Comparing Robot and Animal Behaviour

Bridget Hallam and Gillian Hayes

Abstract

The simulation of animal behaviour on a mobile robot is a useful undertaking for roboticists wishing to develop robots for non-predictable domains or with more generality of application than is currently possible. It can also be valuable to ethologists, as a means of testing theories of animal behaviour in a modular manner, repeatably, and with identically 'naive' subjects.

Not all tasks we would like robots to undertake require high standards of accuracy, repeatability or even reliability. Some require more in the way of adaptability, flexibility of approach and robustness. Given the disparate nature of these requirements and the variety of ways in which some of these words are used, it is not always obvious what is the answer to questions such as: how can these requirements be met? how can the results be measured? how can the performance of different robots be compared? The first question requires answers which are specific to the individual robot and the preferred style of its designer and user. The second and third questions are better answered in a more general way which does not depend so much on the particular robot, environment or task. We address the second and third questions here.

First we briefly mention some of the benefits of trying to get a robot to behave, in some senses, like an animal. Then we outline some of the fundamental capabilities of animals and examine the extent to which these capabilities are available on present-day robots. This provides a mapping between robot and animal capabilities. We then present a list of capabilities which we consider are probably essential to general autonomy, and consider how the headings could be met with current technology. This is given as an aid to designers of robots for tasks requiring adaptability, flexibility of approach and robustness; and also as a standard against which existing robots can be contrasted and compared.

It should be noted that we expect this minimal standard to be revised in the light of advances in robot technology and programming.

1 Introduction

Animal behaviour is adaptable, flexible and robust. This generalisation applies to many different types of animal across most taxonomic boundaries. Assuming evolution to be conservative, this implies that there is a common basic set of behavioural capabilities which are necessary and complete in terms of their ability to produce adaptable, flexible and robust behaviour. The identification of such a set and its use in practice would be a great step forward for many robotics applications, especially those involving environments which change in a non-predictable way or several independent and parallel tasks.

In our work we have attempted to identify behavioural capabilities common to a wide variety of animals and to note which are present on existing robots. We have, as yet, made no attempt to identify which of the items from our list are *essential* for adaptable behaviour. We have found that most mobile robots have some of these behavioural capabilities in some form or other, but we have found none which has all the items.

This list of basic behavioural capabilities includes the following: sensing of internal state, sensing the environment, mechanisms to ensure homeostasis of certain chemicals, energy management (including cyclic behaviour), mechanism(s) for translatory movements, mechanism(s) for bending movements, reactions (including reflexes, taxes and fixed (or modal) action patterns), deliberative behaviour, the ability to return to places seen before, communication, and plasticity of response (including associative learning and latent learning).

This list contrasts slightly with the list given by Beer (1990, pp 21-24) in that, as roboticists, we concentrate more on the physical actions and capabilities required.

1.1 Why Simulate Animals on Robots?

There are several good reasons for attempting to imitate the building blocks of animal behaviour on a robot. One set concerns questions about the completeness and sufficiency of ethological theories of individual animal behaviour. Another concerns insights into the essentials of autonomy and 'common sense', which should prove useful for future generations of robot designers.

From an **ethological perspective** the robot provides an alternative means for testing theories of behaviour. With real animals, it is difficult to know the effect of underlying factors not included in the model; it is also difficult to ensure that the animals being used are truly naive. These particular problems are avoided by the use of some sort of simulation. With computer simulation the level of available computing power and the patience and real-world experience of the programmer mean that simplifications are made in the 'animal', the 'environment', and the interaction between the two (sensing and effects of actions). One of the criticisms of much computer simulation of animal behaviour is that the various components being modelled are oversimplified. Some of this oversimplification can be avoided by the simulation of behaviour not on a computer but on a physical robot acting in the real world. Other simplifications and assumptions are changed by the use of a robot, resulting in a complementary interpretation of whatever is being modelled. Computer simulations have the great advantage of being able to run in faster than real time, and can produce a vast amount of statistics for analysis. Robots can also be made to produce statistics (though in smaller quantity) — but the most obvious result of a robot run is some observable behaviour which can be analysed where desired along lines similar to those used for analysing animal behaviour. This is only possible where some sort of mapping exists between robot actions and their animal counterparts. The material in sections 2 and 3 of this paper is intended to provide such a mapping.

