörincz 1993), (Tyrrell 1994)) also used a discrete grid model of space.

Two of the **R0** papers, and one each of the **R1**, **R2**, and **R4** papers were explicitly concerned with the dynamics of multiple interacting robots. These were MacLennan & Burghardt (1993), Numaoka (1994), Wheeler & de Bourcier (1995), Grefenstette (1992), and Kube & Zhang (1993) respectively.

Only two papers were concerned with robots that could move in three dimensions; the **R2** paper Ram et al. (1994) and the second half of the **R3** paper Ekeberg et al. (1995).

2.2.2 Categories of task or activities

Now let us turn our attention to what it is that the robots do. We will only concern ourselves with robots with some spatial component to their activities, i.e., the robots from papers in classes **R1** through **R4**. There are 30 such papers, but one, Beer & Gallagher (1992), has two major examples so there are a total of 31 examples to consider.

Modulo some clarifications below there seem to be four primary activities to which the examples can be uniquely assigned. The *navigation* activity can be further subdivided into three types of navigation, although some of the papers are a little hard to classify exactly according to this scheme.

Primary Activity		Examples
Locomotion		7
Pole balancing		2
Box pushing		1
Navigation	<i>taxis</i>	7
	<i>predator/prey</i>	3
	<i>moving about</i>	11
		31

The subclassifications within navigation are a little difficult to make as the distinction between some sort of *taxis* and chasing some simulated *prey* is often purely in the words used by the authors rather than any deep semantic meaning in what is presented in the papers. Some of the papers categorized as *moving about* do have the notion of the robot, or creature, capturing and eating prey, but except for the three papers (Szepesvári & Lörincz 1993), Tyrrell (1993), and Tyrrell (1994), there is no specific action that need be taken—simply landing on top of the thing to be eaten is sufficient for it to be eaten.

In order to compare the complexity of the activities of these 31 systems we can talk about the actuator complexity, and the activity complexity.

First, let us consider the actuator complexity.

The actuator models on all the 24 non-locomotion examples are all rather simple. In the discrete worlds the actions are very simple—moving from a grid square to one of its four or eight neighbors (note that this includes the two papers Tyrrell (1993) and Tyrrell (1994)). Many other papers have an actuator model where a heading and range is specified and the robot moves there. The most complex systems have two motors providing a differential drive mechanism, where both motors must go forward to drive the robot forward, and differences in motor velocities result in steering left or right, or in place in the limit; there are two such papers, the **R3** navigation paper Cliff et al. (1993) and the **R4** box-pushing paper Kube & Zhang (1993). The **R4** navigation paper Hallam et al. (1994) appears to use a *synchro-drive* where heading and velocity are independent motors. Thus, none of the 24 examples have more than two actuators, and most have a considerably more abstract version of actuation than even that.

The seven locomotion papers concern robots with many more actuators. Six of the papers are about six legged walking, but three of the four simulation papers do not really even have a third dimension in their leg kinematics and have only one and a half (bang-bang up-down) actuators per leg. The two physically imple-

mented **R4** robots, those of Espenschied et al. (1993) and Ferrell (1994) do have full three-dimensional kinematics and have 12 and 19 actuators respectively. The latter paper is not concerned with locomotion *per se*, but rather with adapting to sensor failure. (Note that Cruse et al. (1995) show a photo of a physical six-legged robot but that is not the actual topic of the experiments in their paper.) The seventh paper (Ekeberg et al. 1995) concerns **R3** simulated Lampreys (a type of eel), and there are a number of actuators per body segment, with potentially an unlimited number of segments. None of these seven papers concerns more than simple walking, so there really is only one primary activity and not any sub-behaviors to be considered (note though, that in other papers (Ferrell 1995*a*), and (Ferrell 1995*b*), Ferrell does discuss multiple underlying behaviors contributing to six legged locomotion over rough terrain).

Now, let us consider activity complexity.

Of the 31 examples listed above, 29 of them have *only* the primary activity as their sole activity. There is no need to select between activities for the robot in those 29 cases. We will consider the other 2 examples separately below. Most of the other robots implement their primary activity as a combination of multiple behaviors, and in this case there is some need to select amongst those behaviors.

There are, however, 7 cases where the primary activity is made up of a number of simpler behaviors, only some of which may be active at any time. The number of behaviors ranges from two to nine, and in order of increasing number of behaviors these cases are described in the papers of Szepesvári & Lörincz (1993) (2 behaviors), Ram et al. (1994) (3 behaviors), Wheeler & de Bourcier (1995), Hallam et al. (1994), Corbacho & Arbib (1995) (all four behaviors), Kube & Zhang (1993) (5 behaviors), and Arkin (1992) (9 behaviors). For all these papers there are still only a handful (or two) of behaviors, and the selection mechanisms are each in their own way fairly uniform over the set of behaviors from which they must choose. Since there is only one major activity that is always ongoing there is no need to have big shifts, rather one can view behavior selection in these cases as simply modulating how the primary activity is being achieved. Note also that in order to implement a behavior, at most two motors need to be coordinated in all these systems, and it is fairly easy to come up with an orthogonal decomposition for a motor control system for even the case of two motors.

