AI and A-Life: Never Mind The Blocksworld*

Dave Cliff
School of Cognitive and Computing Sciences,
University of Sussex,
Brighton BN1 9QH, U.K.
`davec@cogs.susx.ac.uk`

Abstract

This paper discusses the relationship between Artificial Intelligence (AI) and Artificial Life (A-Life). A-Life research addresses a wide range of phenomena, some of which have no obvious bearing on AI research. The work most relevant to AI is sufficiently coherent and distinct that it is best referred to by its own name: it is *Adaptive Behavior* research which is most likely to have significant impact on issues traditionally studied in AI. Some motivations for adaptive behavior research are reviewed, and some of the differences between adaptive behavior and traditional AI are discussed. One significant feature of current adaptive behavior research is a focus on relatively simple and specialised cognitive functions, an approach which invites unfavourable comparisons with the "blocksworld" simplified domains which were popular in AI research of the early 1970's. However, such comparisons usually overlook fundamental differences between the blocksworld-AI and Adaptive Behavior approaches to issues of simplicity and specialisation.

1 Introduction: what is A-Life?

It would be difficult to discuss the relationship between AI and A-Life without attempting to define both fields. It makes sense to start with A-Life, because this newer field is attracting a lot of attention at the moment, and with this attention comes the danger of counterproductive misunderstandings.

Put most simply, A-Life research is concerned with the study of artificial systems which exhibit lifelike behaviors. The rationale for such research is probably best characterised in the words of Chris Langton, writing in the preface to the first international meeting on A-Life, which he organised in 1987:

> Artificial systems which exhibit lifelike behaviors are worthy of investigation on their own rights, whether or not we think that the processes they mimic have played a role in the development or mechanics of life as *we* know it to be. Such systems ... expand our understanding of life as it *could* be.

 To appear in: A. Cohn (editor) *ECAI94: Proceedings of 11th European Conference on Artificial Intelligence* John Wiley & Sons, 1994.

> By allowing us to view the life that has evolved here on earth in the larger context of *possible* life, we may begin to derive a truly general theoretical biology capable of making universal statements about life wherever it may be found and whatever it may be made of. [24, preface, p.*xvi*, original emphasis].

Central to Langton's argument is the notion that 'life' is a property of the *organisation* of matter, and not a property of the matter which is so organised to form living systems [24, introduction, p.2]. This belief allows for the formal study of living systems in the abstract, as a complement to the actual living systems we are familiar with from life on earth. Clearly, there is a heavy dependence here on notions of 'Life' and 'Lifelike', which require some discussion: it would seem that any attempt at evaluating A-Life research requires first an answer to the question 'What is Life?'.

Not surprisingly, this question has long been studied by philosophers of biology, and no definitive answer has yet been devised. Nevertheless, it is intriguing to note the parallels between Langton's distinction of life-as-we-know-it vs. life-as-it-could-be and the following passage taken from the 1975 edition of Helena Curtis's classic biology textbook:

> If we were to be transported through time or space in the search for "life", what would we look for? Scientists concerned with this question, whether for practical or philosophical reasons, appear to agree that there are no simple answers.
>
> ... Even within the confines of our biosphere, ... [living] organisms show astonishing variety. Are these highly varied living organisms distinguishable from nonliving systems? *Can we make distinctions that would apply to unknown living forms as well as to known ones?* [14, p.26, emphasis added]

Curtis [14, pp.27–31] goes on to provide a list of some key characteristics exhibited by living systems, including: that typically they are complex and highly organised; that they take energy from the environment and change it to other forms; that they are homeostatic, thereby ensuring the stability of their highly organised complexity; they respond to stimuli; mostly they reproduce, grow and develop; they are adapted, i.e. they are suited to their environment and the functions required of them; and the information by which they organise and maintain all of these features is contained within the individual organisms.