Using robot simulation of behaviour, only the animal has to be modelled. The environment is also simplified (being usually an indoor environment) and the interactions are also different but many real-world features are retained. These include occlusions, randomly moving objects, lighting variations and sensor inaccuracies. These features can cause the results of a test run with a robot to be completely different from the results predicted by computer simulation. Under some circumstances the results produced by robot may be more realistic than those produced by computer.

From a **robotics perspective**, animals are the only extant examples of truly autonomous systems capable of carrying out parallel tasks in an unstructured environment. Alternatively, one could say that only animals exhibit 'common sense'. Since we wish to produce robots that can act sensibly and autonomously in a real environment and our best efforts to date have had very limited success, it is hoped that understanding imitations of animals will be a useful first step towards understanding the basis of autonomy, and that we can then take this understanding and apply it in more industrial areas of robotics. There has long been an undercurrent of thought suggesting that biological inspiration is as important as good engineering principles in the design of successful robots [Wie48], [Wal50], [Rai86], [Bal88], [MYC92] but this inspiration has often been restricted to the production of new, biologically inspired, principles. This approach has met with considerable success in limited areas. It seems that there is a place for attempting to apply biological *mechanisms* as well as biological principles, if we wish our robots to be better at some of the things that animals do easily. (There is also obviously a place for imitating biological structures too, such as sensory organs and muscle arrangements (see, for example, [PBF90] and [Mea89]).)

1.2 How Can the Results be Assessed?

Currently, most robots are assessed in terms of a specific task in a specific and often tightly constrained environment. However, for adaptable behaviour it is not desirable that the environment should be so constrained and, although the task may be fixed (*e.g.* 'push all boxes to corner x'), the variation possible within such a task means that statistical comparisons are difficult. Moreover, given that robots are normally fairly slow creatures, statistically significant results take too long to obtain. This could also be a problem with using normal ethological behavioural assessment on robots, depending on the rarity of the behaviour being modelled.

In section 4 of this paper we provide a list of basic behavioural capabilities which is intended both as a 'check-list' of functionalities for autonomous robots and as a standard against which robots can be compared. This will provide a means of comparing robots with animals and, perhaps more usefully at present, of comparing robot with robot. In particular, different robot morphologies can be compared in different laboratories without requiring extensive standardisation of task or environment. Also, the behaviour of the same robot under different control structures can be examined. It is not envisaged that any quantitative analysis will be made possible by comparison with this list of abilities, as most headings encompasses a wide range of mechanisms at varying levels of abstraction. However, a qualitative analysis can be made.

The set of abilities given in the standard is obviously rather basic and the mapping between robot and animal abilities is loose. This is because this standard has been designed to be achievable with current robot technology — it is a first step in the evolution of autonomous and maybe even general-purpose robots.

Purposive language is used throughout this document, as it improves the flow of the discussion.

2 Some Basic Animal Characteristics

Animals exhibit a wide range of behavioural characteristics, yet all complex animals have a variety of abilities in common. Many of these basic capabilities are too trivial to mention from an ethologist's point of view, yet even these are the subject of much admiration when they are successfully implemented on a robot. In this section we describe the abilities which we consider fundamental to the actions of most animals. These headings are probably controversial as they are nothing like the headings under which animal behaviour is usually discussed. Instead, these abilities can be considered as the basic building blocks from which actions and ultimately behaviour emerge.

Sensing Internal Factors. Animals sense several types of internal factors: these include various joint and muscle parameters, the concentrations of many chemicals, matted fur, injury and illness. Small injuries may cause actions to be done differently as painful movements are avoided, but will not always affect the choice of action. Matted fur or feathers will eventually cause grooming or preening actions in healthy animals.