The two papers that specifically deal with activity selection are those of Tyrrell (1993) and Tyrrell (1994). The first has 5 major activities, but really only a simple model of motor control underlying the generation of these activities. The second paper has 29 different primary activities and 35 underlying behaviors or actions. While these papers are far and away challenging

the most complex activity control issues they are still both rather simple. They are **R1** and **R2** papers respectively, and both operate in a discrete grid domain. Neither paper has a non-trivial actuator model. And in fact many of the behaviors in the second paper are very atomic, such as *court*, *mate*, and *sleep* which correspond directly to global activities. These need only be activated at discrete times, exactly in the robot's current location. While there are interesting issues in which of the activities, and hence behaviors or actions, to select at any given time, the dynamics of interactions of these activities and the robot and the environment are not very complex.

2.3 Behavior-based to cognitive robotics: second shift

We are now in a position to consider the new second shift in viewpoint which is necessary if we are to build robot systems that have behavior that is of similar complexity to that of humans.

In thinking about building a full human level intelligence that is able to operate and interact in the world in much the way a human would operate and interact. We are led to a decomposition that is different from both the traditional AI approach and the behavior-based approach to mobile robots. This new decomposition is based on fundamentally different concerns at many different levels of analysis.

In order to act like a human⁵ an artificial creature with a human form needs a vastly richer set of abilities in gaining sensor information, even from a single vantage point. Some of these basic capabilities include saccading to motion, eyes responding appropriately to vestibular signals, smooth tracking, coordinating body, head, and eye motions, compensating for visual slip when the head or body moves, correlating sound localization to oculomotor coordinates, maintaining a zero disparity common fixation point between eyes, and saccading while maintaining a fixation distance. In addition it needs much more coordinated motor control over many subsystems, as it must, for instance, maintain body posture as the head and arms are moved, coordinate redundant degrees of freedom so as to maximize effective sense and work space, and protect itself from self-injury.

In using arms and hands some of the new challenges include visually guided reaching, identification of self motion, compliant interaction with the world, self protective reflexes, grasping strategies without accurate geometric models, material estimation, dynamic load balancing, and smooth on-the-fly trajectory generation.

⁵It is difficult to establish hard and fast criteria for what it might mean to *act like a human*—roughly we mean that the robot should act in such a way that an average (whatever that might mean) human observer would say that it is acting in a human-like manner, rather than a machine-like or alien-like manner.

In thinking about interacting with people some of the important issues are detecting faces, distinguishing human voices from other sounds, making eye contact, following the gaze of people, understanding where people are pointing, interpreting facial gestures, responding appropriately to breaking or making of eye contact, making eye contact to indicate a change of turn in social interactions, and understanding personal space sufficiently.

Besides this significantly increased behavioral repertoire there are also a number of key issues that have not been so obviously essential in the previous work on behavior-based robots. When considering cognitive robotics, or *cognobotics*, one must deal with the following issues:

- bodily form
- motivation
- coherence
- self adaptation
- development
- historical contingencies
- inspiration from the brain

We will examine each of these issues in a little more detail in the following paragraphs. While some of these issues have been touched upon in behavior-based robotics it seems that they are much more critical in very complex robots such as a humanoid robot.

Bodily form. In building small behavior-based robots the *overall* morphology of the body has not been viewed as critical. The *detailed* morphology has been viewed as very critical, as very small changes in morphology have been observed to make major changes in the behavior of a robot running any particular set of programs as the dynamics of its interaction with the environment can be greatly affected. But, there has been no particular reason to make the robots morphologically similar to any particular living creature.

In thinking about human level intelligence however, there are two sets of reasons one might build a robot with humanoid form.

If one takes seriously the arguments of Johnson (1987) and Lakoff (1987), then the form of our bodies is critical to the representations that we develop and use for both our *internal thought*⁶ and our language. If we are to build a robot with human like intelligence then it must have a human like body in order to be able to develop similar sorts of representations. However, there is a large cautionary note to accompany this particular

⁶Whatever that might mean.

line of reasoning. Since we can only build a very crude approximation to a human body there is a danger that the essential aspects of the human body will be totally missed. There is thus a danger of engaging in cargo-cult science, where only the broad outline form is mimicked, but none of the internal essentials are there at all.

A second reason for building a humanoid form robot stands on firmer ground. An important aspect of being human is interaction with other humans. For a human-level intelligent robot to gain experience in interacting with humans it needs a large number of interactions. If the robot has humanoid form then it will be both *easy* and *natural* for humans to interact with it in a human like way. In fact it has been our observation that with just a very few human-like cues from a humanoid robot, people naturally fall into the pattern of interacting with it as if it were a human. Thus we can get a large source of dynamic interaction examples for the robot to participate in. These examples can be used with various internal and external evaluation functions to provide experiences for learning in the robot. Note that this source would not be at all possible if we simply had a disembodied human intelligence. There would be no reason for people to interact with it in a human-like way⁷.