Synthetic systems exhibiting such characteristics are studied within A-Life research. For instance, the proceedings of the second international A-Life meeting [25] are divided into the following categories: artificial chemistries, self-organisation, and the origin of life; evolutionary dynamics (i.e. studying the *process* of evolution rather than the *results* of evolutionary processes); development (especially morphogenesis); evolution, learning, and communication; computer life (e.g. software viruses); and philosophical issues.

Clearly, A-Life research encompasses a very broad range of phenomena. The scope of A-Life is *so* broad that the relevance to AI may not be immediately clear: it might appear that, at best, those phenomena studied in AI are a very small subset of the class of 'lifelike behaviors', in virtue of the fact that (so far) the only existent intelligent entities also happen to be living things.

Nevertheless, the broad-scale biologically-inspired work in A-Life *can* be illuminating to AI research: for much of its history, AI has typically been concerned with the activities of just one species: *Homo sapiens*, and notions of biological feasibility in AI research have tended to focus on comparatively narrow architectural issues. There is a new scientific field, which like AI explicitly addresses issues of cognition and intelligence, and like A-Life takes a broad biological perspective: this is the field of *Adaptive Behavior* research. The motivations for and implications of this new style of work are discussed further below.

2 And what is AI?

Artificial Intelligence (AI) is commonly defined as the scientific endeavour of trying to make computers perform tasks which, if performed by humans, would require intelligence. But what is 'intelligence'? Again, this is a question that has troubled philosophers for hundreds of years, yet there is still no general agreement on a definition

For most of its history, work in AI has been based on the assumption that intelligence is a form of *computation* that takes place in people's brains, involving the manipulation of symbolic representations of 'facts' or 'knowledge', in a manner similar to mathematical logic. It was assumed that perceptual systems delivered symbolic representations to general-purpose reasoning mechanisms, which would in turn specify appropriate actions. To limit potential combinatoric explosions, canonical-form objective representation techniques were assumed desirable.[1]

One consequence of this view of intelligence is that, because computation is an abstract process, the details of how 'intelligent computations' are actually *implemented* are rather irrelevant. In particular, for many years, most AI research didn't care too much about how well the AI computer programs corresponded with the actual biological mechanisms operating in the brains of 'intelligent' animals such as ourselves.

At least in part, this view of intelligence led AI researchers to concentrate on advanced human-level intellectual abilities, such as language use, planning complex sequences of tasks, or learning and applying "expert" knowledge for tasks such as mineral prospecting or diagnosing blood diseases. Several research groups have had notable success in making computers perform such tasks, and these successes added support to the belief that the techniques of traditional symbolic AI could be expanded and built upon, until eventually an artefact might one day pass the Turing test.

In recent years, these traditional views have been challenged. A small (but growing) number of researchers have argued that, for AI research to be conducted on firmer foundations, it should be more strongly integrated with biology. It turns out that, when a more rigorously biological approach is applied, many of the assumptions of traditional AI have to be reviewed.

[1] Brevity requires an element of caricature, but these assumptions are clearly expounded in the work of Newell and Simon [34], figured strongly in the work of David Marr and his followers (e.g. [27, 28]), and are clearly exhibited in most undergraduate AI textbooks I am familiar with (e.g. [35, 43, 8]).

3 Biological Issues: Adaptive Behavior

Appeals to biology are not new: in the 1960's [36], and again since the mid-1980's [37, 29, 1], there has been widespread interest within AI in so-called 'artificial neural networks'. These artificial networks are motivated primarily by physiological observations of the brain and other parts of the nervous system. The major influence comes from the observation that the nervous system is composed of large numbers of very simple 'processing units' (i.e. *neurons*) which are richly interconnected and all operate in parallel. Because the overall performance of the network is often the result of the *interactions* between the neurons (that is, there is no clear 'division of labour' with particular individual neurons doing particular specialised jobs), artificial neural networks are therefore examples of *parallel distributed processing.*

However, most artificial neural network researchers still treat the nervous system as a computing device, and still try to tackle the same classes of problems as are tackled in traditional AI. While such networks have interesting properties, they are commonly built without reference to any wider biological context, and (crucially) it is the wider biological perspective that offers realistic prospects of understanding intelligence.

