ANIMAL BEHAVIOUR, 2000, **60**, 545–558
doi:10.1006/anbe.2000.1514, available online at http://www.idealibrary.com on IDEAL®

REVIEW

What does robotics offer animal behaviour?

BARBARA WEBB
Centre for Cognitive and Computational Neuroscience, Department of Psychology, University of Stirling

*(Received 15 February 2000; initial acceptance 24 March 2000;
final acceptance 12 June 2000; MS. number: RV-19)*

There is a growing body of robot-based research that makes a serious claim to be a new methodology for biology. Robots can be used as models of specific animal systems to test hypotheses regarding the control of behaviour. At levels from learning algorithms to specific dendritic circuits, implementing a proposed controller in a robotic device tests it against real environments in a way that is difficult to simulate. This often provides insight into the true nature of the problem. It also enforces complete specifications and combines bodies of data. Current work can sometimes be criticized for drawing unjustified conclusions given the limited evaluation and inevitable inaccuracies of robot models. Nevertheless, this approach has led to novel hypotheses for animal behaviour and seems likely to provide fruitful results in the future.

We have recently seen robots that use 'bee' vision to navigate down corridors and avoid obstacles; a robot 'cricket' that tracks down a male cricket by its song; a 'robolobster' that follows an underwater chemical plume to its source; and a group of robots that construct an 'ant nest' wall. Such examples of engineering can be attention grabbing, but what is their value for biological science? In particular, beyond the 'gimmick' of resemblance to natural systems, is any deeper understanding of how animals behave brought about by the building of such robot systems?

Robots are machines built to perform tasks using actions that are based on, or reminiscent of, humans or other animals. Some features distinguishing these machines are: direct actuation, that is, they have motor devices that allow them to move about in and manipulate their environment independently of human intervention; direct sensing, that is, not just a highly constrained push button or keyboard interface; 'intelligent' control, that is, their actuation is goal-oriented with respect to the variations in the task environment; and 'autonomy', that is, once set in motion, the behaviour is self-determined rather than remote controlled. Clearly the problems roboticists have to solve correspond closely to the problems evolution has solved for real animals.

Attempts to make machines behave in a life-like manner are as old as science. Ingenious mechanical devices have been built to mimic animal behaviours, sometimes with impressive detail, but their clockwork mechanisms did not noticeably resemble the inner workings of biological systems. Fifty years ago the advent of cybernetics saw the building of a series of electromechanical devices more seriously intended to explore aspects of animal behaviour, such as the 'homeostat' machine (Ashby 1952) and the reactive 'turtle' (Walter 1961). However, since then robotics as a research field has been largely dominated by approaches developed in control systems theory, such as methods to find inverse kinematics or establish robust feedback control. As yet there is no general theoretical solution to the problem of building a robot to perform a specific behaviour, but rather a diverse range of mechanisms that can be adapted to specific tasks. It is apparent that if we knew how animals controlled their behaviour this might give us ideas about how to make robots do it (although these may not turn out to be the best methods to adopt from an engineering perspective). In fact, the title of this paper inverts the question posed in a recent robotics textbook 'What does animal behaviour offer robotics?' (Arkin 1998, page 52). Biology can be viewed as a source of existence proofs for what capabilities might be possible for robots, and of ideas for mechanisms for achieving these capabilities.

However, in the present paper, I am not concerned to debate the potential benefits or otherwise of biological inspiration for robot engineering, but rather to explore how robot-based research might be a new methodology for biology. This involves using robots as biological

Correspondence: B. Webb, Department of Psychology, University of Stirling, Stirling FK9 4LA, U.K. (email: b.h.webb@stir.ac.uk).

models: 'robots can be used as physical models of animals to address specific biological questions' (Beer et al. 1998, page 777); 'the goal of this approach is to develop an understanding of natural systems by building a robot that mimics some aspects of their sensory and nervous system and their behaviour' (Lambrinos et al. 1997, page 185); 'hypotheses can . . . be verified or falsified by building an adequate robot, implementing the hypothesized control mechanism, and then observing if it can account for the behaviour' (Lund & Asada 1998, page 303). However as Delcomyn & Nelson (2000, pp. 5–6) have noted, compared to the number of engineers adopting biological ideas in their robot building 'neurobiologists have been slower in establishing collaborations with engineers for the converse purpose, to use robots as models to study biological mechanisms'. Lund & Asada (1998, page 303) further suggested 'the majority of biologists might not accept this methodology . . . the debate . . . concerns the reliability of a robot as the model of the animal under study'.

From this perspective, robot building has the same general steps one would expect in any model building: identifying a target and issue to be explained; offering an explanation; demonstrating that it accounts for observations; deriving further predictions and testing them. Indeed, many of the issues to be discussed, such as underdetermination, accuracy of representation and evaluation, are relevant to other kinds of modelling, such as mathematical or computational approaches. However, there are also unique issues for this particular form of modelling, not always appreciated by biologists when they make claims such as 'a two-eared robot programmed to obey this rule (if suitable noise were incorporated) could be made to track a sound source in a manner like that of the female [cricket]' (Weber & Thorson 1988, page 14).

Building such a robot requires the solution of a large range of problems. The basic body plan and actuation mechanism for the robot have to be determined, for example, how fast must it move, how easily must it manoeuvre? Transduction devices that resemble the animal's sensors have to be found or designed. The exact details of the wiring between the sensors and actuators, and the precise programming specification of the behavioural 'rule' has to be determined. 'Suitable noise' (whatever that may be) may have to be incorporated. If a functional device is produced, the behaviour has to be tested under a range of conditions, and means of collecting data devised. Finally, appropriate ways to compare the robot's behaviour with that of the animal have to be found: given the inevitable differences between them, how similar must the behaviour be to support the original claim? I hope to illustrate in this review how solving some of these problems can help us better understand the behaviour of animals.

Some collections of papers in this area can be found in recent special issues of *Connection Science* (**10**, part 3) and *Robotics and Autonomous Systems* (**18**, part 1 and **30**, parts 1–2). Some discussions of this approach include: for modelling motor control (Beer et al. 1998); in relation to more general 'animat' methods (Dean 1998); and in relation to robot autonomy (Sharkey & Ziemke 1998). In this paper, I focus on work in which a physical robot has been built and tested, and explicitly used to demonstrate understanding or explore hypotheses of animal behaviour. This is not to denigrate biological computer simulations, or robotic investigation for its own sake, but simply to delineate a distinctive area of current research where biology and robotics intersect. The first restriction means that hardware as well as software issues have been involved in building the model and interaction with a real not just simulated world has been examined. The second means that the results of the robot building is intended to produce information of relevance to biologists, not just to engineers.

I refer to this research as 'biorobotics', although the terminology is currently rather mutable, and I argue that indeed such robots are useful for biology. This is partly for the reasons that any implemented model is useful: it enforces clarity and full specification of hypotheses; and provides a means to derive the nonobvious consequences of complex sets of assumptions. But robot models also have particular strengths for understanding animal behaviour. They can test hypotheses of control (at levels from algorithms to subneural properties) under realistic bodily and environmental conditions, thus integrating diverse data. In doing so they can be very effective in helping to characterize the problem that the animal needs to solve. Consequently, they can be productive of new hypotheses, predictions or suggested experiments for biological systems. At the same time there are limitations that should be kept in mind. In particular, the strength of the conclusions that can be drawn from this research depends on the accuracy of representation of the animal and its behaviour, and this is frequently not sufficiently well explained or assessed. Both these strengths and weaknesses are discussed in detail below.

TESTING HYPOTHESES

Robots as models are a means by which hypotheses can be tested for adequacy and sufficiency to explain a set of data, and additional predictions from the hypotheses can be derived. Robots can be useful in exploring hypotheses about the mechanisms of behaviour particularly when behaviour is considered as the outcome of the interaction of an environment, a body and a control system. Computer-simulated models tend to focus investigation on the control system in isolation, or use highly idealized input and output. In robotic modelling, a proposed control system is tested for whether it really produces the behaviour of interest when placed under comparable bodily and environmental conditions to the animal. Hypotheses can be investigated at a variety of different levels, from functional to subneural properties, as the following examples illustrate.

Algorithms

The essential functions performed by the neural system can be explored at an algorithmic level. For example, Srinivasan et al. (1999) have explored what basic rules or cues for visual navigation could be used by the bee, both in experiments on the animal and by showing their efficacy in robot control. An example is using the balance of optical flow on each side to negotiate a corridor of variable direction and width. A mobile robot equipped with a camera and able to derive an estimate of optical flow can use this to maintain a constant distance from walls or ground surfaces (Santos-Victor et al. 1995; Weber et al. 1998). Santos-Victor et al. (1995) noted how running their robot in a real environment raised some novel issues, for example difficulties in using direct reflexes based on optical flow when natural environments can contain patches of untextured surface that provide no information.

