



Biologically-Inspired Computing Approaches To Cognitive Systems: a partial tour of the literature

Dave Cliff
Digital Media Systems Laboratory
HP Laboratories Bristol
HPL-2003-11
April 14th, 2003*

E-mail: dave.cliff@hp.com

cognitive
systems,
biologically-
inspired
computing,
artificial life,
artificial
intelligence,
autonomous
agents

This paper presents a review of the academic literature on biologically-inspired computing approaches to the science and engineering of cognitive systems. This review is intended as a rapid tour through the area (rather than a leisurely wander); and it should be readable in a few hours. The tour is *partial* in both senses of the word: it is only *partially* complete, and it is biased (i.e., it is not an *impartial* tour). Originally written in 1997, but revised and updated since, its primary intention is as background reading for undergraduate and postgraduate teaching in cognitive science and artificial intelligence. Its content is based on courses taught by the author at the University of Sussex (UK); at MIT (USA); and at the University of Southampton (UK).

Contents

1	Introduction.....	3
2	Background Literature	6
2.1	Prehistory	6
2.2	Artificial Life.....	7
2.3	Artificial Neural Networks	11
2.4	Artificial Autonomous Agents and Adaptive Behavior	17
2.5	Evolutionary Computation	23
2.6	Artificial Biochemistries.....	31
3	Applications of Artificial Autonomous Cognitive Agents	34
3.1	State of the Art.....	34
3.1.1	Autonomous Robots	34
3.1.2	Software Agents.....	36
3.2	Future Research Issues and Application Areas.....	39
4	References	42

1 Introduction

Broadly speaking, the study of cognitive systems can be pursued for purposes of *science* or for purposes of *engineering*. Although this is a very coarse division, it can be useful to distinguish between these two purposes. A scientific study will seek to establish novel and falsifiable statements about naturally-occurring cognitive systems (or perhaps about the abstract space of *possible* naturally-occurring cognitive systems), while an engineering approach seeks simply to create artificial systems that reliably exhibit some desired level of cognitive performance or behaviour. Of course, the two approaches are not mutually exclusive, and in principle each can inform the other.

As all known naturally-occurring cognitive systems are biological in origin, scientific studies of cognitive systems ignore biological data at their peril; but for engineering there is no *a priori* requirement to pay attention to biology. The majority of research directed at the engineering approach to creating cognitive systems was conducted under the banner of *artificial intelligence*, a field strongly (but not uniquely) identified with computer science. For much of its history, artificial intelligence (AI) research largely ignored biology. From the outset (commonly identified as the 1956 Dartmouth Conference on AI: McCarthy *et al.*, 1955), up until the mid-1980's, the focus within AI research was almost exclusively on treating cognition as a process involving the manipulation of symbolic representations of facts or knowledge, using techniques inspired by mathematical logic and/or computer programming languages. The dominant view was that the job of sensory systems (such as vision or hearing) was to deliver abstract symbolic representations ("models") of the external world, which would then be reasoned about in order to determine an appropriate plan of action (e.g. Marr, 1982). Most studies of learning within AI were concerned with learning new ways to acquire, manipulate, relate, or classify symbolic representations of knowledge.

This symbolic logic-based approach to AI eclipsed prior biologically inspired approaches, such as early studies in cybernetics that treated networks of neurons as logical devices (McCulloch & Pitts, 1943), Rosenblatt's (1959) *Perceptron* artificial neural network, and the automated learning work of Widrow & Hoff (1960). The demonstration of minimal biologically inspired architectures for mobile robot controllers can be traced back to the cybernetics research of Walter's (1950) *turtle* robots and Ashby's (1952) *homeostat*, and attempts at synthesizing life-like phenomena in mechanical automata can be traced back to the mid-1700s (for a historical review see Langton, 1989).

In the mid-1980's a number of new approaches to the engineering of artificial cognitive systems emerged independently, but with a common theme of paying much more attention to biology. This was a sharp contrast to traditional symbolic logic-based AI, which (at best) treated issues of biological plausibility as irrelevant "implementation details". These new approaches included the renaissance of "sub-symbolic" *artificial neural networks* and *parallel distributed processing* (PDP) architectures, inspired by the nervous systems of animals; the development of "*behaviour-based*" systems as an explicit rejection of the knowledge-based approach, inspired in part by studies in ethology and behavioural ecology; and the rapid growth of research interest in the field of *artificial life*, which seeks to engineer artificial systems that draw inspiration from a diverse range of naturally-occurring systems including developmental biology, the human immune system,

evolving gene-pools, and interacting groups of autonomous agents. These topics are explored in more depth later in this document.