This perspective involves going beyond physiology, and viewing the 'intelligent' activities of nervous systems within the contexts of Ethology, Ecology, and Evolution:

Ethology: Animals have nervous systems that give rise to particular patterns of *behavior*, some of which we would like to describe as 'intelligent'. The study of animal behavior is known as *Ethology*, and the ultimate success or failure of a particular nervous systems is dependent on whether it produces sensible behaviors: running away from predators, or moving towards food, are both generally sensible things to do; running towards predators, or away from food, are not. In traditional AI research, there is a strong tendency to try to identify abstract computations which could be general-purpose mechanisms of intelligence, without worrying about how long it might take to gather and process all the necessary information. But in the real world, animals very rarely have the luxury of enough spare time that they can indulge in the sort of computations traditional AI has studied. The need to perform the right behavior at the right time is often so urgent that any method will do, so long as it works fast enough. In such situations, questions of whether the method is general-purpose, elegant, or optimal, are often irrelevant.

Ecology: Deciding if a particular behavior is sensible or not depends not only on the immediate circumstances of the animal, but also on that animal's *ecological niche*, i.e. the animal's interactions with its habitat and with other life forms including predators, prey, and members of the same species. Many behaviors are highly tailored to particular environments and ecological niches: often, trying to analyse the animal as *separate* from its environment obscures more than it clarifies; the animal and its behaviors are so finely tuned to its natural environment that if it is separated from that environment (e.g. by transporting it to a laboratory where it can be tested), then very little of the original interesting behavior remains.

Evolution: For any animal, the nervous system, behaviors, and ecological niche, are all subject to *evolutionary* processes, such as natural selection ("survival-of-the-fittest") or sexual selection (mate-choice). Almost always, the evolutionary history of an animal plays a very important role in explaining the animal's behaviors or physiology. Furthermore, it is important to remember that evolution proceeds mainly by a continuous process of "tinkering" or making minor changes to established patterns. For this reason, it is important to remember that designs found in biology aren't necessarily the best possible solutions: they are adapted versions of earlier designs, and the earlier designs themselves may not have been particularly wonderful.

A primary claim here is that AI research probably does have much to gain from looking to biology for inspiration, but it is essential to remember that *biology doesn't stop at physiology*: many important lessons are likely to be learned from evolutionary, ecological, and ethological analyses. One of the most important lessons could be that traditional AI, working towards elegant general-purpose computational mechanisms underlying intelligence, is simply misguided. After all, when did you last see a general-purpose animal?

One of the advantages of taking account of these established areas of biological research is that the terminology of such fields offers more concrete definitions of some phenomena which are within the domain of study of traditional AI. One good example of this is the move away from arguing about the definition of "intelligence", towards an agreement in the new approach that the interesting phenomena are *adaptive behaviors*. In the ethology literature, an adaptive behavior is any behavior which, if exhibited by an animal, increases the chance that the animal will survive long enough in its ecological niche to produce viable offspring. Underlying this definition is the assumption that, if the animal does nothing, it will die before it has a chance to reproduce.

In essence, the implication is that "intelligence" is a name we give to a class of adaptive behaviors, all of which are rooted in being able to coordinate our perceptions (seeing, hearing, etc) with our actions, so as to survive in environments which are often hostile, uncertain, and unforgiving. So intelligence can perhaps be best understood as something which makes us better at satisfying our fundamental drives. These fundamental drives are commonly known as the Four F's: feeding, fighting, fleeing, and reproduction. One of the claims of this new style of AI research is that, if we can study and understand the basic mechanisms by which we and other animals satisfy these drives, then we will have gone a very long way toward understanding the biological foundations of intelligent activity.