Sensing the Environment. There are various environmental features that are important to an animal, but at any one time most aspects of the environment are not significant. Features which are normally significant include food, predators, other family members, temperature, and landmarks: anything directly relevant to the animal. When the animal senses that a significant feature may be present (for example, it hears a noise which might indicate a predator), it may stop doing other things in order to concentrate on the sensor concerned; in other words, it *focusses attention* on relevant stimuli.

Novel features need to be noticed, so that they can be considered and classified. Animals generally notice and orient towards new features and then decide what to do next (which is often to go back to what they were doing before). This ability to distinguish changes in the environment is very specific, so that a new shadow causes swift and immediate action whereas the moving shadows of trees waving in the wind rarely cause any response. Animals have *enormous* numbers of sensors of many different types which enable them to distinguish between similar sensory conditions.

Homeostasis. The concentrations of many of the chemicals inside an animal's body have to be kept within non-lethal limits, as do other factors such as internal temperature. The tolerated deviation from optimal may change depending on time of day (for example, body temperature is slightly lower at night for animals active during the day) or conditions (*e.g.* a camel will accept a higher body temperature if there is a shortage of water).

There are various methods of keeping relevant factors within the desired limits. For example, high internal temperature can be avoided by physical methods such as sweating, by a general reduction in activity, or by a change in activity such as to seeking shade or water. Animals have many chemicals which must be homeostatically controlled and many different ways of achieving this control. There is always a fast method of influencing important factors; there may also be other methods which may be more energy efficient: for example blood oxygen levels can be raised by panting (an instant but energy-expensive method) or by increased blood red cell concentration (a slow but efficient method).

Energy Management. Animal lifespans vary from a few days, *e.g.* some insects, to over a hundred years, *e.g.* tortoises. During this time an individual will be constantly using energy, which has to be replaced from somewhere. Finding food is a major component of much animal behaviour and itself requires energy expenditure. Energy management involves solving problems such as when to start looking for food, what to do when food supplies are scarce, and how to choose appropriate times for actions requiring large amounts of energy.

Many animals minimise their energy expenditure when food is less available, less easily found, or when dangers are greatest. This results in **cyclic behaviour** attuned to the time of day or year; and in acyclic behaviour which depends on factors such as weather conditions. Somehow animals know, perhaps by sensing environmental light levels, temperature, *etc.*, else by some sort of internal clock, which part of the day is which. Any internal clock is, however, resettable — as evidenced by the 'recovery' of animals from jet-lag, when their activity cycle slowly adjusts to the change in daylight hours until the individual is active during the same part of the day as it was at home.

Translatory Movements. Healthy animals manage to move around their normal environment without hitting things or falling over very often. They do this at various speeds, over complicated terrain or in turbulent fluid, without apparent difficulty even while concentrating on something else such as a predator behind them. Many animals are very efficient at chasing and/or evading others, and can also avoid objects coming towards them up to a certain speed.

Animals can move in different ways relative to the sensed object — including 'towards', 'away', and 'around'. They can 'keep away' as well as 'move away' — the former being a proscriptive behaviour (defining what *not* to do) rather than a prescriptive one (defining what should be done).

Other Movements. Vertebrates are very supple and are generally excellent at manipulating objects, within the limits of their appendages. Some vertebrates can reach every part of their surface with some appendage, often for grooming purposes.

Reactions. Reactions can be described as 'actions which do not require deliberation'.These include *reflexes* where the response dies away quickly, *taxes*, and also longer non-deliberative sequences of actions such as the *fixed* or *modal* action patterns (*FAPs* or *MAPS*) often mentioned by ethologists.

All animals need to respond quickly in some circumstances, for example, when they feel themselves to be falling. In some cases it is only necessary to stop, perhaps to take time to work out what it is best to do next. In other cases a more active instantaneous response may be required.

Deliberative Behaviour. Much of the behaviour exhibited by vertebrates is not entirely reactive but has both reactive and non-reactive components. These non-reactive components include both learning (see below) and actions requiring deliberation. Actions which are normally classed as deliberative include curiosity responses to new stimuli.