Common to all of these new approaches was the observation that many naturally-occurring systems, at one level of analysis, can be described as being built from components that are individually "simple" and that interact with each other in relatively "simple" ways; yet at another level of analysis these systems exhibit some "complex" overall behaviour that is not readily predictable from the individual components.

Typically the complex overall behaviour is the result of compounded non-linearities in the component interactions. Furthermore, many systems of interest exhibit sophisticated adaptation responses over multiple timescales, and are resilient with respect to variations in component connectivity -- often as a result of self-organization properties. Almost all such naturally-occurring systems are biological in origin, and the phrase *biologically-inspired complex adaptive systems* (BICAS) has been used to generically refer to all such artificial systems.¹

The ultimate aim of BICAS research is to understand and build artificial complex adaptive systems with the same attractive properties of adaptation and resilience and self-organization as are found in naturally-occurring complex adaptive systems. The relevance of the BICAS approach to the science and engineering of cognitive systems is manifest.

Some examples of naturally-occurring complex adaptive systems include:

- An individual nerve-cell (a neuron) can be described at one level of analysis as a simple component that integrates electrical impulses received on its inputs and, if the sum of impulses received within some time-period is sufficiently high, which generates an output impulse. This sounds simple, yet if enough neurons are connected together in the right way, and if this tangle of neurons is exposed to the right environment for sufficiently long, then the end result could be an adult human brain capable of thinking, learning, and acting.
- An individual animal can be viewed as a simple self-interested vehicle for propagating its genes: it just has to survive long enough to find a good mate and produce viable offspring, and its work is done. Yet over a sufficiently long period of time, random genetic variation combined with evolutionary processes such as Darwinian survival-of-the-fittest selection can create "designs" of animals that are exquisitely well-tailored to those animals' environmental niches.
- An individual trader in a marketplace can be viewed as a simple self-interested agent. Sellers try to trade at the highest price possible, while buyers try to trade at the lowest possible price. Yet, in the right conditions, this conflict between groups of traders acting out of naked self-interest can collectively form a market where transaction prices rapidly and repeatedly settle to an equilibrium price that represents the most efficient allocation of scarce resources, without the presence of a central coordinating or synchronising "auctioneer", despite dynamic variations in the underlying market supply and demand.

¹ It should be noted that some complex systems of significant interest are non-biological: for example, *simulated annealing* (e.g., Aarts & Korst, 1988) is a computerised optimisation technique inspired by the cooling of molten solids.

As the last example makes clear, the interpretation of “naturally-occurring” extends to social systems (which are, at root, biological). In all three cases there are local *small-scale* components and interactions that are relatively well understood and predictable; which compound to create global *large-scale* system behaviours that are (hopefully) desirable, but generally hard to predict in advance from knowledge of the small-scale characteristics.

A number of BICAS computing tools and techniques have been developed which have over the last 15 years (as a consequence of the dramatic falls in the real cost of processor power and memory and disk storage) shifted in status from academic curios to powerful methods used in critical industrial applications. Many BICAS techniques are directly applicable to the “bottom-up” engineering of artificial cognitive systems. Furthermore, the same falls in the real costs of computing have opened up the use of these BICAS tools and techniques for advanced computer modelling and simulation studies in the scientific understanding of natural cognitive systems. Thus, in contrast to the top-down knowledge-based AI approach, with this greater emphasis on biology comes a more realistic hope for genuinely productive interplay between the scientific and the engineering approaches in cognitive systems research.