Srinivasan's group has taken a similar algorithmic approach to understanding how the peering behaviour of the locust (Sobel 1990) can serve as a range-finding mechanism (Srinivasan et al. 1999). This peering behaviour has also been implemented in a neurocomputational model on a robot by Lewis & Nelson, using nondirectional selective motion detectors (Douglass & Strausfeld 1996) and producing a 'jump' parameter estimate of object distance from real image inputs as a robot moves back and forth in front of a target (Lewis & Nelson 1998). Other examples of 'algorithm' level research include investigations of navigation (e.g. Lambrinos et al. 1997), collective behaviour (e.g. Holland & Melhuish 1999) and motor control (e.g. Raibert 1986) discussed further below.

Learning Mechanisms

Another group of studies carried out mostly at the algorithmic level is the use of learning mechanisms to train robot behaviour. While most uses of 'learning algorithms' in robotics make few explicit connections to animal learning (e.g. biological parallels are barely mentioned in articles in the special issue 'Reinforcement Learning' of *Robotics and Autonomous Systems*, **15**, part 4), there are exceptions. An example is the implementation of operant conditioning by Saksida et al. (1997) which includes conditioned reinforcers and shifting reinforcement contingencies. They tested the model using the delayed match-to-sample paradigm often used to test working memory in rats. Verschure (1998) has investigated classical conditioning issues such as secondary conditioning and blocking in robot models. Chang & Gaudiano (1998) implemented Grossberg's (1982) model of associative learning and demonstrated its flexible application across different robot devices and environments. Hallam et al. (1994) implemented and tested on a robot a biologically derived model of associative learning in Siamese fighting fish, *Betta splendens*. The last three studies all introduced 'neural network' descriptions of the mechanisms and used variant forms of Hebbian changes of 'connection weights'. Damper & Scutt used more explicit models of neural habituation and sensitization for robot learning (Scutt & Damper 1997; Damper et al., in press), to test the suggestion that the same mechanisms underlie associative learning, as suggested by Hawkins & Kandel (1984).

Savage (1998) reviewed and criticized some of these learning models as not sufficiently sophisticated to address current issues in learning theory. It is true that discussions of learning (and other animal behaviours) in robotics can often seem 'biologically naïve'. It is important that robot researchers wishing to claim relevance of their results for biology make a serious effort to understand the current state of knowledge and to engage appropriate problems. However, this criticism is missing the point if the objection is only that robots fail to reproduce the full diversity of observed animal behaviour. The point of this research is to discover what behaviour does emerge from the algorithms as currently specified, and therefore to test whether they can account for observed animal behaviour; insofar as they fail to do so, it may be the theory that needs revision. Indeed, one use of robot learning studies can be, irrespective of the precise learning mechanism, to explore whether certain environmental contingencies are sufficient for explaining certain animal behaviours; we can be sure what is or is not 'innate' in our robot. An example is the use by S. Nolfi & H. H. Lund (unpublished data) of an evolutionarily adaptive approach to show that a robot (and by implication a rat) does not need a 'cognitive map' to learn to perform oriented behaviour in the box task used by Cheng (1986) and Margules & Gallistel (1988). The conclusion drawn is that sufficient richness of environmental interaction can account for the behaviour without additional internal mechanisms.

Evolution

The idea of testing the power of adaptive mechanisms by using robots in real environments was discussed by Nolfi (1998, page 169) in relation to the use of evolutionary algorithms. As he noted 'only information . . . truly available in the environment can be used for training' or evolving the behaviour of the robot. Evolutionary algorithms implement processes akin to natural selection: that is, random variations in the robot result in differential success on some behavioural task, which is used to preserve differentially the variations in subsequent implementations. This can be an effective 'blind' method for improving the adaptation of the robot to its task and environment, confirming the basic principle of natural selection as a mechanism for evolution. However, it is often not trivial to make this method effective: it may fail to improve the robot in the way desired (Mataric & Cliff 1996). The reasons for difficulties may be revealing about the process of natural selection, for example that the discriminations made by certain sensory devices in certain environments may not be the obvious ones that we as observers assume are available to shape the robot's (or animal's) performance.

It has been suggested that evolutionary robotics might go beyond generically exploring natural selection to test more explicit hypotheses: for example, that an evolutionary simulation based on a real creature may be a way to derive a possible network model for that behaviour (Beer & Gallagher 1992; Cliff et al. 1993). Most evolutionary robotics to date is not, in fact, closely based on specific animals. One exception is Kortmann & Hallam's (1999) investigation of cricket song preference which sought to explore the 'evolutionary plausibility' of a specific neural model for the behaviour. Nolfi (1998) questioned the viability of this approach and argued that robot evolution is better used either to study more general principles of evolutionary theory, such as how predator–prey interactions may lead to increased sensory motor complexity (Nolfi & Floreano 1998), or to develop counterexamples to assumptions about 'necessary' mechanisms in animals by evolving robots capable of comparable behaviour without those mechanisms, as described above.

Neuroanatomy

Rucci et al. (2000) used a somewhat similar approach, a learning mechanism functioning under natural conditions of real stimuli and real sensorimotor responses, to look at azimuth localization of sound sources by a robot head. In this study the system was also constrained by and compared with data on the neural structures subserving this capability in the barn owl, *Tyto alba* (Konishi 1993). Rucci et al. (2000) noted that 'the use of robotic systems has the advantage of introducing phenotypic and environmental constraints similar to those that brains of animals have to face during development . . . particularly important in the light of modern brain theories, that emphasise the importance of the environment and sensorimotor experience during neural development'.

Also at the level of brain neuroanatomy, Arbib & Liaw (1995) have produced both simulation and robot models of the functioning of the tectal region of the frog brain (Ewert 1987), for example in the response to visual looming, expressed in terms of 'schemas'. Arbib & Liaw (1995) described how their robot raised previously unconsidered problems of handling object occlusion and detecting gap size when trying to guide a real device around a real environment. They argued that what makes such research a specifically biological model is when 'explicit hypotheses are offered as to how the constituent schemas are played over particular regions of the brain' (Arbib & Law 1995, page 56). In the context of robot models of hippocampal function, discussed further below, Gaussier et al. (2000, page 175) noted 'our level of modelization cannot account today for the neurobiological details of each brain structure' but can allow the researcher 'to test for the coherence of global brain models and to verify if two functional boxes can really be connected'.

Specific Neural Circuits

At the level of specific neuron function and circuitry, good examples are provided by several robotic investigations of the motion-sensitive circuitry of the fly and its role in optomotor, obstacle-avoiding and visual fixation responses. The first was a hardware model of the motion selectivity of the H1 neuron of the fly (Franceschini et al. 1992). As well as confirming that this mechanism sufficed for visually guided avoidance of obstacles (Pichon et al. 1989), it led to novel experiments on the fly such as the discovery of a mechanism of 'binocular vergence' (Franceschini 1996). The H1 neuron has also been modelled in analogue VLSI (Very Large Scale Integration) technology to produce an optomotor 'chip' (Harrison & Koch 1999). The output of this device has been used to correct to straight-line motion the path of a robot with asymmetric gearing. It has also been directly wired into the experimental equipment used to test the fly (Warzecha & Egelhaaf 1996) to demonstrate that the implemented circuitry suffices to reproduce the fly's behaviour closely, both correcting for imposed rotation and producing microoscillations.

In another interesting model of the fly's motion detection system, Huber et al. (Huber & Bulthoff 1998; Huber et al. 1999) have used a conical mirror to obtain a coarse 360° image comparable to the fly's visual input on the horizontal axis (similar technology has also been used in work by Srinivasan et al. 1999 and Lambrinos et al. 2000). This is processed using a model of the spatial and temporal filtering of the fly's retina and lamina and Reichardt-style motion detectors. By copying the asymmetric properties of lobula plate HSE neurons to large-field progressive and regressive motion (Egelhaaf et al. 1989) a control signal for movement by the robot is derived. Huber et al. (1999) have shown that the same sensory circuit can then produce both optomotor and fixation behaviour comparable to that of the fly depending on the environmental cues, and concluded that the large field cells in the fly participate in both responses. They have thus explicitly shown how 'the process of designing the sensorimotor control of a robot . . . can provide the basis of a critical evaluation of biological models' (Huber et al. 1999, page 227).

Dendritic Circuitry

An example of robotic investigations of the behavioural capabilities consequent on the dendritic structure of specific neurons is the implementation of the locust looming detector by Blanchard et al. (1999, 2000). Studies of the locust LGMD neuron have suggested that a 'race' between inhibitory and excitatory inputs tune it to respond only to rapidly moving edges (Rind & Bramwell 1996). On a robot this model has been used to detect approaching objects and avoid collisions. An interesting finding was that, with relatively natural scenes as input, the response of the model does not predictably encode the rate of approach in the spike rate, as had been taken to be the case from studies of the real LGMD obtained with simpler stimuli (Rind & Simmons 1997). This has led to proposals for further experiments on the animal using more complex stimuli.