Navigation. Most animals can find their way back to previously visited locations. Some appear to do this by retracing their steps but it seems more common for them to remember the overall heading of where they are going or else to use landmarks. This shows that there are many acceptable methods of achieving this important ability. Significant landmarks must be remembered in some way, either 'as they stand' or *via* markers of some sort (often pheromonal).

Animals tend to follow reasonably direct paths despite the perturbations caused by their having to avoid obstacles.

Communication. Most animals communicate within their species, at least to the level of female/male attractants. Many vertebrates have a large repertoire of communicative behaviour which gives information about the individual involved. This aids in mate selection and often enables territorial disputes to be settled without recourse to actual fighting. The communication may be mediated by many factors, including scent markers, song complexity, visual display, or ritualised fighting of some sort.

Communication between species is very limited and often involves deceit, in that one species tries to appear either harmless or dangerous when it is not. Examples include eyespots on moth wings startling birds and the light on an angler fish attracting small prey fish.

Learning. Animals exhibit many different types of learning, including associative learning where a stimulus becomes associated with a reward or punishment (as in classical Pavlovian conditioning), and latent learning where there is no particular reward or punishment.

The ability to learn is beneficial to the survival of any animal. Herd animals are amongst those born with a remarkable repertoire of innate skills; however, even these seem to learn most of the social behaviour appropriate to their species. Many other vertebrates spend a significant proportion of their pre-puberty learning.

3 But a Robot is Not an Animal

Since there are definite differences between the hardware of a robot and the physical characteristics of any animal, it is obvious that the detailed behaviour of the first will not be particularly close to the detailed behaviour of the second without a major revolution in technology. However, there are many correlations which can be made which allow a mapping between the capabilities of robots and animals. In this section we mention current capabilities of state-of-the-art robots under the headings given earlier for animals.

Sensing Internal Factors. Many robots have internal sensors measuring properties such as axle/wheel revolutions, power levels, and joint angles. Encoder counts may be used to estimate external parameters such as the distance travelled as well as internal factors such as joint angle.

Sensing the Environment. Sensors exist that can sense most of the factors that animals can sense: light, sound, acidity, some other chemicals, temperature, surface distortion, pressure, wind direction, magnetic dip and direction, *etc.*. Man-made sensors also exist which can detect factors invisible to animal senses, such as radio waves and short wavelength electromagnetic radiation. While it is possible to duplicate several biological sensory capabilities in isolation, the versatility, dynamic range, and sheer number of sensors in even simple animals defeats engineering. Robots typically have orders of magnitude fewer sensors than animals and often only a rather limited range of different types of sensor. Most robots manage to sense the environment continuously, although the sampling rate often changes with the amount of processing going on, as in a real animal.

Homeostasis. Laboratory robots typically only need to keep their power levels within limits, as other relevant factors such as temperature and humidity are kept within limits by constraints on the environment. 'Too high' a battery level need not occur as some types of battery do not overcharge. The most appropriate 'too low' threshold depends on how much energy it is likely to take for the robot to find a recharging point. The energy taken to find the recharging point can be considered equivalent to the energy used by the animal in finding food.

Many mobiles have no factors which are kept under homeostatic control.

Energy Management. Robots which use mains power directly can keep going until they break down, but these tethered robots are not entirely autonomous. The batteries on mobile robots are more limiting, with few systems capable of running the robot at full activity for more than a few hours; some types of robot may have an operating time of less than an hour. With robots that can recharge without forgetting all they have learned, longer 'lifespans' can be achieved.

Robots with an operating time longer than that of their batteries will automatically have **cyclic behaviour** in much the same way as animals, in that they will need to take 'time off' to (literally) recharge their batteries. Current battery technology is not advanced enough to power a fully active robot for more than a few hours, so the robot 'day' will have to be shorter than 24 hours. However, it would be useful for the robot day to tie in with the 24 hour day so that the researcher can view the robot's activity at a convenient time. Animals which live in desert places often have a crepuscular rhythm, being active in the 'cool of the day' (*i.e.*, morning and evening), which may be a more suitable model for robots.