Motivation. As was illustrated in section 2.2, recent work in behavior-based robots has dealt with systems with only a few sensors and mostly with a single activity of navigation. The issue of the robot having any motivation to engage in different activities does not arise. The robots have some small set of navigation behaviors built-in, some switching between behaviors may occur, and the robot must simply navigate about in the world⁸. There is no choice of activities in these papers, so there needs be no mechanism to motivate the pursuit of one activity over another. The one major activity that these robots engage in, has, naturally, a special place, and so the systems are constructed so that they engage in that activity—there is no mechanism built in for them not to engage in that activity.

⁷One might argue that a well simulated human face on a monitor would be as engaging as a robot—perhaps so, but it might be necessary to make the face appear to be part of a robot viewed by a distant TV camera, and even then the illusion of reality and engagedness might well disappear if the interacting humans were to know it was a simulation. These arguments, in both directions are speculative of course, and it would be interesting, though difficult, to carry out careful experiments to determine the truth. Rather than being a binary truth, it may well be the case that the level of natural interaction is a function of the physical reality of the simulation, leading to another set of difficult engineering problems. Our *experience*, a terribly introspective and dangerous thing in general, leads us to believe that a physical robot is more engaging than a screen image, no matter how sophisticated.

⁸Navigation of course is not simple in itself, as was pointed out by Moravec (1984), but here we are talking about robots which have almost no other high level or emergent activities besides navigation.

When we move to much more complex systems there are new considerations:

- When the system has many degrees of freedom, or actuators, it can be the case that certain activities can be engaged in using only some subsystems of the physical robot. Thus it can be the case that it may be possible for the system to engage in more than one activity simultaneously using separable sensory and mechanical subsystems.
- The system may have many possible activities that it can engage in, but these activities may conflict—i.e., if engaging in activity A it may be impossible to simultaneously engage in activity B.

Thus for more complex systems, such as with a full human level intelligence in a human form, we are faced with the problem of *motivation*. When a humanoid robot is placed in a room with many artifacts around it, why should it interact with them at all? When there are people within its sensory range why should it respond to them? Unlike the mobile robot case where an implicit unitary motivation sufficed, in the case of a full humanoid robot there is the problem of a confrontation with a myriad of choices of what or who to interact with, and how that interaction should take place. The system needs to have some sort of motivations, which may vary over time both absolutely and relatively to other motivations, and these motivations must be able to reach some sort of expression in what it is that the humanoid does.

Coherence. Of course, with a complex robot, such as a humanoid robot with many degrees of freedom, and many different computational and mechanical and sensory subsystems another problem arises. Whereas with small behavior-based robots it was rather clear what the robot had to do at any point, and very little chance for conflict between subsystems, this is not the case with a humanoid robot.

Suppose the humanoid robot is trying to carry out some manipulation task and is foveating on its hand and the object with which it is interacting. But, then suppose some object moves in its peripheral vision. Should it saccade to the motion to determine what it is? Under some circumstances this would be the appropriate behavior, for instance when the humanoid is just fooling around and is not highly motivated by the task at hand. But when it is engaged in active play with a person, and there is a lot of background activity going on in the room this would be entirely inappropriate. If it kept saccading to everything moving in the room it would not be able to engage the person sufficiently, who no doubt would find the robot's behavior distracting and even annoying.

This is just one simple example of the problem of coherence. A humanoid robot has many different subsystems, and many different low level reflexes and behavioral patterns. How all these should be orchestrated, especially without a centralized controller, into some sort of coherent behavior will become a central problem.

Self adaptation. When we want a humanoid robot to act fluently in the world interacting with different objects and people it is a very different situation to that in classical robotics (Brooks 1991c) where the robot essentially goes through only a limited number of stereotyped actions. Now the robot must be able to adapt its motor control to changes in the dynamics of its interaction with the world, variously due to changes in what it is grasping, changes in relative internal temperatures of its many parts—brought about by the widely different activities it engages in over time, drift in its many and various sensors, changes in lighting conditions during the day as the sun moves, etc. Given the wide variety of sensory and motor patterns expected of the robot it is simply not practical to think of having a fully calibrated system where calibration is separate from interaction. Instead the system must be continuously self adapting, and thus self calibrating. The challenge is to identify the appropriate signals that can be extracted from the environment in order to have this adaptation happen seamlessly behind the scenes.

Two classical types of learning that have been little used in robotics are habituation and sensitization. Both of these types of learning seem to be critical in adapting a complex robot to a complex environment. Both are likely to turn out to be critical for self adaptation of complex systems.

Development. The designers of the hardware and software for a humanoid robot can play the role of evolution, trying to instill in the robot the resources with which evolution endows human individuals. But a human baby is not a developed human. It goes through a long series of developmental stages using scaffolding (see (Rutkowska 1994) for a review) built at previous stages, and expectation-based drives from its primary care givers to develop into an adult human. While *in principle* it might be possible to build an adult-level intelligence fully formed, another approach is to build the baby-like levels, and then recapitulate human development in order to gain adult human-like understanding of the world and self. Such coupling of the development of cognitive, sensory and motor developments has also been suggested by Pfeiffer (1996). Cognitive development is a completely new challenge for robotics, behavior-based or otherwise. Apart from the complexities of building development systems, there is also the added complication that in humans (and, indeed, most animals) there

is a parallel development between cognitive activities, sensory capabilities, and motor abilities. The latter two are usually fully formed in any robot that is built so additional care must be taken in order to gain the natural advantage that such lockstep development naturally provides.