The scope for potential applications is very broad, and we will not attempt a comprehensive review in this document. One promising general BICAS application area is in the construction of artificial *autonomous agents*. Autonomous agents are entities that are capable of co-ordinating perception and action, for extended periods of time, and without human intervention, in the pursuit of some set of goals. Biological autonomous agents are better known as animals: their “goals” are generally to stay alive long enough to mate. Artificial autonomous agents are animal-like artefacts: they may be physical mobile robots, or purely virtual entities. Autonomous mobile robots have many potential applications, from remote operation in hazardous environments, through to more mundane applications such as office cleaning or robotic toy “pets”. Virtual agents might be used as simulations of real agents within real environments: for example to predict the behaviours of animals or people in certain circumstances, perhaps to scientifically evaluate some hypothesis; or they may be the product of an artist’s imagination, finding uses in computer games and animations. Alternatively, virtual agents might sit in a “cyberspace” of electronic data, performing useful roles in business and industry (such as stock-market trading, controlling a manufacturing process, or as a plausible synthetic actor in an interactive education application). These (and other) applications may involve a single agent, or may involve groups of agents interacting either competitively or co-operatively. In all application areas, there are indications that the traditional (“von Neuman”) computer architecture that has served so well for so long may well not be the most applicable for studying or creating biologically-inspired complex adaptive systems: some new alternative approaches appear to offer greater promise.

Section 2 of this document surveys the background literature for BICAS approaches to the science and engineering of cognitive systems; and Section 3 discusses potential applications in more depth.

2 Background Literature

As was stated above, the general aim of developing techniques for computer simulation of intelligent or adaptive behaviours, creating systems capable of perception, learning, or language use, has traditionally been the domain of artificial intelligence (AI) research, first established as a distinct field in the late 1950's. In the latter half of the 1980's, a number of distinct new (or revived) styles of research started to attract significant attention, all which placed significant emphasis on biological inspiration. These included artificial life, artificial neural networks, artificial autonomous agents, artificial evolution, and artificial biochemistries; these are the topics that are reviewed here.

Section 2.1 gives a very brief indication of AI research prior to the mid-1980's, then Section 2.2 discusses artificial life. Following that, Sections 2.3, 2.4, 2.5, and 2.6 respectively discuss research in artificial neural networks, autonomous agents, artificial evolution, and artificial biochemistries. Each of these covers the historical background, core concepts, and gives pointers to further reading. Each of these sections reveals how metaphors and ideas from biological systems have influenced these computer science technologies. Finally in Section 2.7 we

2.1 Prehistory

Up until the mid-1980s, the vast majority of research in AI had concentrated on the development of computer systems that emulated or modelled high-level cognitive functions, often uniquely human functions such as language understanding and translation, forming representations of visual scenes suitable for generating linguistic descriptions from, or the application of degree-level knowledge in tasks such as mineral prospecting from geological data or diagnosing bacterial blood diseases from patient symptoms.

Much of this research, and almost all undergraduate AI textbooks of the time (such as those by Nilsson (1982), Winston (1984), and Charniak & McDermott (1985)) were based on the premise that intelligence is fundamentally concerned with the representation and manipulation of 'facts' or 'knowledge'. The representations and manipulations are often rooted in mathematical logic (e.g., Hodges, 1977). For example, given a database containing the two facts:

- '*Jo is taller than Mary*' represented as `tallerthan(jo,mary)`

and

- '*Mary is taller than Jane*' represented as `tallerthan(mary,jane)`

along with this general rule involving three variables A, B, and C:

- `'if (tallerthan(A,B) and tallerthan(B,C)) then tallerthan(A,C)`

logic-based reasoning allows the new fact `tallerthan(jo,jane)` (i.e. '*Jo is taller than Jane*') to be added to the database automatically. Symbol-manipulating and logic-based programming

languages such as LISP (Winston & Horn, 1980) and PROLOG (Clocksin & Mellish, 1984) were developed to facilitate the engineering of so-called “knowledge-based” or “rule-based” systems (e.g., Gonzalez & Dankel, 1993).

This style of AI research is commonly referred to as *symbolic AI*, a name that stems from a paper by Newell & Simon (1976). There, they argued for “computer science as empirical inquiry”, using computer programs to test the *Physical Symbol System Hypothesis*. This hypothesis was essentially a formal and precise statement of the assumptions that had been widespread but implicit in prior AI research: that systems which manipulated symbolic representations of knowledge exhibited the necessary and sufficient conditions for the generation of intelligent activity. For a review of the development of AI, relative to the other cognitive sciences (i.e., psychology, neuroscience, anthropology, linguistics, and philosophy), up to the mid-1980’s, see Gardner (1985).