Robots in Animal Experiments

Another somewhat different role for robots in testing biological hypotheses is the use of robot devices to carry out behavioural experiments on animals. That is, the robot can be used to interact with the animals to test particular predictions of animal behaviour. One example is the electromechanical 'bee' used by Michelsen et al. (1992) to study critical cues in the communication of the location of nectar sources by honeybees, *Apis mellifera*. A more recent example is discussed in Takanishi et al. (1998a) where a remote-controlled electromechanical rat was used to look at imitation behaviour in real rodents. Both these devices are under fairly direct experimenter control, but a more autonomous device has been developed as a robot 'sheepdog' whose interaction with groups of ducks tests models of flocking behaviour in these animals (Vaughan et al. 2000).

CHARACTERIZING THE PROBLEM AND UNDERSTANDING THE ENVIRONMENT

An advantage of real robot implementation is that the researcher is forced to confront assumptions about the nature of the stimulus and the possible actions given real characteristics of the environment. What is thereby brought to the fore is the critical importance of environmental and bodily factors in explaining behavioural capabilities. The significance of these considerations for a variety of biological examples was reviewed in Chiel & Beer (1997). Flynn & Brooks (1989) have discussed this point in relation to simulation within robotics research, noting how building systems in the real world often shows one's prior intuitions about requirements to be completely wrong. As Grasso (in press) argued: 'The bio-robot offers a way of systematically evaluating the contribution and influence of environmental structure on behaviour patterns'. For example, in work on the 'robolobster' described further below, Grasso et al. (2000) noted that in a sense the work they are doing is a study of the characteristics of the water-borne plume that the lobster tracks, as much as it is a study of the lobster itself.

Stimulus and Sensor Characteristics

The airborne plume that is the stimulus for pheromone tracking in the moth is a similarly complex environmental signal and several groups are using robots to model the moth's tracking behaviour. Obviously the nature of a 'signal', for an animal, is relative to its sensors. Kuwana et al. (1995) have wired the dissected antennae of the male silkworm moth, *Bombyx mori*, directly to their robot to investigate the plume-following behaviour. This biological sensor is much more sensitive than any available artificial one. Using the robot, they have shown that a control network based on the 'flip-flop' hypothesis (Kanzaki 1996) produces behaviour more comparable to the moth than a simple 'reflex'-based network (that has no memory) or an evolved recurrent network. The latter worked in simulation but not when faced with the real environment.

Ishida et al. (1999) used standard gas sensors in their robot model of moth pheromone tracking, but have mimicked, with a fan, the way that the moth uses wing vibration as a means to draw air over sensors to improve the sensors' directional accuracy. They have shown their device can localize odour directions in three dimensions. M. Willis (web site http://flightpath.neurobio.arizona.edu/Model/index.html) has used ionization detection in a robot model, as having more similar temporal response properties to the moth than available gas sensors. This robot is intended to be used in parallel studies with the real moth and detailed simulations. Chemical sensing has also been used in robots that mimic ant trail-following behaviour (Russell 1998; Sharpe & Webb 1998). The latter study tested the effects of different chemical concentrations, antennae spans and speeds on robot performance, in replication of biological and simulation studies (Calenbuhr & Deneubourg 1992), using a simple neural controller. The results qualitatively resembled those for ants: for example, showing a peak in performance at a particular concentration, improvement with larger antennae span, and improvement with slower movement.

Substrate Characteristics

The importance of environmental characteristics can also be illustrated by current investigations of robot 'fish'. The 'robotuna' (Triantafyllou & Triantafyllou 1995) showed how tail propulsion was dependent on the production of a series of vortices in the water. More generally the aim in this work is to explore the 'complex fluid dynamics that fish use to propel themselves' (J. T. Davis, web site http://web.mit.edu/towtank/www/pike). Mojarrad & Shahinpoor (1997) developed novel polymeric muscles to explore similar issues, such as the relationship of undulation frequency, production of vortices and speed of propulsion. Kato & Inaba (1998) described a different form of fish propulsion, the pectoral fin motions of 'feathering' and 'lead-lag'. Their robotic implementation was a motivating factor for the close study and characterization of the animal mechanism.

These studies have parallels in studying flight of insects, birds and bats, where understanding the aerodynamic characteristics is critical. Preliminary robot investigations of animal flight include Fearing et al. (2000) and Pornsin-Sirirak et al. (2000). Another obvious 'substrate' is realistic natural terrain in which terrestrial animals locomote. Many of the robots described thus far use wheeled propulsion and require a flat floor to run on. Although many animal experiments may have comparable conditions, natural environments do not. The aim of using robots in such natural environments is one of the inspirations for research on legged biorobots.

Body Characteristics

A feature of research on legged robots is the emphasis placed on the equal importance of mechanical structure and control systems for explaining how locomotion is achieved. This leads to an emphasis on descriptions of the dynamics of the body's interaction with the substrate. Raibert's (1986) analysis of dynamical stability in running machines also raised a number of critical questions for biological locomotion, such as the potential importance of symmetry, or the multiple means for achieving foot placement. Dynamical ideas have been directly applied to the cockroach (Kubow & Full 1999). Similar ideas have led investigators of four-legged robot walking to include a foot actuated by springy tendons (Berkemeier & Desai 1996), and to show how this can be compared to an abstracted cat hindleg to evaluate the spring-mass model of muscle function. Indeed, biomechanics is an area where there is a well-established crossover of robotic theory into biology.

These ideas apply of course to other motor control systems. Williamson (1998) has explored, using a two-armed robot torso, how problems of actuator co-ordination (for example in playing with a slinky) can be solved by relatively simple oscillators 'hooked' to the world and each other through implementation on the robot. In investigating a model of human oculomotor control, Shibata & Schaal (1999) emphasized the importance of matching the motor properties of the system, such as the effective 'spring' that returns the eye to the zero position without applying torque.

INTEGRATING DATA AND ENFORCING COMPLETENESS

One advantage of robot models is that a robot builder is forced to be complete in their specification. There are two aspects to this completeness. One is that all terms and mechanisms have to be precisely specified if the simulation is to be run to produce data. The other is that the complete 'problem' of behaviour, from an environment, through sensing and processing to actuation and its consequence in the world, has to be tackled. Unlike simulation, a robot model cannot choose an arbitrary form of input to avoid the sensing problem, or have an interpreted output that skips the actuator problem. The behaviour has to be addressed as the integration of all these factors. A consequence of this is that it becomes particularly evident where existing hypotheses are incomplete. This has led to the adoption of what could be called a methodology of incremental modelling.

Incremental Models

The cricket phonotaxis system is an example often cited in neuroethology as one of the better explored and understood systems. However, attempting to build a robot model revealed that, although much is known about the auditory system and identified neurons (Wohlers & Huber 1981; Schildberger 1988; Stumpner et al. 1995) the connectivity and means by which the behaviour is controlled is not at all clear. The strategy adopted was to build a series of models that included progressively more of the biological detail as previous models made clear what aspects of this are crucial to test. Thus the first model was relatively crude but sufficed to show that a single mechanism could underlie both the approach behaviour and the selectivity to song pattern (Webb 1995). The influence of the peripheral auditory system was then further explored in a more complex robot model that attempted to maintain the same spatial/temporal scale as the cricket, making experiments with real cricket song as stimuli possible (Lund et al. 1998). In a third implementation, this improved robot base was augmented with a dynamic spiking neuron model to carry out a more direct exploration of the function of timing in neural interactions, allowing both behavioural and 'physiological' tests on the model (Webb & Scutt 2000). One result of this work has been the suggestion that what, in the biological studies, appeared to be a critical 'recognition' neural response (Schildberger 1984), that is, the firing rates of a brain neuron that correlated with performance to a varying stimulus, may in fact play no direct role in explaining the behaviour.

A similar sequence, from relatively inaccurate but 'complete' systems to those with an increasingly close resemblance to the biological system, can be seen in the research on six-legged walking mechanisms by the biorobotics group at Case Western. Their Robot I had simple legs with two degrees of freedom and they investigated the effectiveness of a neural controller for generating gait patterns (Quinn & Espenscheid 1993) based on Pearson's flexor burst-generator model for cockroach walking (Pearson 1976). Robot II (Espenschied et al. 1996) had legs with three degrees of freedom and implemented the control circuit devised by Cruse (1990) that uses local interactions to coordinate leg movements (this stick-insect-derived controller has also been tested on robots by Pfeiffer et al. (1995), Kindermann et al. (1998) and Ferrell (1995). By incorporating several reflex mechanisms the robot was able to negotiate uneven terrain. Robot III is based on close studies of cockroach kinematics, with rear legs with three degrees of freedom, middle legs with four and front legs with five. Posture control of this robot based on proprioceptive data has been implemented (Nelson & Quinn 1998). Attempting to devise locomotion control has led productively to further detailed experimentation on the cockroach: 'many of the problems that arise in controlling the robot lead to new understanding of the animal' (Quinn & Ritzmann 1998, page 252) for example the need to stiffen leg joints prior to loading. Delcomyn & Nelson (2000) reported work on a similar robot model with cockroach-like kinematics, controlled by a pulsed pneumatic system that provides a parallel for spike-impulse control of real muscle.