4 Implications: Autonomous Agents

In order to study adaptive behaviors, it is necessary to conduct research in a manner different from that found in most traditional AI research. Rather than working on computer programs that appear to mimic some limited aspect of high-level human intelligence, or effect some supposedly vital transformation between internal representations, the new approach concentrates instead on studying *complete autonomous agents*. An autonomous agent is any self-governing system which is capable of coordinating perceptions and actions to produce adaptive behaviors, without human intervention, for extended periods of time. Nature is full of autonomous agents: we call them animals. However, *artificial* autonomous agents are much rarer things. Nevertheless they have a name: the word 'animat' was coined for them by Stewart Wilson [42].

Animats may be real physical things, in which case they are typically autonomous mobile robots, or they may be simulated on a computer, where they go about their business in some 'virtual reality'; either way, there are difficult problems to be faced in constructing whole animats. The strong emphasis on building *complete* animats is a consequence of needing to ensure that the agents really are autonomous: if perception is studied in isolation from action (or vice versa), there is a much increased risk of failing to address important tasks by assuming they will be dealt with elsewhere in the agent, or wasting resources on tasks that would be unnecessary in a complete system: both are common problems in traditional AI systems.

Although traditional AI has concentrated on advanced human-level behaviors, it would at the moment be too ambitious to try to build a complete animat with the capabilities of a human or other ape. Nevertheless, this is no reason for abandoning hope of progress. Valuable lessons can be learned from studying autonomous agents slightly less complex than apes: for instance, *insects* are pretty good at being autonomous agents.

The view that AI should proceed by studying insects is strongly associated with the work of Rodney Brooks and his colleagues at MIT. His views are probably best expressed in his own words:[2]

> Insects are not usually thought of as intelligent. However, they are very robust devices. They operate in a dynamic world, carrying out a number of complex tasks... No human-built systems are remotely as reliable... Thus I see insect level behavior as a noble goal for artificial intelligence practitioners. I believe it is closer to the ultimate right track than are the higher level goals now being pursued. [5, p.7]

So, perhaps AI should stand for "Artificial Insects" rather than "Artificial Intelligence". Certainly, building or simulating insect-like animats presents a number of very challenging problems, and many of these are problems traditionally studied in AI and Cognitive Science. For example, there are many AI researchers working on artificial neural networks that imitate some aspect of human language processing, even though the neural mechanisms underlying the "dance-language" of bees are not yet fully understood by biologists. If we can't yet understand how the bee's nervous system allows for communication, then surely the prospects of understanding the (much larger and probably more complex) neural networks which enable humans to use language are not good. Maybe if we understood the neural basis of language in bees first, we would find some clues for understanding the neural basis of language in humans.

Moreover, it is important to note that the animat approach to AI isn't restricted solely to the study of intelligence in the natural world: the notions borrowed from ethology, of intelligence as adaptive behavior, can also be applied to purely artificial systems: David McFarland, an ethologist at Oxford, has developed [30] a direct mathematical analogy

[2]For further details of Brooks' work, see [7, 6].

between adaptive behaviors in animals and adaptive behaviors in robots (e.g. the research-and-development time for a robot is analogous to the gestation period of an animal); and Mark Tilden, currently at Los Alamos, has proposed [39] three laws of robotics inspired by studies of animals. Tilden's three laws, to be obeyed by any interesting robot, can be stated as: protect yourself (i.e. avoid damage); 'feed' yourself (i.e. try to keep enough power available for running the robot's electronics and motors, etc); and find a better place to be (e.g. try to find a less dangerous place with more 'food'). Any agent obeying these three laws while operating in the same world as we humans (e.g. a world populated by dogs, inquisitive children, and clumsy adults) will have solved most of the hard problems of coordinating perception and action. In certain circumstances, such an agent may even be judged by observers to be 'intelligent', whatever that means.