Symbolic AI attracted some vociferous criticism, both on the grounds of philosophical objections (e.g., Searle, 1980; Dreyfus, 1979, 1981), and, as time progressed, on the grounds of a mismatch between ambitious promises and the actual results that came to pass (e.g., McDermott, 1981, 1987; Winograd & Flores, 1986; Leith, 1990). But research continued nonetheless. The premier journal in the field, *Artificial Intelligence*, continues to publish papers describing advances in symbolic search algorithms, or new developments in advanced logics for reasoning about time, uncertainties, monotonicities, and so on. However, little or any of this is of any consequence to the development of biologically-inspired approaches. Most work in AI prior to the mid-1980’s is largely of historical interest, and will be discussed no further.

2.2 Artificial Life

Artificial Life (or “A-Life”) emerged as a new scientific field in the late 1980’s. The first international workshop on A-Life, organised by Chris Langton, was held in September 1987, with the proceedings being published two years later (Langton, 1989b). Langton is widely acknowledged as the founder of the field, although he appears to have withdrawn from active research in recent years. The proceedings of that meeting commences with an introductory article that includes a historical review (Langton, 1989a), and concludes with an annotated bibliography of nearly 500 papers on prior research related to A-Life. Further international workshops (and, later, conferences) in this series were held biennially; the most recent being the meeting in Sydney in December 2002. In December 1991, the first biennial European Conference on Artificial Life (ECAL) was held in Paris: see Varela & Bourgine (1992). In 1993, MIT Press commenced publication of the international journal *Artificial Life*, the first three issues of which contained overview articles by members of the editorial board, giving their perspectives on past and current A-Life research. This collection of articles was subsequently published as a book (Langton, 1995).

A-Life rapidly achieved a high media profile, helped by the publication of a very readable popular science book devoted to it (Levy, 1993), and by the general increased interest in the notion of studying “complexity”, such as that exhibited by natural “complex systems” including economies and other social structures, brains, immune systems, interacting genes, and evolving species. The interdisciplinary Santa Fe Institute in New Mexico, USA, where Langton and several other prominent early A-Life workers held research posts, is widely viewed as the premier institution

for the study of complex systems. For excellent popular accounts of the genesis of complexity research, see Waldrop (1993) and Coveney & Highfield (1995).

Much work in A-Life concentrates on so-called *emergent* phenomena. This is a name that is given to complex but coherent global phenomena that arise from the interaction of many small constituents of a system, where the constituents are simple in relation to the global phenomena. Co-ordinated behaviors in groups of animals provide one example of emergence: flocks of birds, schools of fish, and herds of animals may all appear to move as a co-ordinated whole, while their behavior can be modelled (and perhaps explained) by each agent in the group following a small number of simple rules (e.g., Reynolds, 1987); the co-ordinated behavior of the group *emerges* from the interaction of these simple agent behaviors. As real systems are often too complex to allow analysis, many A-Life studies of emergent phenomena have concentrated on relatively simple formal systems such as one-dimensional cellular automata (e.g. Langton, 1991). See Forrest (1991) for a collection on emergent computation, a special issue of the journal *Physica D* reissued as a book. The journal *Complex Systems* has also acted as a forum for the publication of such research papers.

Despite this rapid growth in research activity and the attendant media coverage, formulating a precise definition of the scope of A-Life research appears to have caused some difficulties. This is made clear by the following quote from Langton's *Editor's Introduction* in the first issue of the first volume of the *Artificial Life* journal:

“The term artificial life literally means ‘life made by man rather than by nature’. As you will see, ‘artificial life’ means many things to many people, and I will not attempt to give a concise definition of it here, for in reality, artificial life is not yet ready to be constrained by quick and short definitions. Artificial life is still in the process of defining itself, as is proper for any new discipline. The articles in this initial volume carefully stake out claims to certain areas of study, but there is far more intellectual territory out there waiting to be discovered and laid claim to. Perhaps in ten years or so, we will be able to look back and provide brief and succinct characterizations of artificial life with the benefit of hindsight.”

(Langton, 1993b, page v).

This seems a curious statement. Six years after the first workshop, Langton was unwilling to attempt a definition of the field he founded, and suggested that it may take another decade before a definition could be arrived at.² There is a worrying lack of precision here, which will be discussed further below.

² That decade is now up. Looking at recent issues of the *Artificial Life* journal and the proceedings volumes of the recent *A-Life* and *ECAL* conferences, the range and nature of the topics studied do not appear to have changed significantly over the last ten years. If anything, the striking issue is the *lack* of additional “intellectual territory” laid claim to in the past decade. Encouragingly, the one topic-area showing growth is in papers reporting on technological applications of A-Life tools and techniques.