Limiting Models

A number of models of hippocampal function in rat navigation have looked to improve understanding of this biological system by testing hypotheses for function within the complete system, thus setting a kind of minimum limit on the viability of the models. Burgess et al. (1997, 1998, page 292) implemented their neural net model of the hippocampus on a robot as this 'forces it to use real world inputs, tests whether it can indeed direct navigation under these conditions, and allows a fairer comparison with behavioural and electrophysiological data'. This is provided by a robot with a camera using visual cues to create 'place' cells to locate itself with respect to features of the environment.

Some alternative models of hippocampal function have also been implemented on robots (Recce & Harris 1996; Gaussier et al. 2000). Etienne (1998) critically reviewed work in this area. She argued that biological knowledge in this field is too weak for such models to be more than speculative, and that the robot's sensorimotor systems are too different from rats in the input and feedback they provide. The robot researchers argue that, while different, the input and feedback are at least 'comparably bad' and thus provide a valid test of the robustness of the model. The robot can be seen as a limiting test on the controller: if it suffices for a system of clearly inferior sensorimotor capability, then it is at least plausible for the animal. Etienne (1998, page 286) agreed that it is not trivial to be able to show a mechanism can work in the 'real physical world' and also suggested the nature of robotic interaction 'may lead to discovery of new functional principles that may influence our understanding of naturally evolved behaviour and neural events'.

Integrative Models

In this section I have so far discussed robot models of increasingly complex biological behaviours, from reactive sound source localization, through motor patterns for six-legged walking, to issues of mammalian brain function. There is also extensive research in 'humanoid' robotics (e.g. *Advanced Robotics* special issue on humanoid robots, **11**, part 6, 1997). While by no means all this work is directed at understanding human sensorimotor and behaviour systems, some of it is explicitly undertaken with that end. Examples include the anthroform arm designed by Hannaford et al. (1995) to explore spinal control circuits; the biped walking system developed by Yamaguchi & Takanishi (1997); the head–eye co-ordination systems studied by Takanishi et al. (1998b) and Clark (1998); and the hand–eye coordination systems explored by Taddeucci & Dario (1998) and Marjanovic et al. (1996). The latter group has also used their humanoid robot to investigate issues in social development in humans (Scassellati 1998), such as the Baron-Cohen model of shared gaze and imitation (Baron-Cohen 1995). A particular motivation in many of these studies is to use the robot as a way of integrating a range of data and theories under the constraints of a single device. Thus Hannaford et al. (1995, page 399) described their robot as a platform 'designed for comparing, integrating and testing diverse pieces of knowledge about motor control and spinal cord function' and argued that 'Physical modelling as opposed to computer simulation is used to enforce self consistency among co-ordinate systems, units and kinematic constraints'.

PRODUCING NEW HYPOTHESES

Biorobotic investigations have also contributed new hypotheses to biology. Some of the factors already discussed drive this production of new ideas. The problem of producing complete working systems often requires existing hypotheses to be augmented and adapted. The need to deal with real environments can change the emphasis on what the problem to be solved is, suggesting a novel solution. In addition, the robot builder is always under pressure to find a simple solution, one that involves less processing power, or is easier to implement. A particular philosophy behind much current work in biorobotics incorporates the biological ideas of 'matched filters' (Wehner 1987) and 'active perception' (Ballard 1991). That is, it looks to building the right kind of active interface between the system and the environment to provide relatively simple solutions to the control problem, as illustrated in the following examples.

Navigation

Current research related to the robot model of the desert ant *Catyglyphis* has led to several new hypotheses. The 'Sahabot' was initially constructed to test the efficacy of a compass based on sky polarization as thoroughly explored by Wehner (1994). Critical aspects of the biological system were adopted, such as using a cross-analyser configuration to obtain a relative polarization response that is independent of the ambient light intensity. The robot was tested in the same desert environment as the ants (Lambrinos et al. 1997). While the previously hypothesized 'scanning' mechanism, in which the animal uses a maximum in the response as the reference, could be successfully deployed to control behaviour, a more efficient 'simultaneous' model, using change in polarization to estimate the angle turned was equally effective.

Similarly, in investigating the use of 'snapshot' landmark navigation (Cartwright & Collett 1983), Lambrinos et al. (2000) found that a mathematically equivalent but in processing terms much simpler model for finding a heading vector from the current viewpoint to the home position could be devised. This mechanism has been implemented in analogue hardware and tested in experiments based on those performed on bees and ants (Moller, in press). It has been used to suggest what critical experiments might determine whether these animals really store the full information of a snapshot or this simplified 'average landmark vector'.

Recognition

Many animal behaviours seem to have obvious task decompositions, which have guided research efforts in both behavioural and neurophysiological studies. An example is cricket phonotaxis behaviour, in which the female cricket needs to recognize the conspecific song and approach it. In attempting to design a minimal controller for the robot cricket mentioned above, it became evident that an alternative view could be taken: that recognition could consist only in ability to localize; and 'unrecognized' signals were in fact those that could not adequately stimulate the localization device. Experiments taken as demonstrations of the 'recognition' mechanism in the cricket could be reproduced in the robot model without such a mechanism (Webb & Scutt 2000).

An analogous result has emerged in robotic investigation of bat echolocation (Walker 1997; Walker et al. 1998). Bats hunting specific targets, in this case moths with fluttering wings, are assumed first to recognize the desired target (via a remembered 3-dimensional model) and then locate it. However, by using a localization mechanism that is inherently tuned to specific patterns of energy in the echo, because it performs binaural comparison of invariant spectral cues produced by targets with periodic motion, the robot effectively 'ignores' stationary reflectors or those with different motion signatures. The results of the robot implementation of this idea show that 'target selective behaviour [in bats] need not be mediated by formal recognition of targets, but, rather, can be underpinned by a coupling between echolocator and environment to the extent that the signature signal provides a steering signal which is used directly in localisation' (Walker 1997, page 93).

Collective Behaviour

The effectiveness of relatively simple control mechanisms has also been investigated in robot models of collective behaviours such as piling and sorting in ants. In this case, the capabilities are closely related to how the environment is used to provide the cues for appropriate behaviour, a principle dubbed 'stigmergy' by Grassé (cited in Holland & Melhuish 1999). For example, by using a small set of behaviour rules relating to locally sensed events, small groups of robots have been shown to collect and sort objects successfully in a way that mimics ant capabilities (Holland & Melhuish 1999). The result emerges without any global coordination or goal states, simply from the interaction of the robots with the structures they are building. This line of robot work is now being developed in close coordination with studies of nest building in *Leptothorax tuberinterruptus* (Franks & Deneubourg 1997), to determine how they might exploit environmental 'templates' such as chemical gradients or other heterogeneities in the environment to build nest walls using only simple stimulus–response behaviours (Melhuish et al. 1999).

Kube & Bonabeau (2000) have implemented robot models of cooperative transport in ants. This work has similarly demonstrated the sufficiency of some minimal assumptions to reproduce ant capabilities such as repositioning dependent on prey size. Kube & Bonabeau (2000, page 87) noted that prior to their work 'no formal description of the biological phenomenon has been developed' and that 'surprisingly, roboticists went further than biologists in trying to model co-operative transport: perhaps the only convincing model so far is one that has been introduced and studied by roboticists'.

DISCUSSION

The range of examples presented here establishes that there are already fruitful outcomes from the use of robots for understanding animal behaviour. Essentially, biorobotics is an extension of conventional modelling methods already used widely in biology, but it has several specific strengths. In particular it is a powerful way to integrate data from different levels of investigation, to understand how the relation of brain, body and environment produce behaviour, to clarify the essential problems posed, and to devise and test hypotheses under realistic conditions.

Why not just Simulate?

For many problems posed in explaining animal behaviour, computer simulation rather than robotic implementation would seem easier to carry out, and sufficient to answer the questions. Undoubtedly this can be the case, and indeed many of the robot implementations described above followed on from simulations of the same systems. So what 'added value' might a robot model provide?

One important factor already referred to is the difference between real input and simulated input: this often reveals limitations in what a hypothesis can explain. Thus, for example, Blanchard et al. (2000, pp. 24–25) noted 'traces [from the robot] show the theory that LGMD spike rate increases . . . until collision . . . is not necessarily true when the variability of real images is considered'. Other considerations raised in robotic implementation concern the need for successfully controlling ongoing activity in real time. Blanchard et al. (2000) described how, when looming visual input leads to an avoidance response, then some suppression of the usual response to visual motion is needed during the avoidance action, but the time course of this is critical to avoid ignoring the next potential collision. Saksida et al. (1997) noted difficulties arising in their learning mechanism because the time course of reward for actions is critical and the experimenter shaping the behaviour may not manage to reinforce the right action. Voegtlin & Verschure (1999) argued that the time criticality can also be exploited in a 'self-supervising' scenario, where the robot's actions themselves determine when the reward occurs.