So the adaptive behavior approach to AI, as a consequence of paying more attention to biology, takes a path that is almost the opposite of the traditional work in the field: instead of worrying about how humans use expert knowledge or plan complex tasks, researchers are now studying basic competences, such as an individual agent being able to explore an unknown environment without bumping into things and still make it home safely, or a group of agents being able to communicate using grunts so they can act in a cooperative manner. Such tasks require robust sensory-motor coordination, and can be surprisingly difficult. Other researchers have taken particular insects or other "lowly" animals and performed extensive computer modelling experiments, as attempts at identifying the basic neural sensory-motor mechanisms underlying the generation of adaptive behaviors in these animals. Examples include studies of cockroaches [4], crickets [41], eels [16], frogs and toads [2], houseflies [17], hoverflies [9], and stick-insects [33]. In some of these cases, working robots have been built as the final test of successful computer modelling.[3]

The aim of such research is not only to further our understanding of animals in the natural world: there is a huge potential sales market for semi-intelligent highly autonomous robots. At the moment, because such robots are only ever built in small numbers, they cost quite a lot of money (e.g. over U.S.$100,000 for an insect-like six-legged walking robot about 1 metre long and weighing a few kilograms). Such expensive price-tags mean that they are only cost-effective for use in space missions or military applications. But if mass-produced, they would probably cost less than a family saloon car, and could possibly eventually cost no more than the price of a home video-cassette recorder, in which case such robots could realistically be employed for helping in household cleaning duties (e.g. [13]), security patrols (not guard-dogs but guard-animats), or even as "pet robot" toys.

It should be clear that artificial autonomous agents will need to be able to sense their surroundings, using mixtures of vision, sound, 'touch' (in the form of whiskers or bumpers), and possibly also less obvious techniques such as 'smell' (e.g. chemical sensors that detect dangerous gases), electronic compasses, and infra-red or pyroelectric sensors that allow the robot to detect warm bodies. In biological terminology, such environment-sensing capabilities are referred to as *exteroception*. However, it is also essential that the

[3] For further details of these and other examples, see e.g. [26, 32, 40, 31, 12] and the MIT Press journal *Adaptive Behavior*.

agents have a capability for *interoception*: that is, the capability to monitor important internal values such as the level of power in their batteries, and levels of damage or mechanical strain, and so on. All of these diverse sources of information need to be dealt with and integrated in order to generate appropriate behaviors *in real time* (that is, fast enough for the behaviors to be considered adaptive in real-world tasks). Moreover, for real robots, it is practically impossible to build sensors and motors that are not affected by *noise*, in the form of random variations in response, which introduces uncertainties about the accuracy of the measurements made in both exteroception and interoception. Currently the most suitable candidate technology for dealing with all these issues in the generation of adaptive behaviors is probably still artificial neural networks, or some similar parallel distributed processing architecture.

5 Blocksworlds Revisited?

One criticism of animat-style adaptive behavior research, often voiced by people unfamiliar with the field, goes something like this: all the successful autonomous agent research results seems to concern specialised systems operating in restricted domains; and this is disturbingly reminiscent of the "blocksworld" research paradigm popular in mainstream AI during the first half of the 1970's. In this paradigm, specially contrived and simplified ("micro-world") problem domains were used for testing supposedly general or extensible AI representation and reasoning mechanisms — in several notable cases, mechanisms which were encouragingly successful in their toy test domains failed miserably in more complex domains. Almost always, combinatoric explosion caused significant problems, and in several cases this was compounded by the 'general' mechanism having hidden or implicit specialisations to the toy domain. At best, so the argument goes, all that was achieved was a collection of ad-hoc solutions to fairytale problems. (See [15] for further discussion).

Such arguments have a number of problems. First, they often seem to overlook the fact that, throughout the history of science (or since Galileo at least), the study of simple restricted systems as precursors to more general and complex systems has been a principled and successful approach. The problem with micro-world AI research was not due just to the use of restricted domains. Rather, there were problems in the *nature* of the restrictions and the claims made for the extensibility of the approach. The nature of the restrictions were problematic because (as Dreyfus [15] points out) in several cases researchers assumed that it would be possible to carve out isolated domains, such as a micro-theory of flattery, which would be understandable in isolation from the rest of human existence. And the claims made of the future of micro-worlds centered on a belief that, if sufficient work was done, all of the micro-worlds would coalesce to form a single super-world, corresponding to the objective reality of the "real world"; with the combined expertise from the different domains coming close to resembling human intelligence. Results over the intervening years have not offered much support for this belief.