Historical contingencies. In trying to understand the human system enough to build something which has human like behavior one always has to be conscious of what is essential and what is accidental in the human system. Slavishly incorporating everything that exists in the human system may well be a waste of time if the particular thing is merely a historical contingency of the evolutionary process, and no longer plays any significant role in the operation of people. On the other hand there may be aspects of humans that play no visible role in the fully developed system, but which act as part of the scaffolding that is crucial during development.

Inspiration from the brain. In building a human-level intelligence, a natural approach is to try to understand all the constraints that we can from what is known about the organization of the human brain. The truth is that the picture is entirely fuzzy and undergoes almost weekly revision as any scanning of the weekly journals indicates. What once appeared as matching anatomical and functional divisions of the brain are quickly disappearing, as we find that many different anatomical features are activated in many different sorts of tasks. The brain, not surprisingly, does not have the modularity that either a carefully designed system might have, nor that a philosophically pure, wished for, brain might have. Thus skepticism should be applied towards approaches that try to get functional decompositions from what is known about the brain, and trying to apply that directly to the design of the processing system for a humanoid robot.

3 Progress on Cog

Work as proceeded on Cog since 1993⁹

⁹The original paper on the rationale for Cog, (Brooks & Stein 1993), was rather optimistic in its timescale, as that was based on the intent that we would be able to get funding to tackle the project head on. We had a plan which included 5 full time engineers working on the project, plus graduate students and post doctoral students doing the programming. We were not successful at raising that money so we have had to live with a greatly scaled back project. In particular, all the engineering has been done by graduate students and one part time engineer. This note is meant to explain why the original timescale has not been met. There have not been technical problems that have unexpectedly held things up, rather there has been less effort than originally hoped for.



Figure 1: The robot Cog, looking at its left arm stretched out

3.1 The Cog robot

The Cog robot, figure 1, is a human sized and shaped torso from the waist up, with an arm (soon to be two), an onboard claw for a hand (a full hand has been built but not integrated), a neck, and a head.

The torso is mounted on a fixed base with three degree of freedom hips. The neck also has three degrees of freedom. The arm(s) has six degrees of freedom, and the eyes each have two degrees of freedom.

All motors in the system are controlled by individual servo processors which update a PWM chip at 1000Hz. The motors have temperature sensors mounted on them, as do the driver chips—these, along with current sensors, are available to give a kinesthetic sense of the body and its activities in the world. Those motors on the eyes, neck, and torso all drive joints which have limit switches attached to them. On power up, the robot subsystems must drive the joints to their limits to calibrate the 16 bit shaft encoders which give position feedback.

The eyes have undergone a number of revisions, and will continue to do so. They form part of a high performance active vision system. Each eye has a separate pan and tilt motor¹⁰. Much care has been taken to reduce the drag from the video cables to the cameras. The eyes are capable of performing saccades with human level speeds, and similar levels of stability. A vestibular system using three orthogonal solid state gyroscopes and two inclinometers will soon be mounted within the head.

¹⁰In the next revision a single tilt motor will drive the two eyes.

The arm demonstrates a novel design in that all the joints use series elastic actuators (Williamson 1995). There is a physical spring in series between each motor and the link that it drives. Strain gauges are mounted on the spring to measure the twist induced in it. A 16 bit shaft encoder is mounted on the motor itself. The servo loops turn the physical spring into a virtual spring with a virtual spring constant. As far as control is concerned one can best think of controlling each link with two opposing springs as muscles. By setting the virtual spring constants one can set different equilibrium points for the link, with controllable stiffness (by making both springs proportionately more or less stiff).

The currently mounted hand is a touch sensitive claw.

The main processing system for Cog is a network of Motorola 68332s. These run a multithreaded lisp, **L** (written by the author), a downwardly compatible subset of Common Lisp. Each processing board has eight communications ports. Communications can be connected to each other via flat cables and dual ported RAMs. Video input and output is connected in the same manner from frame grabbers and to display boards. The motor driver boards communicate via a different serial mechanism to individual 68332 boards. All the 68332's communicate with a Macintosh computer which acts as a file server and window manager.

3.2 Experiments with Cog

Engineering a full humanoid robot is a major task and has taken a number of years. Recently, however, we have started to be able to make interesting demonstrations on the integrated hardware. There have been a number of Master theses, namely Irie (1995), Marjanović (1995), Matsuoka (1995a), and Williamson (1995), and a number of more recent papers Matsuoka (1995b), Ferrell (1996), Marjanović, Scassellati & Williamson (1996), and Williamson (1996), describing various aspects of the demonstrations with Cog.

Irie has built a sound localization system that uses cues very similar to those used by humans (phase difference below 1.5Khz, and time of arrival above that frequency) to determine the direction of sounds. He then used a set of neural networks to learn how to correlate particular sorts of sounds, and their apparent aural location with their visual location using the assumption that aural events are often accompanied by motion. This learning of the correlation between aural and visual sensing, along with a coupling into the oculomotor map is akin to the process that goes on in the superior colliculus.

Marjanović started out with a simple model of the cerebellum that over a period of eight hours of training was able to compensate for the visual slip induced in the eyes by neck motions.