It could be argued that sufficiently sophisticated simulation environments would also raise these issues. However, such issues are often missed in simulations, whereas they are unavoidably encountered in real world implementations. Moreover, matching the complexity of the world in a simulation is difficult, and in general, we include only what we think is important, and thus cannot discover what is really important. This point was made by Beckers et al. (1996, page 183): 'the complexity of interactions available for exploitation cannot be matched by any practical simulation environment'; similarly 'properties of the environment are usually difficult to reproduce in simulations. Wrong assumptions about these properties may severely misguide the development of models' (Lambrinos et al. 2000, page 39).

Thus robotics provides a tough testing ground for biological hypotheses. Dean (1998, page 64) described the potential results as 'a proof in principle, which in [the case of a robot rather than simulation] can be considered to be stronger because such a closed model captures the physical context and the loop through the environment [i.e. the consequences of action]'. Indeed it can be possible to provide the hardware device with the identical environmental conditions to those used for behavioural experiments on the animal, for example the examples given above for the testing of optomotor responses (Harrison & Koch 1999; Huber et al. 1999). However, it is important to avoid leaping to the conclusion that a robot is necessarily a good model of an animal, simply because real-world deployment is attempted. It is still necessary to demonstrate that it is a valid model, through appropriate evaluation.

Evaluation?

If a control mechanism can reproduce the animal's behaviour in a robot, then it seems plausible to believe it could explain behaviour in the animal. However, there are problems with drawing this conclusion. One is that we cannot conclude from similar performance that the same causal mechanism has been found. This problem of underdetermination is recognized as a general constraint on model building in science. Strictly speaking, no matter how well our current explanation seems to fit the facts, there logically exists an infinite number of other theories that could also account for them, so we cannot conclude that we have found the true explanation. Laudan (1998) has usefully laid out variants of this idea in the philosophy of science. He pointed out that, in practical terms, we are generally lucky to have even a few good explanations to choose between, let alone an infinitude; and their logical equivalence does not correspond to there being no rational way to choose between them. Nevertheless there does remain some difficulty in drawing strong conclusions from the successful results of a biorobotic (or any other) modelling study.

Hannaford et al. (1995, page 399) argued that 'A putative theory of motor control which has been programmed into the [robot] replica will either replicate the behaviour of the system or it will not. If it does, it is thereby shown to be consistent with all of the ideas that went into the replica's design'. The majority of studies in biorobotics to date have confirmed (to this extent) rather than refuted hypotheses. In particular, it is hard to find examples in which competing hypotheses have been tested and one shown to fail. In general, it seems, hypotheses can be more or less made to work, provided they are complete and coherent (which is not always the case). This may reflect the implicit limit expressed by Hannaford et al. (1995): if we have enough flexibility in designing the robot replica, we can make sure that it does work with any given theory of motor control. However, it more likely reflects the problem that many biorobots are not subjected to sufficiently rigorous evaluation.

Proper experimental evaluation is needed to determine fully the real strengths or limitations of the implemented hypothesis. Behaviour qualitatively similar to the animal in a few trials, while encouraging, cannot be taken as confirmation, yet too few studies do more. The more strong correspondences between the model behaviour and the animal behaviour can be demonstrated, particularly where those correspondences go beyond the behaviours the model was initially built to explain, the less likely it is that the behaviour is actually being produced by an entirely different mechanism, although logically that possibility remains. Thus even when successful correspondences are found, the significance should be conservatively interpreted, for example in concluding that mechanisms shown to be sufficient to reproduce the animal's behaviour in the robot 'perhaps should be added to the list of candidates for study' (Holland & Melhuish 1999, page 200) in biology.

An example of 'good practice' in this regard, which indeed leads to the rejection of hypotheses, is given by the study of the 'robolobster' (Grasso et al. 2000). First, the robot has been designed to be tested in the same experimental flow tank as that used for the lobster; it has the same scale, speed and sensors that can be positioned at the same relative position as the lobster's antennules. Second, the analysis methods used on the lobster behaviour are directly applied to the behaviour produced by the robot. Several different 'tropotaxis' control schemes based on chemical concentrations were implemented and tested in a series of controlled trials. Although some of these resulted in successful approach to the source of the plume, several characteristics of this behaviour differed significantly from the lobster. The lobster must therefore use more than the instantaneous concentration levels to determine its movement.

Accuracy?

The other main constraint on drawing conclusions about animals from the behaviour of robot models is the effect on those conclusions of limitations in the accuracy of the models. That is, the technological limits on how well biological mechanisms can be copied may limit the force of the conclusions that can be drawn from robot models. 'Animal systems are made of very different components from robotic systems and there are no a priori grounds for supposing that the system trade-offs are identical . . .' (J. Hallam, unpublished). Problems include differences in scale, the lack of good muscle-like actuators, and an imbalance in the accuracy of various parts of the system, for example a complex controller used in a simple wheeled robot or a highly accurate sensorimotor replication controlled by an arbitrary mechanism. Finding a productive and plausible level of accuracy is one of the most difficult issues in modelling. As Brooks (1997, page 296) suggested in the context of building a humanoid robot '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 [and] only the broad outline form is mimicked'.

Several strategies can be discerned in biorobotics research in reaction to this problem. One is to limit the claims made, for example that studies of collective foraging in robots 'contribute a piece of heuristic evidence that complex social systems may be organised on decentralised organisational principles' but because 'our robots do not mimic any specific social insect species . . . no binding conclusion can be drawn by the comparative study of our robot's behaviour' (Krieger & Billeter 2000, page 67).

However this seems rather self-defeating; the question is then why not mimic a specific species, so that conclusions can be drawn? Another approach is to argue that the robot is really only being used to investigate the general character of the problem, for example 'at present we do not know enough about the actual processing mechanisms that underlie insect vision to produce a carbon or silicon copy . . . our aim is to reveal the cues that insects use to navigate' (Srinivasan et al. 1999, page 214).

However, it is not true that because a biorobot cannot be an exact replica of its target animal, it is not possible to draw any substantive conclusions about the target. All modelling involves simplification and substitution. The issue is whether the simplifications and substitutions in a particular case of modelling mean that the conclusions cannot validly be transferred. In other words, are the arbitrary accretions involved in building the robot in fact the main determinants of the behaviour, not the original hypothesis supposedly under test? To answer this researchers need to be as explicit as possible about how their robot differs from the animal, why it so differs, and what effect those differences might have.

An example is provided by the robot model of the nematode *Caenorhabditis elegans*. Morse et al. (1998) noted that there are 'snake' or 'worm'-like robots available (e.g. Hirose 1993), but for the purposes of their investigation, it was sufficient to have a robot with a forward propulsive force and a steering mechanism based on the relative contraction of opposing muscles on each side of the 'head'. This was based on experimental observations that the speed of *C. elegans* during taxis is almost constant and the direction determined by the angle of the head. Similarly, a comparable sensory gradient to the usual substrate of a two-dimensional Gaussian distribution of chemical concentration on a petri dish can be generated using a light source and a light sensor. Morse et al. (1998) were thus able to show that a particular neural connectivity (also expressed in terms of differential equations) can produce the kind of approach paths seen in the animal under conditions of real noise in motor control. Thus observationally grounded abstractions can be used to establish the correspondence between the robot and the animal with respect to testing particular hypotheses.

Interaction with Biological Investigations?

The issues of evaluation and accuracy both raise the necessity of going back and forth between the robot system and the real system. There can be a pernicious view that the appropriate methodology for biorobotics is (1) to find out everything about an animal, then (2) replicate it as closely as possible in a machine. Nolfi (1998, page 180) described Franceschini's work thus: 'by carefully studying the vision system of the fly he came up with a model so detailed it was easy for him to implement in a mobile robot able to navigate . . .'. There are several fallacies here. First, is the claim that it was 'easy'. Even for the best-studied examples, it is simply not the case that robotics researchers can directly implement the biological system: a straightforward translation of what is known about the animal into hardware and software is rarely if ever possible. Second, it implies that we cannot begin building a useful biorobot until by careful study such a detailed and easy-to-implement model has been arrived at. In fact the main use of biorobots may be to suggest what careful study might eventually lead to a detailed model. Third, Franceschini's model is not simply a direct implementation of the fly but involves several abstractions and substitutions motivated by engineering considerations based on the final task constraints.