Certainly, Langton hedged less in his earlier writings, e.g.:

“Artificial Life is the study of man-made systems that exhibit behaviors characteristic of natural living systems. It complements the traditional biological sciences concerned with the *analysis* of living organisms by attempting to *synthesize* life-like behaviors within computers and other artificial media. By extending the empirical foundation upon which biology is based *beyond* the carbon-chain life that has evolved on Earth, Artificial Life can contribute to theoretical biology by locating *life-as-we-know-it* within the larger picture of *life-as-it-could-be*.”
(Langton, 1989a, p.1, original emphasis).

...

“Only when we are able to view *life-as-we-know-it* in the larger context of *life-as-it-could-be* will we really understand the nature of the beast. Artificial Life ... is a relatively new field employing a *synthetic* approach to the study of *life-as-it-could-be*. It views life as a property of the *organization* of matter, rather than a property of the matter which is so organised.”
(Langton, 1989a, p.2, original emphasis).

Thus defined, A-Life research involves the study of synthetic systems that exhibit ‘life-like’ behaviors (Langton, 1989a, p.5). The synthetic systems could be physical hardware, computer software, or biological ‘wetware’ such as laboratory preparations of molecules or cell cultures. A-Life research aims to study *life-as-it-could-be*, in order to advance our understanding of *life-as-we-know-it*. Rather than viewing ‘life’ as a property peculiar to carbon-chain chemistry, life is viewed as a property of the organisation of matter, no matter what the matter may be.

Under this view, it is possible that computer software agents, competing to survive and reproduce in some virtual environment, could be considered to be alive. In the late 1980’s, Thomas Ray (a professor of biology at the University of Delaware) argued forcefully that such software systems are not models or simulations of life on earth; rather, they are independent instances of life (Ray, 1994). To further explore this possibility, Ray developed a system called *Tierra*, a simple virtual environment where primitive software ‘agents’ could replicate, compete for limited resources, and (potentially) evolve. In a landmark paper, Ray (1992) described how, by seeding the *Tierra* environment with a simple self-replicating agent, a rich evolutionary process with complex dynamics was unleashed, giving rise to diverse groups (or ‘species’) of agents, including species that were parasitic.

Clearly, work such as Ray’s raises the philosophical question of “what is life?”. For Ray, an appropriately configured virtual environment is sufficient to sustain life. Other researchers have disagreed, but such disagreement is common at the leading edge of any scientific field. Papers by authors with varying positions on the philosophical issues raised by (and important to) A-Life were gathered in a collection edited by Boden (1996). Although at some point in the future there may be serious debate over the ethics of using synthetic living things in industrial applications (i.e. “animal rights” for *synthetic* animals), such questions are at present largely irrelevant to engineering uses of A-Life technologies. What matters is what they do, not whether they are alive.

The lack of precision in the definition of Artificial Life, and the emphasis in much A-Life research on abstract computational models of biological phenomena, have offered opportunities for criticisms of the field. Such criticisms have been further fuelled by the publication, in the proceedings of A-Life meetings, of papers that make extreme claims or predictions for the future of synthetic life-forms: see e.g., Moravec (1989) and Farmer & Belin (1992). An unsympathetic article published in *Scientific American* (Horgan, 1995) implied that much A-Life research bore little relevance to reality. Terms such as “eye-candy” (referring to the impressive, but possibly irrelevant, computer-graphic visualisations often produced to illustrate A-Life results), “fact-free science”, and “science-as-it-could-be” have been used by critics of the field. In a scholarly but somewhat damning paper, Miller (1995) forcefully argued that much A-Life research is poor-quality science, being biology research performed by mathematicians and computer scientists with little understanding of past research in, or current working practices of, the biological sciences. In essence, Miller’s argument is that if the aim of A-Life is to further our understanding of biological systems, then it should really be considered as a sub-branch of theoretical biology, in which case there is no need for a separate field with its own conferences and journals, other than as an outlet for publishing low-quality papers that would not stand the rigorous assessment and selection criteria of the established biology publications.