Translatory Movements. Robots generally move in a single plane on fairly even terrain. (The exceptions are mostly missiles and underwater robots which have to be able to move and navigate in 3D.) Robots rarely fall over mostly because they are designed to be extremely stable, many using wheels which are always supposed to touch the ground. Even legged robots are normally designed for great static stability *e.g.* CMU's Planetary Rover robot 'Ambler' [SK91], although there are some exceptions, notably Raibert's pogo stick robots [Rai86].

Careful design of mobile robot control systems has ensured that modern robots generally do well at avoiding stationary obstacles. Most robots avoid moving objects in their environment as a side effect of being programmed not to run into obstacles themselves. It works because the robots don't take account of whether they themselves are moving.

Many robots can move 'towards', 'away from', or 'along beside' objects. However robots are rarely programmed to do proscriptive actions; even obstacle avoidance movements are often implemented in a prescriptive way where the robot is told which direction to take, rather than in a proscriptive way where the robot is allowed to go anywhere except towards the object.

Other Movements. Some mobile robots carry a manipulator which enables the collection and manipulation of objects. Other mobiles are restricted to pushing objects about by running into them. The robots themselves have rigid bodies: some multijoint manipulators can be flexible, but in general these are not mounted on mobile robots. Robot sensors are sometimes capable of being rotated and translated separately from the robot body, *e.g.* the cameras on the robot head of [MYC92].

Reactions. All robots have *reflexes*, at least to the extent of becoming stationary whenever their emergency stop button is depressed! *Taxes* are also often represented; many robots can follow a slowly moving target and/or can be 'shepherded' by stimulation of some low-level avoidance routines. Larger scale reactive behaviour more analogous to *fixed action patterns* is shown by some robots such as the MIT robot Herbert [Con88] which does a small and uninterruptible set of actions (a 'behaviour' in Brooks' terminology) in response to specific environmental conditions.

Deliberative Behaviour. Traditionally programmed robots often use a 'stop-think-act' paradigm which results in all their actions being deliberative. A few robots do not stop while they think, which can cause problems *e.g.* see [Peb91].

Navigation. Robot navigation and map building has been (and still is!) the subject of considerable study. Various methods have met with success; including gradient field techniques [AD90], vector summation or 'dead reckoning' [For92], feature-based navigation [HFH89] and self-organising nets [NSH91]. The 'correct' representation for maps has also been the subject of much debate, partly because it depends largely on how the maps are intended to be used.

Despite the changes in direction caused by obstacles, animals generally manage to go fairly directly to their goal. Robots often have more difficulty and tend to retrace their steps, which only works in environments where most obstacles are generally stationary. Some 'return to heading' competence can be done by re-sensing — *e.g.* 'go towards beacon A' will keep the agent heading towards the same point and will cause animal-like behavioural anomalies under certain circumstances (such as when beacon A is beyond the bottom of a 'U'-shaped obstacle).

Communication. All robots communicate with their researchers. Most of this is extremely one-sided, with the human telling the robot what to do. However some communication from robot to human also takes place in that the robot may power lights or emit sounds which inform the researcher as to what state the robot is in or which sensors are active — an important aid to debugging its program.

Robots which cooperate to achieve tasks often need to communicate with each other. This can be done using any of the principal sensing modalities open to robots — *e.g.* a particular light frequency may act as a signal [Mei91].

Learning. Considerable research has been carried out into learning in robots by the inclusion of self-organising nets *e.g.* see the review in [NSH90]. Other approaches include the reinforcement learning architectures reviewed in [Sut90] and the chaos based learning described in [Ver91].

4 The Proposed Standard

From consideration of the factors discussed above it can be seen that robots already possess most of the different types of ability outlined. However there are very few robots which have all the capabilities mentioned. In this section we consider what capabilities might be necessary in a robot designed for anything like animal-quality behaviour. We believe that true autonomy will require at least a large subset of the different abilities described below.

1. **Sensors.** The robot should be sensor-rich with as many sensors of as many different modalities as is feasible. These should include directional sensors and whatever else is necessary to equip the robot for the functionalities outlined here. The robot should sense the environment continuously; preferably all sensors should be read continuously.