The difference between AI and Adaptive Behavior approaches to simplicity and specialisation should be fairly clear. First, however simple, animats are typically complete 'holistic' systems, operating under closed-loop conditions in their own environments, which provide them with their own subjective reality (however simple that may happen to be). So, under this view, *all* worlds are 'micro-worlds', insofar as there is no sensible way in which they can be combined to recreate a single overall objective reality. Even we humans exist in a micro-world, albeit a highly complex one. Consider the human visual world: however rich we find it, we have to acknowledge that we cannot distinguish the characters on a car registration-plate at a distance of one mile, whereas other animals have sufficient visual resolution to do so; nor can we see the ultraviolet and near-infra-red wavelengths, or detect polarisation orientations, which other animals use to great effect in their visual worlds [38]. Once again, it is currently a matter of belief whether successes in simple (i.e. low-complexity) domains will scale up to more complex environments, but empirical answers should be more readily available than they were for AI micro-world studies.

Furthermore, as was noted earlier in this paper: the fact that all animals are (to varying degrees) specialised to their ecological niche offers support for the notion that *exploitation* of niche-specialisation is a profitable approach to developing artificial autonomous agents. Rather than attempting to develop general-purpose domain-independent solutions, the animat approach could be characterised as seeking to develop powerful techniques for creating domain-specific solutions in a principled manner. For example, Horswill [20] has developed a method for creating specialised autonomous robot control architectures from more general ones, where the change from the general to the specialised is achieved by means of provably performance-improving transformations, some of which are dependent on particular environmental regularities. Once a specialised controller has been arrived at, the sequence of transformations applied in its development can be used to make explicit any assumptions about the robot's environment, and to predict how it will perform in other environments.

Because most niches of practical interest are liable to change over time, forms of plasticity and adaptation are necessary, so as to maintain a satisfactory degree of specificity: this need is reflected in the large research efforts currently being directed at various adaptation mechanisms based on self-organisation, reinforcement learning, and evolutionary adaptation in the form of genetic algorithms (see e.g. [23, 3, 22, 21, 19, 18]). Specialised architectures arrived at by these methods can still be analysed and understood (see e.g. [10, 11]).

6 Conclusion

It is tempting to draw one more comparison between biology and AI. This concerns the history of the theory of the origin of the species. According to [14], pre-evolutionary theories centered on the work of Georges Cuvier, who was a vociferous and influential opponent of evolutionary theory. Cuvier was a founder of paleontology, and his studies of the fossil record led him to recognise the extinctions of species. Such extinction events were explained by positing a series of *catastrophies*: God populated the world with many varied and wonderful species, and occasionally a catastrophe would occur (the most recent being the biblical flood), which rendered many species extinct; and the surviving species repopulated the planet. The proponents of catastrophism were divided into two camps: the *deluvianists*, who believed all catastrophes were floods, and the *vulcanists*, who believed all catastrophes were due to volcanic explosions and lava flows. Members of both camps agreed that the fossil record shows the remains of once-living species which met violent ends, but there was vituperative debate between the deluvianists and the vulcanists. Nevertheless, both camps rose as one against Lamarck, who proposed the theory of evolution. Unfortunately, Lamarck's theory included the inheritance of acquired characteristics, and posited that all animals had a god-given in-built urge to evolve to higher things (i.e. to approach the 'perfection' seen in *Homo sapiens*). Darwin's revolutionary contribution was the suggestion that the combination of chance variation and natural selection could account for the evolution of species, without the intervention of any god.