While proponents of A-Life would undoubtedly want to take issue with these criticisms, it seems undeniable that the lack of a precise definition of the field causes problems of focus. Definitions of the field based on the aim of replicating “life-like” phenomena in synthetic systems fall foul of the problem that almost any research in the life sciences, from molecular biology, through neuroscience and psychology, to evolutionary biology, population dynamics, and even social sciences such as economics and political science, can be viewed as studying “life-like” phenomena. Indeed, the logic-based theorem-proving of symbolic AI could be classed as “life-like” because theorem-proving is a phenomenon exhibited by humans, who are alive. For a more positive discussion of the A-Life roots of AI, see Steels (1993).

However, this lack of focus can be viewed charitably. Just as an ‘umbrella group’ is formed by representatives of small parties or organisations to act for them all where they have common interests, so A-Life has often acted as an ‘umbrella term’ for a variety of research interests which have not yet received sufficient attention or followers to establish their own workshops, conferences, and journals. Despite the often-valid criticisms levelled at A-Life as a field, the journal and proceedings volumes have undoubtedly provided a valuable forum for the publication of research papers that cross traditional subject-boundaries and that would otherwise encounter difficulties if attempts were made to publish them in better-established but more traditional outlets. While it is true that some of the papers published in the A-Life literature are of dubious scientific merit, the majority of work is of a standard just as high as in other scientific disciplines. Indeed, no other scientific field is entirely free from the problem of bad work being published by careless or mendacious researchers.

In the remaining subsections of this review, the literature on Artificial Neural Networks, Autonomous Agents, Artificial Evolution, and Artificial Biochemistries is examined. For the reasons just given, all of these could plausibly be considered as aspects of, or sub-fields within, Artificial Life.

2.3 Artificial Neural Networks

Without doubt, one of the most significant publications in the cognitive science literature of the 1980's was the 1986 two-volume collection reviewing developments in *Parallel Distributed Processing* (PDP), edited by Rumelhart and McClelland (Rumelhart & McClelland, 1986; McClelland & Rumelhart, 1986). PDP involved the revival of an unusual style of computational architecture, inspired by biological neural networks, where the combined activity of many simple processing units operating in parallel could give rise to an overall behavior that could be interpreted as a complex computation. In more traditional models, the computation would have been effected by a sequential program running on a standard (von Neuman architecture) single-processor computer.

In typical PDP models each of the simple processing elements take as input some number of scalar numeric values, and applies a nonlinear mathematical function to those input values, one which provides a single scalar output value. The nonlinear function typically has a sigmoidal ('S'-shaped) relationship between input and output. For example, in a system with binary output values, if the sum of input values received is less than some threshold, an output value of zero would be given by a unit; otherwise, the output is one. The input values for a unit could come from the external world (i.e., sensors or other input devices) or from other units within the PDP system. Similarly, the output value of a unit could form part of the output of the overall system, or could form part of the input to one or more other units within the system. The output value of a unit would be passed to form the input value of another unit via a weighted connection: the weight on the connection determines the influence one unit's output would have in deciding the output value of another unit; with differing connection weights, a particular single unit could have a large effect on some units while having only a small influence others.

This emphasis on a connected network of units, where the connection-strengths determined the nature of the computation performed by the network, gave rise to an alternative name for PDP systems: they are also often referred to as *connectionist* systems. In contrast to symbolic AI systems, the 'knowledge' in a PDP system is 'represented' in the relative strengths of the connections between units: a concept referred to as *distributed representation* (Hinton, McClelland, & Rumelhart 1986). A general mathematical framework for describing and analysing PDP systems was described by (Rumelhart, Hinton, and McClelland, 1986).

Different PDP researchers had different motivations for studying such networks. One strong motivation was the then-nascent development of massively parallel computer architectures such as the Connection Machine (Hillis, 1985), which escaped the 'von Neuman bottleneck' of requiring all processing to be executed by a single central processing unit. Another attractive consideration was that many PDP systems exhibit so-called 'graceful degradation': if connections or units are deleted from the system, or noise is added to the inputs or the internal weights or signals, the performance of a PDP network may decline gradually. This is in preference to the system suffering catastrophic failures such as would be expected from, for example, deleting lines of code in a conventional sequential computer program. Furthermore, some PDP systems exhibit smooth continuous transitions between states that would otherwise be discrete states in a symbolic AI system: for instance, Rumelhart, Smolensky, McClelland, & Hinton (1986) demonstrated a PDP network that combined inputs representing 'microfeatures' of different types of rooms, such as the

presence of a bed, a desk, a refrigerator, etc, to smoothly vary between global states representing 'bedroom' or 'kitchen' or 'office'. While a symbolic AI system could easily have represented these states by predicates such as `roomtype(room1, office)` or `roomtype(room2, kitchen)`, the PDP system was capable, via varying the combinations of input microfeatures, of passing through intermediate 'emergent' states that could be interpreted as analogues of 'bed-sit' or 'study-bedroom', which had not been explicitly coded into the system. This was something that would have been much more difficult for a symbolic AI system to achieve.