2. **Homeostasis.** The robot should have internal factors which it keeps within preset bounds. Power level is one such suitable factor.

3. **Energy Management.** The robot should have an active life of several times its 'one-charge' operating time. Preferably it should be able to recharge itself without any human intervention.

   Cyclic behaviour occurs automatically in robots with rechargeable batteries. The cycle should preferably not be regulated entirely by energy level; instead an internal clock which is recalibrated periodically (at sunrise? when recharging?) would be better. Note that the robot should still be able to respond to stimuli when recharging.

4. **Translatory Movements.** The robot is expected to move around without hitting things, with the ability to keep fairly close to the object it is circumnavigating if required. It should be able to follow or evade an object which is moving in a non-predictable way, preferably with reasonable speed and accuracy.

   The robot should be able to make both proscriptive ('don't go there') as well as prescriptive ('go there') movements.

5. **Other Movements.** Robots do not yet need to be supple, but should be able to manipulate objects at least to the extent of pushing them from place to place and changing the object orientation.

6. **Reactions.** There should be some stimuli to which the robot responds immediately, whatever else it is doing. Taxes and some longer sequences of actions which are virtually non-interruptible should also be present.

7. **Deliberative Behaviour.** There should be some stimuli to which the response of the robot changes with the number of presentations, *e.g.* new environmental features which are not of interest could be studied and then ignored.

   The robot should be able to learn significant features of its environment and stimulus combinations which should result in particular actions. If possible, the robot should decide for itself what makes a feature or stimulus significant.

8. **Navigation.** The robot should have a 'home base' to which it can return whenever circumstances demand or the researcher commands.

9. **Communication.** The robot should be able to communicate in a way that at least the researcher can understand and explain with reasonable plausibility!

10. **Learning.** One candidate for latent learning in robots is the learning of some sort of map for navigation. Associative learning would also be desirable, but should maybe not be a requirement until version 2 of this list.

5 Summary

In this paper we have argued that simulating animal behaviour on a robot is sensible and interesting both from the point of view of ethologists wanting to test theories of animal behaviour and of roboticists wanting to investigate robot autonomy. Given such simulations, it can be seen that some means of comparing the resultant robot behaviours is essential in order for discussion of the successes and limitations of the various theories and robots to be meaningful across different research laboratories.

The list of abilities described in section 4 forms the outline of the behaviour of a generalised, small, solitary vertebrate without any opportunity for reproduction. This list is primitive in biological terms as it is appropriate for the current state of robot technology; it is proposed as a standard or target against which the capabilities of individual mobile robots can be compared and contrasted. It is not intended to give any indication of the absolute worth of any individual robot, but to give an idea of the extent to which the robot can be expected to be autonomous. Since this list is not tied to any particular environment, task or control architecture, the details of how these abilities are achieved will vary enormously from robot to robot and it may well not be possible to say to what extent one robot is better than another under any particular heading. Most implementations will only be suitable for specified environments or tasks. If an implementation can be found which is less constraining, then a successful robot so produced should be capable of operating in a wide variety of ways, *i.e.* it will be to some extent a general-purpose robot.

We hope that this standard will prove useful to those ethologists with well-specified models of individual animal behaviour, in that computational versions of the models can be tested and their sufficiency examined under different headings by reference to the set of targets given.

We hope that this standard will also inspire robot builders aiming for autonomy in that it will provide a target to aim for, abilities to achieve, and a reminder of the range of abilities which is possible. We intend using this standard ourselves to compare the behaviour of mobile robots using disparate control systems.

It is anticipated that the standard outlined in this paper will be expanded and redefined in the light of further experimentation and discussion with other mobile robot researchers. It will also change as technology improves — here we have restricted the proposed standard to a simple level that should be realistically achievable with current technology. These extensions may involve the inclusion of further parts of an animal's behavioural repertoire or more detailed definition of the abilities already included. Some of the extensions to the standard which can be foreseen involve the inclusion of a requirement for map-building, for communication between robots, and for some mechanism enabling self-maintenance.