This story could bear a resemblance to the history of AI. If we view catastrophism as analogous to computationalism, then the deluvianist/vulcanist debate becomes analogous to the symbolism/sub-symbolism debate that has been ongoing over the last decade. The fruit of this comparison is that the catastrophists were right, in that extinctions *do* occur and *do* affect the origin of species (by freeing or creating new niches), but focusing purely on extinctions missed the wider issues. The deluvianists and the vulcanists were both partly right, but not to the exclusion of the opposing view. Similarly, computational perspectives on intelligence are unlikely to be proved entirely wrong, and both symbolic and subsymbolic accounts have a role to play. But, just as Lamarck and then Darwin widened the debate, so A-Life and Adaptive Behavior widen our perspectives on intelligence and cognition. And, as with any other science, we should allow A-Life and Adaptive Behavior to take some time in refining concepts, explanations, and techniques.

It is probably too early to reliably evaluate the impact of A-Life and Adaptive Behavior research on AI. Depending on who you listen to, the impact will either be a damp squib drowned out by the combinatorial explosion which usually follows from early successes on simple examples, or it will provoke a revolution of Copernican or Darwinian proportions. As always, what actually happens is likely to be somewhere in between. Time will tell.

References

[1] J. A. Anderson and E. Rosenfield, editors. *Neurocomputing: Foundations of research.* MIT Press, Cambridge MA, 1988.

[2] M. A. Arbib. Levels of modelling of mechanisms of visually guided behaviour. *The Behavioral and Brain Sciences*, 10:407–465, 1987.

[3] H. G. Barrow. Learning receptive fields. In *IEEE First Annual International Conference on Neural Networks*, volume IV, pages 115–121, June 1987.

[4] R. D. Beer. *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology.* Academic Press, 1990.

[5] R. A. Brooks. Achieving artificial intelligence through building robots. A.I. Memo 899, M.I.T. A.I. Lab, May 1986.

[6] R. A. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 139–159, San Mateo, California, 1991. Morgan Kaufmann.

[7] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[8] E. Charniak and D. McDermott. *Introduction to Artificial Intelligence*. Addison Wesley, 1985.

[9] D. Cliff. Neural networks for visual tracking in an artificial fly. In F. J. Varela and P. Bourgine, editors, *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, pages 78–87. MIT Press Bradford Books, Cambridge, MA, 1992.

[10] D. Cliff and S. G. Bullock. Adding 'foveal vision' to Wilson's animat. *Adaptive Behavior*, 2(1):47–70, 1993. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP263.

[11] D. Cliff, I. Harvey, and P. Husbands. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):71–108, 1993.

[12] D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors. *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*, Cambridge MA, 1994. M.I.T. Press Bradford Books.

[13] J. H. Connell. *Minimalist Mobile Robotics: A Colony-style Architecture for an Artificial Creature*. Academic Press, 1990.

[14] H. Curtis. *Biology*. Worth Publishers Inc., New York, second edition, 1975.

[15] H. L. Dreyfus. From micro-worlds to knowledge representation: AI at an impasse. In J. Haugeland, editor, *Mind Design: Philosophy, Psychology, Artificial Intelligence*, pages 161–204. M.I.T. Press Bradford Books, Cambridge MA, 1981.

[16] O. Ekeberg. *Computer Simulation Techniques in the Study of Neural Systems*. PhD thesis, Royal Institute of Technology, Department of Numerical Analysis and Computing Science, Stockholm, Sweden, 1992.

[17] N. Franceschini, J.-M. Pichon, and C. Blanes. From insect vision to robot vision. *Philosophical Transactions of the Royal Society of London B*, 337(1281):283–294, 1992.

[18] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley, 1989.

[19] J. H. Holland. *Adaptation in Natural and Artificial Systems.* University of Michigan Press, Ann Arbour, 1975.

[20] I. D. Horswill. *Specialization of Perceptual Processes.* PhD thesis, MIT AI Lab, May 1993.

[21] L. P. Kaelbling. *Learning in Embedded Systems.* MIT Press Bradford Books, 1993.