Matsuoka built a three fingered, one thumbed, self contained hand (not currently mounted on Cog). It is fully self contained and uses frictionally coupled tendons to automatically accommodate the shape of grasped objects. The hand is covered with touch sensors. A combination of three different sorts of neural networks is used to control the hand, although more work remains to be done to make them all incremental and on-line.

Williamson controls the robot arm using a biological model (from Bizzi, Giszter, Loeb, Mussa-Ivaldi & Saltiel (1995) based on postural primitives. The arm is completely compliant and is safe for people to interact with. People can push the arm out of the way, just as they could with a child if they chose to.

The more recent work with Cog has concentrated on component behaviors that will be necessary for Cog to orient, using sound localization to a noisy and moving object, and then bat at the visually localized object. Ferrell has developed two dimensional topographic map structures which let Cog learn mappings between peripheral and foveal image coordinates and relate them to oculomotor coordinates. Marjanović, Scassellati and Williamson, have used similar sorts of maps to relate hand coordinates to eye coordinates and can learn how to reach to a visual target.

These basic capabilities will form the basis for higher level learning and development of Cog.

4 Conclusion

We have discussed shifts in viewpoints on how to organize intelligence and their immediate impact on how we might build robots with human level intelligence. Let us assume that these shifts are in the right direction, and will not immediately be shown to be totally wrong. Perhaps the best we could hope for, however, is only that they are in approximately the right direction, and that over the next few years there will be many refinements and adjustments to these particular approaches.

But, there is a more troublesome possibility. Perhaps it is the case that all the approaches to building intelligent systems are just completely off-base, and are doomed to fail. Why should we worry that this is so? Well, certainly it is the case that all biological systems are:

- Much more robust to changed circumstances than our artificial systems.
- Much quicker to learn or adapt than any of our machine learning algorithms¹¹.

¹¹ The very term *machine learning* is unfortunately synonymous with a pernicious form of totally impractical but theoretically sound and elegant classes of algorithms.

- Behave in a way which just simply seems *life-like* in a way that our robots never do.

Perhaps we have all missed some organizing principle of biological systems, or some general truth about them. Perhaps there is a way of looking at biological systems which will illuminate an inherent necessity in some aspect of the interactions of their parts that is completely missing from our artificial systems. This could be interpreted to be an *elixir of life* or some such, but I am not suggesting that we need go outside the current realms of mathematics, physics, chemistry, or bio-chemistry. Rather I am suggesting that perhaps at this point we simply do not *get it*, and that there is some fundamental change necessary in our thinking in order that we might build artificial systems that have the levels of intelligence, emotional interactions, long term stability and autonomy, and general robustness that we might expect of biological systems. In deference to the elixir metaphor, I prefer to think that perhaps we are currently missing the *juice* of life.

Final note. Why the title of this paper? “From Earwigs to Humans”. It comes from a paper written by Kirsh (1991), titled “Today the earwig, tomorrow man?”, written as a response to long delayed publication of my first manifesto on behavior-based robotics (Brooks 1991*b*). The title was meant in a somewhat contemptuous spirit, arguing that behavior-based approaches, while perhaps adequate for insect-level behavior, could never scale to human-level behavior. The Cog project, and in a little way, this paper, are my response. Or, more precisely “Yes, exactly!”.

Acknowledgements

The bulk of the engineering on Cog has been carried out by graduate students Cynthia Ferrell, Matthew Williamson, Matthew Marjanović, Brian Scassellati, Robert Irie, and Yoky Matsuoka. Earlier Michael Binard contributed to the body and neck, and Eleni Kapogiannis built the first version of the brain. They have all labored long and hard. Various pieces of software have been contributed by Mike Wessler, Lynn Stein, and Joanna Bryson. A number of undergraduates have been involved, including Elmer Lee, Erin Panttaja, Jonah Peshkin, and Milton Wong. James McLurkin and Nicholas Schectman have contributed engineering assistance to the project. Rene Schaad recently built a vestibular system for Cog.

References

Agre, P. E. & Chapman, D. (1987), Pengi: An Implementation of a Theory of Activity, in ‘Proceedings

of the Sixth Annual Meeting of the American Association for Artificial Intelligence’, Morgan Kaufmann Publishers, Seattle, Washington, pp. 268–272. Also appeared in (Luger 1995).

Arkin, R. C. (1992), ‘Behavior-Based Robot Navigation for Extended Domains’, *Adaptive Behavior* **1**(2), 201–225.

Ashby, W. R. (1952), *Design for a Brain*, Chapman and Hall, London, United Kingdom.

Baird III, L. C. & Klopff, A. H. (1993), ‘A Hierarchical Network of Provably Optimal Learning Control Systems: Extensions of the Associative Control Process (ACP) Network’, *Adaptive Behavior* **1**(3), 321–352.

Beer, R. D. (1995*a*), ‘A Dynamical Systems Perspective on Agent-Environment Interaction’, *Artificial Intelligence Journal* **72**(1/2), 173–215.

Beer, R. D. (1995*b*), ‘On the Dynamics of Small Continuous-Time Recurrent Neural Networks’, *Adaptive Behavior* **3**(4), 469–509.