If biorobotics is to be effective, there needs to be good two-way communication between those working on the robot and the animal. The most productive robotic groups have developed close and direct connections with the biologists studying the systems of interest; and more and more commonly biologists are taking the lead in initiating robot research programmes. Nevertheless, communication problems remain, perhaps because of the different outlooks on problem solving that result from engineering versus biological training. One manifestation of this difference is that engineers are generally good at (and keen about) generating solutions to posed problems. In the current context, this means that they will propose explicit hypotheses, where biologists tend to be more cautious. Indeed, many of the examples discussed in previous sections represent novel proposals and are the most complete models of their respective animal systems developed to date. This is a natural strength of the methodology, and as such should be seen to complement, not challenge, more traditional research approaches. As Raibert (1986, page 189) suggested in the particular context of legged locomotion 'detailed knowledge of working locomotion algorithms, like those embodied in running machines, should help to formulate good experimental questions to ask of biological legged systems'.

Biologists can justifiably be cautious in accepting as correct the alternative hypotheses put forward by roboticists. As discussed already, the fact that a mechanism is successfully deployed to imitate animal behaviour in a robot is not direct evidence that the animal functions the same way. However, it would seem biologists can only gain by consideration of such alternative hypotheses, insofar as these have been shown capable of explaining current observations on the animal. 'Consideration' includes taking up the challenge to find the critical biological evidence to rule out such alternatives, or being equally explicit in specifying and testing their own hypotheses regarding the animal systems in question.

Conclusion

Franceschini et al. (1992, page 291) listed five advantages of hardware implementation of biological models: they enforce concreteness of specification; save the difficulty of simulating complexity of real device and environment; require working in parallel and analogue; demonstrate, when a working device is produced, that the 'embodied principles [are] viable'; and verify 'beyond all expectation that the construction step itself [provides] us with deeper insights into the real problems'. The final factor was echoed by Grasso et al. (2000, page 127): 'experience has made clear a number of the intricacies of the problem . . . that were not clear at the outset'. To those involved in biorobotics projects, it is these insights derived during the construction process that are commonly felt to be the most important contribution, although they can be the hardest to express to those outside the field. I hope that the examples discussed here will have communicated some of the scope and variety of insight that biorobotics can provide.

Most current examples of biorobots are justified at least partly in terms of potential engineering advances. The aims of engineering are not incompatible with doing biologically relevant work; trying to make an artificial system work can provide a driving force for trying to understand the natural system. However, tension can be induced. Franz & Mallot (2000, pp. 148–149) argued that 'both interests cannot be pursued in the same system: either, realistic modelling of animal behaviour restricts the technical application to very specialised cases . . . or the investigated mechanism is so highly abstracted that predictions of actual animal behaviour are difficult'. Insofar as biologists become directly involved in these investigations, the 'engineering motivation' will become less important, and the focus on biological relevance will be improved.

One further way that robot models of animals may be useful to the study of animal behaviour relates to the impression made by robots as described at the beginning of this paper. It is common for observers of robot behaviour to be quick to attribute capabilities to the robot that go well beyond what it is actually doing, both in behavioural terms and especially in terms of the mechanism supporting that behaviour. Goals and intentions are readily ascribed to any directed behaviour, a stalled robot is sometimes described as 'confused' or 'thinking what to do next' and so on. Brooks (1997, page 296) noted 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 human'. Systematic study of human reactions to robot behaviour could be a useful way to explore the inherent prejudices of human observers of animal behaviour, and thus lead to improvements in traditional methodology in this field.

References

Arbib, M. & Liaw, J. 1995. Sensorimotor transformations in the worlds of frogs and robots. *Artificial Intelligence*, **72**, 53–79.

Arkin, R. C. 1998. *Behaviour-based Robotics*. Cambridge, Massachusetts: MIT Press.

Ashby, W. R. 1952. *Design for a Brain*. London: Chapman & Hall.

Ballard, D. 1991. Animate vision. *Artificial Intelligence*, **48**, 57–86.

Baron-Cohen, S. 1995. *Mindblindness*. Cambridge, Massachusetts: MIT Press.

Beckers, R., Holland, O. E. & Deneubourg, J. L. 1996. From local actions to global tasks: stigmergy and collective robotics. *Artificial Life*, **4**, 181–189.

Beer, R. D. & Gallagher, J. C. 1992. Evolving dynamical neural networks for adaptive behaviour. *Adaptive Behaviour*, **1**, 91–122.

Beer, R. D., Chiel, H. J., Quinn, R. D. & Ritzmann, R. E. 1998. Biorobotic approaches to the study of motor systems. *Current Opinion in Neurobiology*, **8**, 777–782.

Berkemeier, M. & Desai, K. 1996. Design of a robot leg with elastic energy storage, comparison to biology, and preliminary experimental results. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 22–28 April, Minneapolis*, pp. 213–218. Piscataway, New Jersey: IEEE.

Blanchard, M., Verschure, P. F. M. J. & Rind, F. C. 1999. Using a mobile robot to study locust collision avoidance responses. *International Journal of Neural Systems*, **9**, 405–410.

Blanchard, M., Rind, F. C. & Verschure, P. F. M. J. 2000. Collision avoidance using a model of the locust LGMD neuron. *Robotics and Autonomous Systems*, **30**, 17–38.

Brooks, R. A. 1997. From earwigs to humans. *Robotics and Autonomous Systems*, **20**, 291–304.

Burgess, N., Donnett, J. G., Jeffery, K. J. & O'Keefe, J. 1997. Robotic and neuronal simulation of the hippocampus and rat navigation. *Philosophical Transactions of the Royal Society of London, Series B*, **352**, 1535–1543.

Burgess, N., Donnett, J. G. & O'Keefe, J. 1998. Using a mobile robot to test a model of the rat hippocampus. *Connection Science*, **10**, 291–300.

Calenbuhr, V. & Denuebourg, J. L. 1992. A model for osmotropotatic orientation I. *Journal of Theoretical Biology*, **158**, 359–393.

Cartwright, B. & Collett, T. 1983. Landmark learning in bees. *Journal of Comparative Physiology A*, **151**, 521–543.

Chang, C. & Gaudiano, P. 1998. Application of biological learning theories to mobile robot avoidance and approach behaviours. *Journal of Complex Systems*, **1**, 79–114.

Cheng, K. 1986. A purely geometric module in the rat's spatial representation. *Cognition*, **23**, 149–178.

Chiel, H. & Beer, R. 1997. The brain has a body: adaptive behaviour emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, **20**, 553–557.

Clark, J. J. 1998. Spatial attention and saccadic camera motion. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 16–20 May 1998, Leuven*, pp. 3247–3252. Piscataway, New Jersey: IEEE.

Cliff, D., Husbands, P. & Harvey, I. 1993. Explorations in evolutionary robotics. *Adaptive Behaviour*, **2**, 73–110.

Cruse, H. 1990. What mechanisms coordinate leg movement in walking arthropods? *Trends in Neurosciences*, **13**, 15–21.

Damper, R. I., French, R. L. B. & Scutt, T. W. In press. ARBIB: an autonomous robot based on inspirations from biology. *Robotics and Autonomous Systems*.

Dean, J. 1998. Animats and what they can tell us. *Trends in Cognitive Sciences*, **2**, 60–67.

Delcomyn, F. & Nelson, M. E. 2000. Architectures for a biomimetic hexapod robot. *Robotics and Autonomous Systems*, **30**, 5–15.

Douglass, J. K. & Strausfeld, N. J. 1996. Visual motion-detection circuits in flies: parallel direction and non-direction sensitive pathways between the medulla and the lobular plate. *Journal of Neuroscience*, **16**, 4551–4562.

Egelhaaf, M., Borst, A. & Reichardt, W. 1989. The nonlinear mechanism of direction selectivity in the fly motion detection system. *Naturwissenschaften*, **76**, 32–35.

Espenschied, K., Quinn, R., Beer, R. & Chiel, H. 1996. Biologically based distributed control and local reflexes improve rough terrain locomotion in a hexapod robot. *Robotics and Autonomous Systems*, **18**, 59–64.

Etienne, A. S. 1998. Mammalian navigation, neural models and robotics. *Connection Science*, **10**, 271–289.

Ewert, J. P. 1987. Neuroethology of releasing mechanisms: prey catching in toads. *Behavioural and Brain Sciences*, **10**, 337–368.

Fearing, R. S., Chiang, K. H., Dickenson, M., Pick, D. L., Sitti, M. & Yan, J. 2000. Wing transmission for a micromechanical flying insect. In: *Proceedings of the IEEE International Conference on Robotics and Automation, April 2000, San Francisco,* pp. 1509–1516. Piscataway, New Jersey: IEEE.

Ferrell, C. 1995. Comparison of three insect-inspired locomotion controllers. *Robotics and Autonomous Systems*, **16,** 135–159.

Flynn, A. M. & Brooks, R. A. 1989. *Battling Reality. A.I. Memo 1148*. Publications Office, 545 Technology Square, Cambridge, Massachusetts, U.S.A.: MIT Artificial Intelligence Laboratory.