We are not claiming that a robot which has all the capabilities on our list (in its current form) will automatically be a good general-purpose robot. Our intention is that this list be refined and developed as our experience progresses. We do claim that a list of this sort could function as a general standard with reference to which the behaviour of real robots of different types and abilities can be measured and compared. Consideration of such a standard will allow designers of robots for tasks requiring adaptability, flexibility of approach and robustness to proceed in a principled manner, and adoption of a common currency will facilitate the testing of theories of animal behaviour on robots.

Acknowledgements

Thanks go to the many people who reviewed this paper and commented on its intelligibility: Janet Halperin, John Hallam, Peter Ross, Toby Tyrrell, Barbara Webb, Joanna Bryson, Ulrich Nehmzow, Ian Franks, Mandy Haggith, Glenn Reece and Weiru Liu.

One of us (BH) acknowledges receipt of SERC grant number 9130847X.

References

[AD90] Tracy Anderson and Max Donath. Animal behaviour as a paradigm for developing robot autonomy. *Robotics and Autonomous Systems*, 6, 1990.

[Bal88] D.H. Ballard. Eye movements and spatial cognition. In *AAAI spring symposium series on physical and biological approaches to computational vision*, Stanford, CA, March 1988.

[Con88] Jonathon Connell. A behaviour-based arm controller. AI memo 1025, MIT, 1988.

[For92] Peter Forster. Personal demonstration. 1992.

[HFH89] John Hallam, Peter Forster, and Jim Howe. Map-free localisation in a partially moving 3D world : the Edinburgh Feature-Based Navigator. In *Proceedings of Intelligent Autonomous Systems 2*, volume 2, Amsterdam, 1989.

[Mea89] C. A. Mead. *Analog VLSI and neural systems.* Addison–Wesley, 1989.

[Mei91] Richard Mein. Cooperative behaviour in uniformly and differentially programmed Lego vehicles. Master's thesis, Department of Artificial Intelligence, 1991.

[MYC92] J Mayhew, Z Ying, and S Cornel. The adaptive control of a 4 degree of freedom stereo camera head. AIVRU. Technical report, University of Sheffield, 1992. In preparation.

[NSH90] Ulrich Nehmzow, Tim Smithers, and John Hallam. Location recognition in a mobile robot using self-organising feature maps. In *Proceedings of International Workshop on Information Processing in Autonomous Mobile Robots*, 1990.

[NSH91] Ulrich Nehmzow, Tim Smithers, and John Hallam. Location recognition using self-organising feature maps. In *Proceedings of International Workshop on Information Processing in Autonomous Mobile Robots*, Munich, March 1991. Springer-Verlag. Also available as DAI Research Report no. 520.

[PBF90] Jean-Marc Pichon, Christian Blanes, and Nicolas Franceschini. Visual guidance of a mobile robot equipped with a network of self-motion sensors. In *SPIE Mobile Robots IV*, Philidelphia, Pennysylvania, 1990. Proceedings of the Int Soc Optical Eng.

[Peb91] Miles Pebody. How to make a Lego robot do the right thing. Master's thesis, Department of Artificial Intelligence, 1991.

[Rai86] Marc Raibert. *Legged Robots that Balance.* MIT Press, Cambridge, MA, 1986.

[SK91] Reid Simmons and E Krotkov. An integrated walking system for the Ambler planetary rover. In *Proceedings of the 1991 IEEE International Conference on Robotics and Automation*, pages 2086–2091. IEEE Society Press, April 1991.

[Sut90] Richard Sutton. Reinforcement learning architectures for animats. In Jean-Arcady Meyer and Stewart Wilson, editors, *From Animals to Animats*, pages 288–296, 1990.

[Ver91] Paul Verschure. Chaos based learning. In *Complex Systems*, volume 5, 1991. In press.

[Wal50] W.G. Walter. An imitation of life. *Scientific American*, 182(5):42–45, May 1950.

[Wie48] N. Wiener. *Cybernetics, or Control and Communication in the Animal and the Machine.* John Wiley, New York, 1948.