Beer, R. D. & Gallagher, J. C. (1992), ‘Evolving Dynamical Neural Networks for Adaptive Behavior’, *Adaptive Behavior* **1**(1), 91–122.

Bizzi, E., Giszter, S. F., Loeb, E., Mussa-Ivaldi, F. A. & Saltiel, P. (1995), ‘Modular Organization of Motor Behavior in the Frog’s Spinal Cord’, *Trends in Neurosciences* **18**, 442–446.

Brooks, R. A. (1986*a*), Achieving Artificial Intelligence Through Building Robots, Memo 899, Massachusetts Institute of Technology, Artificial Intelligence Lab, Cambridge, Massachusetts.

Brooks, R. A. (1986*b*), ‘A Robust Layered Control System for a Mobile Robot’, *IEEE Journal of Robotics and Automation* **RA-2**, 14–23.

Brooks, R. A. (1987), Intelligence without representation, in ‘Proceedings of the Workshop on the Foundations of AI’, Endicott House.

Brooks, R. A. (1989), ‘A Robot That Walks: Emergent Behavior from a Carefully Evolved Network’, *Neural Computation* **1**(2), 253–262.

Brooks, R. A. (1990), Elephants Don’t Play Chess, in P. Maes, ed., ‘Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back’, MIT Press, Cambridge, Massachusetts, pp. 3–15.

- Brooks, R. A. (1991a), Intelligence Without Reason, in 'Proceedings of the 1991 International Joint Conference on Artificial Intelligence', pp. 569-595. Also appeared in (Steels & Brooks 1995).
- Brooks, R. A. (1991b), 'Intelligence Without Representation', *Artificial Intelligence Journal* **47**, 139-160. Also appeared in (Luger 1995).
- Brooks, R. A. (1991c), 'New Approaches to Robotics', *Science* **253**, 1227-1232.
- Brooks, R. A. & Maes, P., eds (1994), *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, MIT Press, Cambridge, Massachusetts.
- Brooks, R. A. & Stein, L. A. (1993), Building Brains for Bodies, Memo 1439, Massachusetts Institute of Technology, Artificial Intelligence Lab, Cambridge, Massachusetts.
- Cliff, D. & Bullock, S. (1993), 'Adding "Foveal Vision" to Wilson's Animat', *Adaptive Behavior* **2**(1), 49-72.
- Cliff, D. & Ross, S. (1994), 'Adding Temporary Memory to ZCS', *Adaptive Behavior* **3**(2), 101-150.
- Cliff, D., Harvey, I. & Husbands, P. (1993), 'Explorations in Evolutionary Robotics', *Adaptive Behavior* **2**(1), 73-110.
- Cliff, D., Husbands, P., Meyer, J.-A. & Wilson, S. W., eds (1994), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, MIT Press, Cambridge, Massachusetts.
- Colombetti, M. & Dorigo, M. (1994), 'Training Agents to Perform Sequential Behavior', *Adaptive Behavior* **2**(3), 247-275.
- Corbacho, F. J. & Arbib, M. A. (1995), 'Learning to Detour', *Adaptive Behavior* **3**(4), 419-468.
- Cruse, H., Brunn, D., Bartling, C., Dean, J., Dreifert, M., Kindermann, T. & Schmitz, J. (1995), 'Walking: A Complex Behavior Controlled by Simple Networks', *Adaptive Behavior* **3**(4), 385-418.
- Ekeberg, Ö., Lansner, A. & Grillner, S. (1995), 'The Neural Control of Fish Swimming Studied Through Numerical Simulations', *Adaptive Behavior* **3**(4), 363-384.
- Espenschied, K. S., Quinn, R. D., Chiel, H. J. & Beer, R. D. (1993), 'Leg Coordination Mechanisms in the Stick Insect Applied to Hexapod Robot Locomotion', *Adaptive Behavior* **1**(4), 455-468.
- Ewert, J.-P., Beneke, T., Buxbaum-Conradi, H., Dinges, A., Fingerling, S., Glasgow, M., Schürg-Pfeiffer, E. & Schwippert, W. (1992), 'Adapted and Adaptive Properties in Neural Networks for Visual Pattern Discrimination: A Neurobiological Analysis Toward Neural Engineering', *Adaptive Behavior* **1**(2), 123-154.
- Fagg, A. H. & Arbib, M. A. (1992), 'A Model of Primate Visual-Motor Conditional Learning', *Adaptive Behavior* **1**(1), 3-37.
- Ferrell, C. (1994), 'Failure Recognition and Fault Tolerance of an Autonomous Robot', *Adaptive Behavior* **2**(4), 375-398.
- Ferrell, C. (1995a), 'A Comparison of Three Insect-Inspired Locomotion Controllers', *Robotics and Autonomous Systems* **16**(2-4), 135-159.
- Ferrell, C. (1995b), 'Global Behavior via Cooperative Local Control', *Autonomous Robots* **2**(2), 105-125.
- Ferrell, C. (1996), Orientation Behavior Using Registered Topographic Maps, in 'Fourth International Conference on Simulation of Adaptive Behavior', Cape Cod, Massachusetts. To appear.
- Grefenstette, J. J. (1992), 'The Evolution of Strategies for Multiagent Environments', *Adaptive Behavior* **1**(1), 65-90.
- Gruau, F. (1994), 'Automatic Definition of Modular Neural Networks', *Adaptive Behavior* **3**(2), 151-183.
- Hallam, B. E., Halperin, J. R. & Hallam, J. C. (1994), 'An Ethological Model for Implementation in Mobile Robots', *Adaptive Behavior* **3**(1), 51-79.
- Halperin, J. R. P. & Dunham, D. W. (1992), 'Postponed Conditioning: Testing a Hypothesis About Synaptic Strengthening', *Adaptive Behavior* **1**(1), 39-63.
- Horswill, I. D. & Brooks, R. A. (1988), Situated Vision in a Dynamic World: Chasing Objects, in 'Proceedings of the Seventh Annual Meeting of the American Association for Artificial Intelligence', St. Paul, Minnesota, pp. 796-800.
- Irie, R. E. (1995), Robust Sound Localization: An Application of an Auditory Perception System for a Humanoid Robot, Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Johnson, M. (1987), *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*, The University of Chicago Press, Chicago, Illinois.