Franceschini, N. 1996. Engineering applications of small brains. *FED Journal*, **7,** 38–52.

Franceschini, N., Pichon, J. M. & Blanes, C. 1992. From insect vision to robot vision. *Philosophical Transactions of the Royal Society of London, Series B*, **337,** 283–294.

Franks, N. & Deneubourg, J.-L. 1997. Self-organizing nest construction in ants: individual worker behaviour and the nest's dynamics. *Animal Behaviour*, **54,** 779–796.

Franz, M. O. & Mallot, H. A. 2000. Biomimetic robot navigation. *Robotics and Autonomous Systems*, **30,** 133–153.

Gaussier, P., Joulain, C., Banquet, J. P., Lepretre, S. & Revel, A. 2000. The visual homing problem; an example of robotics biology cross-fertilization. *Robotics and Autonomous Systems*, **30,** 155–180.

Grasso, F. W. In press. Environmental information, animal behavior and biorobot design: reflections on locating chemical sources in marine environments. In: *Biorobotics* (Ed. by T. Consi & B. Webb). Menlo Park, California: AAAI Press.

Grasso, F., Consi, T., Mountain, D. & Atema, J. 2000. Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: progress and challenges. *Robotics and Autonomous Systems*, **30,** 115–131.

Grossberg, S. 1982. A psychophysiological theory of reinforcement, drive, motivation and attention. *Journal of Theoretical Neurobiology*, **1,** 286–369.

Hallam, B., Halperin, J. R. & Hallam, J. C. 1994. An ethological model for implementation in mobile robots. *Adaptive Behaviour*, **3,** 51–79.

Hannaford, B., Winters, J., Chou, C.-P. & Marbot, P.-H. 1995. The anthroform biorobotic arm: a system for the study of spinal circuits. *Annals of Biomedical Engineering*, **23,** 399–408.

Harrison, R. R. & Koch, C. 1999. A robust analog VLSI motion sensor based on the visual system of the fly. *Autonomous Robotics*, **7,** 211–224.

Hawkins, R. D. & Kandel, E. R. 1984. Is there a cell biological alphabet for simple forms of learning? *Psychological Review*, **91,** 375–391.

Hirose, S. 1993. *Biologically Inspired Robotics*. Oxford: Oxford University Press.

Holland, O. & Melhuish, C. 1999. Stigmergy, self-organization and sorting in collective robotics. *Artificial Life*, **5,** 173–202.

Huber, S. A. & Bulthoff, H. H. 1998. Simulation and robot implementation of visual orientation behaviour of flies. In: *From Animals to Animats Vol. 5* (Ed. by R. Pfeifer, B. Blumberg, J.-A. Meyer & S. W. Wilson), pp. 77–85. Cambridge, Massachusetts: MIT Press.

Huber, S., Franz, M. O. & Bulthoff, H. H. 1999. On robots and flies: modeling the visual orientation behaviour of flies. *Robotics and Autonomous Systems*, **29,** 227–242.

Ishida, H., Kobayashi, A., Nakamoto, T. & Moriisumi, T. 1999. Three dimensional odor compass. *IEEE Transactions on Robotics and Automation*, **15,** 251–257.

Kanzaki, R. 1996. Behavioral and neural basis of instinctive behavior in insects: odor-source searching strategies without memory and learning. *Robotics and Autonomous Systems*, **18,** 33–43.

Kato, M. & Inaba, T. 1998. Guidance and control of fish robot with apparatus of pectoral fin motion. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 16–20 May, Leuven*, pp. 446–451. Piscataway, New Jersey: IEEE.

Kindermann, T., Cruse, H. & Dean, J. 1998. Biologically motivated controller for a six-legged walking system. In: *Proceedings of the 24th Annual Conference of the IEEE Industrial Electronics Society, 31 August–4 September 1998, Aachen*, pp. 2168–2173. Piscataway, New Jersey: IEEE.

Konishi, M. 1993. Listening with 2 ears. *Scientific American*, **268,** 66–73.

Kortmann, R. & Hallam, J. 1999. Studying animals through artificial evolution: the cricket case. In: *Advances in Artificial Life. Lecture Notes in Computer Science, Vol. 1674* (Ed. by D. Floreano, J.-D. Nicond & F. Mondada), pp. 215–224. Berlin: Springer-Verlag.

Krieger, M. J. B. & Billeter, J.-B. 2000. The call of duty: self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, **30,** 65–84.

Kube, C. R. & Bonabeau, E. 2000. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, **30,** 85–101.

Kubow, T. M. & Full, R. J. 1999. The role of the mechanical system in control: a hypothesis of self-stabilisation in the cockroach. *Philosophical Transactions of the Royal Society of London, Series B*, **354,** 849–861.

Kuwana, Y., Shimoyama, I. & Miura, H. 1995. Steering control of a mobile robot using insect antennae. In: *IEEE International Conference on Intelligent Robots and Systems, 5–9 August, Pittsburg*, pp. 530–535. Piscataway, New Jersey: IEEE.

Lambrinos, D., Maris, M., Kobayashi, H., Labhart, T., Pfeifer, R. & Wehner, R. 1997. An autonomous agent navigating with a polarized light compass. *Adaptive Behaviour*, **6,** 175–206.

Lambrinos, D., Moller, R., Labhart, T., Pfeifer, R. & Wehner, R. 2000. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, **30,** 39–64.

Laudan, L. 1998. Demystifying undetermination. In: *Philosophy of Science: the Central Issues* (Ed. by M. Curd & J. A. Cover), pp. 320–353. New York: W. W. Norton.

Lewis, M. A. & Nelson, M. E. 1998. Look before you leap: peering behaviour for depth perception. In: *From Animals to Animats Vol. 5* (Ed. by R. Pfeifer, B. Blumberg, J.-A. Meyer & S. W. Wilson), pp. 98–103. Cambridge, Massachusetts: MIT Press.

Lund, H. H. & Asada, M. 1998. Embodied artificial life: editorial. *Artificial Life*, **4,** 303–307.

Lund, H. H., Webb, B. & Hallam, J. 1998. Physical and temporal scaling considerations in a robot model of cricket calling song preference. *Artificial Life*, **4,** 95–107.

Margules, J. & Gallistel, C. R. 1988. Heading in the rat: determination by environmental shape. *Animal Learning and Behavior*, **16,** 404–410.

Marjanovic, M., Scassellati, B. & Williamson, M. 1996. Self-taught visually-guided pointing for a humanoid robot. In: *From Animals to Animats Vol. 4* (Ed. by P. Maes, M. Mataric, J.-A. Meyer, J. Pollack & S. W. Wilson), pp. 35–44. Cambridge, Massachusetts: MIT Press.

Mataric, M. & Cliff, D. 1996. Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, **19,** 67–83.

Melhuish, C., Welsby, J. & Edwards, C. 1999. Using templates for defensive wall-building with autonomous ant-like robots. In: *TIMR 99: Towards Intelligent Mobile Robots* (Ed. by U. Nehmzow & C. Melhuish), pp. 70–78. Manchester: University of Manchester Press.

Michelsen, A., Andersen, B. B., Storm, J., Kirchner, W. H. & Lindauer, M. 1992. How honeybees perceive communication dances, studied by means of a mechanical model. *Behavioral Ecology and Sociobiology*, **30,** 143–150.

Mojarrad, M. & Shahinpoor, M. 1997. Biomimetic robotic propulsion using polymeric artificial muscles. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 20–25 April, Albuquerque*. pp. 2152–2157. Piscataway, New Jersey: IEEE.

Moller, R. In press. Insect visual homing strategies in a robot with analog processing. *Biological Cybernetics*.

Morse, T. M., Ferree, T. C. & Lockery, S. R. 1998. Robust spatial navigation in a robot inspired by chemotaxis in *Caenorhabditis elegans*. *Adaptive Behaviour*, **6,** 393–410.

Nelson, G. M. & Quinn, R. D. 1998. Posture control of a cockroach-like robot. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 16–20 May 1998, Leuven*, pp. 157–162. Piscataway, New Jersey: IEEE.

Nolfi, S. 1998. Evolutionary robotics: exploiting the full power of self-organization. *Connection Science*, **10**, 167–184.

Nolfi, S. & Floreano, D. 1998. Co-evolving predator and prey robots: do 'arm races' arise in artificial evolution? *Artificial Life*, **4**, 311–335.

Pearson, K. G. 1976. The control of walking. *Scientific American*, **235**, 72–86.

Pfeiffer, F., Etze, J. & Weidemann, H. 1995. Six-legged technical walking considering biological principles. *Robotics and Autonomous Systems*, **14**, 223–232.

Pichon, J.-M., Blanes, C. & Franceschini, N. 1989. Visual guidance of a mobile robot equipped with a network of self-motion sensors. In: *Proceedings: Mobile Robots, Vol. IV. November 1989, Philadelphia* (Ed. by W. J. Wolfe & W. H. Chun), pp. 44–53. Brellingham: Society of Photo-optical Instrumentation Engineers.