Finally, and perhaps most importantly, a number of mathematically rigorous automatic learning procedures had been developed, allowing the PDP systems to 'learn from experience' (Hinton, 1992). The availability of learning algorithms was important, as typically it is laborious or difficult or impossible to determine by hand the desired settings for the weights of the connections between the units, and the value of each unit's threshold. Learning algorithms allow networks to be initialised with random values for these parameters: subsequently, the learning algorithm changes the values of the weights and thresholds on the basis of results of a series of learning trials ('experience'), tuning the values until a desired input-output mapping is produced by the network. In the machine learning literature, learning algorithms are commonly divided into three classes: *supervised*, *reinforcement* (or *semi-supervised*), and *unsupervised*.

Supervised learning involves 'training' the system to replicate some given set of input-output mappings by presenting it with inputs taken from a pre-established *training set* of input-output pairs. Typically, each input-output pair will consist of a list of values to be supplied to the network's input units (the *input vector*), and a list of desired values required on the output units (the target *output vector*). The actual output vector produced by the network for a given input vector is compared with the target output vector, and the network parameters are altered in such a way that next time the same input vector is presented, differences (errors) between the actual and desired output vectors will be reduced. The intention is that the network not only learns to replicate the data in the training set, but also that it *generalises* so that it can produce desired output vectors for input vectors that it has not previously been exposed to. For this reason, it is also common to reserve a *test set* of input-output pairs, which the network is not exposed to during training. Once the network is performing well on the training set, its generalisation performance can be evaluated by recording its responses to the 'unseen' data in the test set. One of the simplest and most popular methods of supervised learning is the *back-propagation* algorithm (Rumelhart, Hinton, & Wilson, 1986; Werbos, 1994).

The other two classes of learning algorithm, reinforcement and unsupervised, differ from supervised learning in the style of feedback. Rather than employ a unit-by-unit comparison of values in the actual and target output vectors, using each of the individual errors, reinforcement learning operates on the basis of a single value, the *payoff*, which may be thought of as a method of giving the network 'reward' or 'punishment'; rigorous learning algorithms have been developed which attempt to maximise the expected future payoff received (e.g., Sutton & Barto, 1987; Sutton, 1988; Watkins, 1989; Watkins & Dyan, 1992). In unsupervised learning, there is no desired output or feedback to the network: typically there is a competitive process within the learning network such that certain output units win the right to be active in the presence of certain input vectors (e.g., Rumelhart & Zipser, 1986; Barlow, 1989). Often, unsupervised learning algorithms have strong links to multivariate statistical techniques such as principal components analysis: for further discussion of a statistical perspective on learning in artificial neural networks, see White

(1989); and for a good overview of analytic multivariate statistical techniques, see Chatfield & Collins (1980). Artificial neural networks have proven extremely useful in the development of more sophisticated multivariate techniques, with Independent Component Analysis receiving significant attention (for reviews, see Hyvarinen *et al.*, 2001; Stone, 2002).

Whatever the style of learning, for a network with n numeric parameters (e.g., unit-thresholds and connection-weights), the learning problem can be thought of as a directed search in the n -dimensional *parameter space* of thresholds and weights. At any one time, the parameters of the network have particular values, corresponding to a single *point* in that parameter space: after a single trial, if the learning algorithm adjusts some of the parameters, the network ‘moves’ to a new point in the parameter space. The result of a series of learning trials can thus be viewed as a *trajectory* through the n -dimensional parameter space, hopefully ending at a point where the parameter-values result in the network reproducing the desired input-output mapping.