- Kehoe, E. J., Horne, A. J. & Macrae, M. (1995), 'Learning to Learn: Features and a Connectionist Model', *Adaptive Behavior* **3**(3), 235–271.
- Kirsh, D. (1991), 'Today the earwig, tomorrow man?', *Artificial Intelligence Journal* **47**, 161–184.
- Klopf, A. H., Morgan, J. S. & Weaver, S. E. (1993), 'A Hierarchical Network of Control Systems that Learn: Modeling Nervous System Function During Classical and Instrumental Conditioning', *Adaptive Behavior* **1**(3), 263–319.
- Koza, J. R., Rice, J. P. & Roughgarden, J. (1992), 'Evolution of Food-Foraging Strategies for the Caribbean *Anolis* Lizard Using Genetic Programming', *Adaptive Behavior* **1**(2), 171–199.
- Kube, C. R. & Zhang, H. (1993), 'Collective Robotics: From Social Insects to Robotics', *Adaptive Behavior* **2**(2), 189–218.
- Lakoff, G. (1987), *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*, University of Chicago Press, Chicago, Illinois.
- Liaw, J.-S. & Arbib, M. A. (1993), 'Neural Mechanisms Underlying Direction-Selective Avoidance Behavior', *Adaptive Behavior* **1**(3), 227–261.
- Ling, C. X. & Buchal, R. (1994), 'Learning to Control Dynamic Systems with Automatic Quantization', *Adaptive Behavior* **3**(1), 29–49.
- Luger, G. F., ed. (1995), *Computation & Intelligence: Collected Readings*, AAAI Press, MIT Press, Menlo Park, California.
- MacLennan, B. & Burghardt, G. M. (1993), 'Synthetic Ethology and the Evolution of Cooperative Communication', *Adaptive Behavior* **2**(2), 161–188.
- Maes, P. (1990a), Situated Agents Can Have Goals, in P. Maes, ed., 'Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back', MIT Press, Cambridge, Massachusetts, pp. 49–70. Also published in the special issue of the *Journal of Robotics and Autonomous Systems*, Spring '90, North Holland.
- Maes, P., ed. (1990b), *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, MIT Press, Cambridge, Massachusetts.
- Marjanović, M. J. (1995), Learning Maps Between Sensorimotor Systems on a Humanoid Robot, Master's thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Marjanović, M., Scassellatti, B. & Williamson, M. (1996), Self-Taught Visually-Guided Pointing for a Humanoid Robot, in 'Fourth International Conference on Simulation of Adaptive Behavior', Cape Cod, Massachusetts. To appear.
- Matsuoka, Y. (1995a), Embodiment and Manipulation Learning Process for a Humanoid Hand, Technical Report TR-1546, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Matsuoka, Y. (1995b), Primitive Manipulation Learning with Connectionism, in M. C. D.S. Touretzky & M. Hasselmo, eds, 'Advances in Neural Information Processing Systems 8', MIT Press, Cambridge, Massachusetts, pp. 889–895.
- Meyer, J.-A. & Wilson, S. W., eds (1991), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior*, MIT Press, Cambridge, Massachusetts.
- Meyer, J.-A., Roitblat, H. L. & Wilson, S. W., eds (1993), *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, MIT Press, Cambridge, Massachusetts.
- Morán, F., Moreno, A., Merelo, J. & Chacón, P., eds (1995), *Advances in Artificial Life: Third European Conference on Artificial Life*, Springer Verlag, Berlin, Germany.
- Moravec, H. P. (1984), Locomotion, Vision and Intelligence, in Brady & Paul, eds, 'Robotics Research 1', MIT Press, Cambridge, Massachusetts, pp. 215–224.
- Müller-Wilm, U., Dean, J., Cruse, H., Weidemann, H., Eltze, J. & Pfeiffer, F. (1992), 'Kinematic Model of a Stick Insect as an Example of a Six-Legged Walking System', *Adaptive Behavior* **1**(2), 155–169.
- Nilsson, N. (1969), A Mobile Automation: An Application of Artificial Intelligence Techniques, in 'Proceedings of the 1st International Joint Conference on Artificial Intelligence', Washington, D.C., pp. 509–520.
- Nilsson, N. J., ed. (1984), *Shakey the Robot*, Stanford Research Institute AI Center, Technical Note 323.
- Nolfi, S., Elman, J. L. & Parisi, D. (1994), 'Learning and Evolution in Neural Networks', *Adaptive Behavior* **3**(1), 5–28.
- Numaoka, C. (1994), 'Phase Transitions in Instigated Collective Decision Making', *Adaptive Behavior* **3**(2), 185–223.