Pornsin-Sirirak, T., Lee, S. W., Nassef, H., Grasmeyer, J., Tai, Y. C., Ho, C. M. & Keennon, M. 2000. MEMS wing technology for a battery-powered ornithopter. In: *Proceedings of the IEEE International Conference in Micro Electro Mechanical Systems, 23–27 January 2000, Miyazaki*, pp. 799–804. Piscataway, New Jersey: IEEE.

Quinn, R. D. & Espenscheid, K. S. 1993. Control of a hexapod robot using a biologically inspired neural network. In: *Biological Neural Networks in Invertebrate Neuroethology and Robotics* (Ed. by R. D. Beer, R. E. Ritzmann, & T. McKenna). London: Academic Press.

Quinn, R. D. & Ritzmann, R. E. 1998. Construction of a hexapod robot with cockroach kinematics benefits both robotics and biology. *Connection Science*, **10**, 239–254.

Raibert, M. H. 1986. *Legged Robots that Balance*. Cambridge, Massachusetts: MIT Press.

Recce, M. & Harris, K. D. 1996. Memory for places: a navigational model in support of Marr's theory of hippocampal function. *Hippocampus*, **6**, 735–748.

Rind, F. C. & Bramwell, D. I. 1996. Neural network based on the input organization of an identified neuron signalling impending collision. *Journal of Neurophysiology*, **75**, 967–985.

Rind, F. C. & Simmons, P. J. 1997. Signalling of object approach by the DCMD neuron of the locust. *Journal of Neurophysiology*, **77**, 1029–1033.

Rucci, M., Wray, J. & Edelman, G. M. 2000. Robust localization of auditory and visual targets in a robotic barn owl. *Robotics and Autonomous Systems*, **30**, 181–193.

Russell, R. 1998. Odour sensing robot draws inspiration from the insect world. In: *Proceedings of the 2nd International Conference on Bioelectromagnetism, 15–18 February 1998, Melbourne*, page 1998. Piscataway, New Jersey: IEEE.

Saksida, L. M., Raymond, S. M. & Touretzky, D. S. 1997. Shaping robot behavior using principles from instrumental conditioning. *Robotics and Autonomous Systems*, **22**, 231–249.

Santos-Victor, J., Sandini, G., Curotto, F. & Garibaldi, S. 1995. Divergent stereo in autonomous navigation: from bees to robots. *International Journal of Computer Vision*, **14**, 159–177.

Savage, T. 1998. Shaping: the link between rats and robots. *Connection Science*, **10**, 321–340.

Scassellati, B. 1998. Imitation and mechanisms of joint attention: a developmental structure for building social skills on a humanoid robot. In: *Computation for Metaphors, Analogy and Agents: Springer Lecture Notes in Artificial Intelligence Vol. 1562* (Ed. by C. Nehaniv), pp. 176–195. Berlin: Springer-Verlag.

Schildberger, K. 1984. Temporal selectivity of identified auditory interneurons in the cricket brain. *Journal of Comparative Physiology*, **155**, 171–185.

Schildberger, K. 1988. Behavioural and neuronal methods of cricket phonotaxis. *Experientia*, **44**, 408–415.

Scutt, T. & Damper, R. 1997. Biologically-motivated learning in adaptive mobile robots. In: *Proceedings of the SMC '97 IEEE International Conference on Systems, Man and Cybernetics, October 1997, Orlando*, pp. 475–480. Piscataway, New Jersey: IEEE.

Sharkey, N. E. & Ziemke, T. 1998. A consideration of the biological and psychological foundations of autonomous robots. *Connection Science*, **10**, 361–391.

Sharpe, T. & Webb, B. 1998. Simulated and situated models of chemical trail following in ants. In: *From Animals to Animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behaviour, Zurich*, pp. 195–204. Cambridge, Massachusetts: MIT Press.

Shibata, T. & Schaal, S. 1999. Robot gaze stabilisation based on mimesis of oculomotor dynamics and vestibulocerebellar learning. *Advanced Robotics*, **13**, 351–352.

Sobel, E. C. 1990. The locust's use of motion parallax to measure distance. *Journal of Comparative Physiology A*, **167**, 579–588.

Srinivasan, M. V., Chahl, J. S., Weber, K. & Venkatesh, S. 1999. Robot navigation inspired by principles of insect vision. *Robotics and Autonomous Systems*, **26**, 203–216.

Stumpner, A., Atkins, G. & Stout, J. F. 1995. Processing of unilateral and bilateral auditory inputs by the ON1 and L1 interneurons of the cricket *Acheta domesticus* and comparison to other cricket species. *Journal of Comparative Physiology A*, **177**, 379–388.

Taddeucci, D. & Dario, P. 1998. Experiments in synthetic psychology for tactile perception in robots: steps toward implementing humanoid robots. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 16–20 May 1998, Leuven*, pp. 2262–2267. Piscataway, New Jersey: IEEE.

Takanishi, A., Aoki, T., Ho, M., Ohkawa, Y. & Yamaguchi, J. 1998a. Interaction between creature and robot: development of an experiment system for rat and rat-robot interaction. In: *IEEE/RSJ International conference on Intelligent Robotics and Systems*, pp. 1975–1980. Los Alamitos, California: IEEE Computer Society Press.

Takanishi, A., Hirano, S. & Sato, K. 1998b. Development of an anthropomophic head–eye system for a humanoid robot. In: *Proceedings of the IEEE International Conference on Robotics and Automation, 16–20 May 1998, Leuven*, pp. 1308–1314. Piscataway, New Jersey: IEEE.

Triantafyllou, M. S. & Triantafyllou, G. S. 1995. An efficient swimming machine. *Scientific American*, **272**, 40–48.

Vaughan, R., Sumpter, N., Henderson, J., Frost, A. & Cameron, S. 2000. Experiments in automatic flock control. *Robotics and Autonomous Systems*, **31**, 109–117.

Verschure, P. F. M. J. 1998. Distributed adaptive control: explorations in robotics and the biology of learning. *Informatik/Informatique*, **1**, 25–29.

Voegtlin, T. & Verschure, P. F. M. J. 1999. What can robots tell us about brains? A synthetic approach towards the study of learning and problem solving. *Reviews in the Neurosciences*, **10**, 291–310.

Walker, A. 1997. One tone, two ears, three dimensions: an investigation of qualitative echolocation strategies in synthetic bats and real robots. PhD. thesis, University of Edinburgh.

Walker, A., Peremans, H. & Hallam, J. 1998. One tone, two ears, three dimensions: a robotics investigation of pinnae movements used by rhinolophid and hipposiderid bats. *Journal of the Acoustical Society of America*, **104**, 569–579.

Walter, W. G. 1961. *The Living Brain*. Harmondsworth, Middlesex: Penguin.

Warzecha, A.-K. & Egelhaaf, M. 1996. Intrinsic properties of biological motion detectors prevent the optomotor control system from getting unstable. *Philosophical Transactions of the Royal Society of London, Series B*, **351,** 1579–1591.

Webb, B. 1995. Using robots to model animals: a cricket test. *Robotics and Autonomous Systems*, **16,** 117–134.

Webb, B. & Scutt, T. 2000. A simple latency dependent spiking neuron model of cricket phonotaxis. *Biological Cybernetics*, **82,** 247–269.

Weber, K., Venkatesh, S. & Srinivasan, M. 1998. An insect-based approach to robotic homing. In: *Fourteenth International Conference on Pattern Recognition* (Ed. by A. K. Jain, S. Venkatash & B. Lovell), pp. 297–299. Piscataway, New Jersey: IEEE.

Weber, T. & Thorson, J. 1988. Auditory behaviour in the cricket. Interaction of direction of tracking with perceived temporal pattern in split-song paradigms. *Journal of Comparative Physiology A*, **163,** 13–22.

Wehner, R. 1987. Matched filters: neural models of the external world. *Journal of Comparative Physiology A*, **161,** 511–531.

Wehner, R. 1994. The polarization-vision project: championing organismic biology. In: *Neural Basis of Behavioural Adaptations* (Ed. by K. Schildberger & N. Elsner), pp. 103–143. Stuttgart: Gustav Fischer Verlag.

Williamson, M. W. 1998. Rhythmic robot arm control using oscillators. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 77–83. Los Alamitos, California: IEEE Computer Society Press.

Wohlers, D. W. & Huber, F. 1981. Processing of sound signals by six types of neurons in the prothoracic ganglion of the cricket *Gryllus campestris* L. *Journal of Comparative Physiology*, **146,** 161–173.

Yamaguchi, J. & Takanishi, A. 1997. Design of biped walking robots having antagonistic driven joints using non-linear spring mechanism. In: *IEEE/RSJ International Conference on Intelligent Robotics and Systems, 1996, Grenoble*, pp. 251–259. Los Alamitos, California: IEEE Computer Society Press.