[22] A. H. Klopf. A neuronal model of classical conditioning. *Psychobiology*, 16(2):85–125, 1988.

[23] T. Kohonen. *Self-Organisation and Associative Memory.* Springer-Verlag, 1984.

[24] C. G. Langton, editor. *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems, Los Alamos, NM, Sept, 1987.* Addison-Wesley: volume VI in the series of the Santa Fe Institute Studies in the Sciences of Complexity, 1989.

[25] C. G. Langton, J. D. Farmer, S. Rasmussen, and C. Taylor, editors. *Artificial Life II: Proceedings Volume of Santa Fe Conference Feb. 1990.* Addison Wesley: volume XI in the series of the Santa Fe Institute Studies in the Sciences of Complexity, 1991.

[26] P. Maes, editor. *Designing autonomous agents: theory and practice from biology to engineering and back.* MIT/Elsevier, Cambridge MA, 1990.

[27] D. Marr. Artificial intelligence: A personal view. *Artificial Intelligence*, 9:37–48, 1977.

[28] D. Marr. *Vision.* W. H. Freeman, New York, 1982.

[29] J. L. McClelland and D. E. Rumelhart, editors. *Parallel Distributed Processing. Volume 2: Psychological and biological models.* MIT Press Bradford Books, Cambridge MA, 1988.

[30] D. McFarland. What it means for robot behaviours to be adaptive. In J.-A. Meyer and S. W. Wilson, editors, *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90)*, pages 22–28, Cambridge MA, 1990. MIT Press Bradford Books.

[31] J.-A. Meyer, H. Roitbalt, and S. W. Wilson, editors. *From Animals to Animats: Proceedings of the Second International Conference on Simulation of Adaptive Behavior (SAB92).* MIT Press Bradford Books, Cambridge MA, 1993.

[32] J.-A. Meyer and S. W. Wilson, editors. *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90).* MIT Press Bradford Books, Cambridge MA, 1991.

[33] U. Muller-Wilm, J. Dean, H. Cruse, H. Weidemann, J. Eltze, and F. Pfeiffer. Kinematic model of a stick insect as an example of a six-legged walking system. *Adaptive Behavior*, 1(2):155–170, 1992.

[34] A. Newell and H. A. Simon. Computer science as empirical enquiry: Symbols and search. *Communications of the Association for Computing Machinery*, 19, March 1976.

[35] N. J. Nilsson. *Principles of Artificial Intelligence.* Springer Verlag, 1982.

[36] F. Rosenblatt. *Principles of Neurodynamics.* Spartan, New York, 1962.

[37] D. E. Rumelhart and J. L. McClelland, editors. *Parallel Distributed Processing, Volume 1: Foundations.* M.I.T. Press Bradford Books, Cambridge MA, 1986.

[38] J. Stone. Evolutionary robotics: our hands in their brains? In R. Brooks and P. Maes, editors, *Proceedings A-Life IV*, Cambridge, MA, 1994. MIT Press. In press.

[39] M. W. Tilden. The second international BEAM robot olympics and micro-mouse competition: Event rules and general guidelines, July 1992. Contact mwtilden@watmath.uwaterloo.ca.

[40] F. J. Varela and P. Bourgine, editors. *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91).* M.I.T. Press Bradford Books, Cambridge MA, 1992.

[41] B. Webb. Modeling biological behavior, or 'dumb animals and stupid robots'. In *Preprint Proceedings of the Second International Conference on Artificial LIfe (ECAL93)*, pages 1090–1103, 1993.

[42] S. W. Wilson. Knowledge growth in an artificial animal. In J. J. Grefenstette, editor, *Proceedings of an International Conference on Genetic Algorithms and their Applications (ICGA85)*, pages 16–23, Hillsdale: New Jersey, July 1985. Lawrence Erlbaum Associates.

[43] P. H. Winston. *Artificial Intelligence.* Addison-Wesley Publishing Co. Inc., second edition, 1984.