- Peng, J. & Williams, R. J. (1993), 'Efficient Learning and Planning Within the Dyna Framework', *Adaptive Behavior* **1**(4), 437-454.
- Pfeiffer, R. (1996), Building "Fungu Eaters": Design Principles of Autonomous Agents, in 'Fourth International Conference on Simulation of Adaptive Behavior', Cape Cod, Massachusetts. To appear.
- Ram, A., Arkin, R., Boone, G. & Pearce, M. (1994), 'Using Genetic Algorithms to Learn Reactive Control Parameters for Autonomous Robotic Navigation', *Adaptive Behavior* **2**(3), 277-305.
- Rosenschein, S. J. & Kaelbling, L. P. (1986), The Synthesis of Machines with Provable Epistemic Properties, in J. Halpern, ed., 'Proc. Conf. on Theoretical Aspects of Reasoning about Knowledge', Morgan Kaufmann Publishers, Los Altos, California, pp. 83-98.
- Rutkowska, J. C. (1994), 'Scaling Up Sensorimotor Systems: Constraints from Human Infancy', *Adaptive Behavior* **2**(4), 349-373.
- Schmajuk, N. A. & Blair, H. T. (1993), 'Place Learning and the Dynamics of Spatial Navigation: A Neural Network Approach', *Adaptive Behavior* **1**(3), 353-385.
- Schölkopf, B. & Mallot, H. A. (1995), 'View-Based Cognitive Mapping and Path Planning', *Adaptive Behavior* **3**(3), 311-348.
- Smithers, T. (1994), What the Dynamics of Adaptive Behaviour and Cognition Might Look Like in Agent-Environment Interaction Systems, in 'DRABC: On the Role of Dynamics and Representation in Adaptive Behaviour and Cognition', Euskal Herriko Unibertsitatea, Universidad del País Vasco, Spain.
- Staddon, J. E. (1993), 'On Rate-Sensitive Habituation', *Adaptive Behavior* **1**(4), 421-436.
- Steels, L. & Brooks, R., eds (1995), *The Artificial Life Route to Artificial Intelligence: Building Embodied, Situated Agents*, Lawrence Erlbaum Associates, Hillsdale, New Jersey.
- Steinhage, A. & Schoener, G. (1996), 'Self-Calibration Based on Invariant View Recognition: Dynamic approach to Navigation', *Robotics and Autonomous Systems*. To appear.
- Szepesvári, C. & Lörincz, A. (1993), 'Behavior of an Adaptive Self-organizing Autonomous Agent Working with Cues and Competing Concepts', *Adaptive Behavior* **2**(2), 131-160.
- Todd, P. M. & Miller, G. F. (1992), 'Parental Guidance Suggested: How Parental Imprinting Evolves Through Sexual Selection as an Adaptive Learning Mechanism', *Adaptive Behavior* **2**(1), 5-47.
- Tyrrell, T. (1993), 'The Use of Hierarchies for Action Selection', *Adaptive Behavior* **1**(4), 387-420.
- Tyrrell, T. (1994), 'An Evaluation of Maes's Bottom-Up Mechanism for Behavior Selection', *Adaptive Behavior* **2**(4), 307-348.
- Walter, W. G. (1950), 'An Imitation of Life', *Scientific American* **182**(5), 42-45.
- Walter, W. G. (1951), 'A Machine That Learns', *Scientific American* **185**(5), 60-63.
- Wang, D. (1993), 'A Neural Model of Synaptic Plasticity Underlying Short-term and Long-term Habituation', *Adaptive Behavior* **2**(2), 111-129.
- Wheeler, M. & de Bourcier, P. (1995), 'How Not to Murder Your Neighbor: Using Synthetic Behavioral Ecology to Study Aggressive Signaling', *Adaptive Behavior* **3**(3), 273-309.
- Williamson, M. (1996), Postural Primitives: Interactive Behavior for a Humanoid Robot Arm, in 'Fourth International Conference on Simulation of Adaptive Behavior', Cape Cod, Massachusetts. To appear.
- Williamson, M. M. (1995), Series Elastic Actuators, Technical Report 1524, Massachusetts Institute of Technology, Cambridge, Massachusetts.
- Wilson, S. W. (1985), Knowledge Growth in an Artificial Animal, in J. J. Grefenstette, ed., 'Proc. International Conference on Genetic Algorithms and Their Applications, Pittsburgh, PA', Erlbaum, Hillsdale, NJ.
- Yamauchi, B. M. & Beer, R. D. (1994), 'Sequential Behavior and Learning in Evolved Dynamical Neural Networks', *Adaptive Behavior* **2**(3), 219-246.