A complementary way of viewing the learning process is to observe the errors between the target output-vectors and the actual output-vectors of the network. At different points in the network parameter space, there will be different amounts of errors. Assume that, for a given point in the network’s parameter space, the errors for all the vector pairs in the training set can be summarised by a single number.³ Then each point in the parameter space has an associated error value, and so an *error surface* has been defined over the parameter space. The task of learning can then be viewed as one of *gradient descent*; trying to move to nearby (or *local*) points in parameter space that are at lower altitudes on the error surface (i.e., give output vectors that are closer to the target vectors in the training set). Many learning algorithms, including back-propagation, work by attempting to identify the direction from the current point in network parameter space where the error surface has the steepest negative gradient, and moving the network’s parameter-values in that direction. The learning algorithm thus directs the system toward a *minimum* in the error surface. If the error surface has zero height at the bottom of the minimum, the training set of input-output mappings is being replicated perfectly. Unfortunately, real error surfaces are often plagued by *local minima*: points on the error surface that are lower than all nearby points, but still too high to yield acceptable network performance. For many learning algorithms, much research has been directed at trying to identify when the system is trapped in a local minima, and specifying appropriate ‘escape’ procedures (e.g., Hanson, 1991).

In addition to the computational advantages offered by PDP systems, for many researchers a much stronger motivation is the similarity between PDP network architectures and the nervous systems of animals. Many PDP systems are referred to as “artificial neural networks”. The similarities are manifest: a typical nerve cell (or *neuron*) receives electrical stimulation from other neurons. This electrical stimulation can (to a first approximation) be represented by a scalar numeric ‘input’ value: in the limit, a binary value can represent ‘stimulation’ or ‘no stimulation’. A neuron’s ‘output’, the electrical stimulation it provides to other cells, varies (typically in a nonlinear fashion) as a function of the ‘inputs’ received. Again, in the limit, a binary value could be used to distinguish between ‘stimulation’ and ‘no stimulation’. Neurons are typically classified

³ For example, in a system where the output vectors are of length k , the errors over the training set could be summarised by calculating the Euclidian (straight-line) distance in k -dimensional space between each actual output vector and its corresponding training output vector; and then calculating the root mean square of these distances for the entire training set to yield a single number.

as *sensory neurons*, as *motor neurons*, or as *interneurons*. An interneuron receives stimulation only from other neurons, and it only stimulates other neurons. This contrasts with sensory neurons, which produce 'output' electrical activity on the basis of stimulation from other physical phenomena, such as light, physical pressure, or chemical concentrations; and with motor neurons, the 'output' activity of which stimulates muscle cells to contract, producing physical forces. This has obvious similarities with the distinction made in artificial neural networks between 'input units', 'output units', and 'hidden units' (also referred to as 'internal units').

In many cases, the researchers are careful to present their work as computational models of cognitive processes *inspired by* natural nervous systems, or using *neuron-like* processing units. This caution acknowledges that real neurons are much more subtle and complex objects than typical PDP processing elements. Moreover, in the period that PDP was becoming established, the continuing rapid decline in the real cost of computing power led to an increase in the use of computationally intensive techniques to model biological neural networks via simulation; an approach referred to as *computational neuroscience* (e.g., Sejnowski, Koch, & Churchland, 1988). Computational neuroscience spans a range of approaches, from complex mathematical models of the electrodynamics of patches of neuronal cell membranes, through simulations of small cell assemblies (e.g., less than ten neurons) with less detail in each simulated neuron, to models of networks involving large numbers of coarsely simulated neurons. MacGregor (1987) described such a range of models, with accompanying program code; while the collection edited by Koch & Segev (1989) spans a similar range, ending with a model by Lehky & Sejnowski (1988) that takes results from employing standard PDP back-propagation learning to cast new light on the interpretation of neuroscience results such as Hubel and Weisel's Nobel-prize-winning work (e.g., Hubel, 1982, 1988).

Within artificial neural network research there are many types of systems and approaches, varying in their organisational principles, in the types of neuron used, in the degree of biological realism, and so on. Many of the learning algorithms require particular styles of interconnection. For instance, the networks popularised by Hopfield (1984, 1985) require symmetric connection patterns, and back-propagation learning requires that networks are *feed-forward*, i.e. have no internal-feedback (or *recurrent*) connections; often such restrictions are necessary to simplify the mathematics of deriving a learning algorithm. Furthermore, in many cases the number of units in the network has to be determined in advance. This is significant because the number of internal units affects the number of input-output mappings that can be learnt by a network. If the network has insufficient internal units, it will fail to reproduce the desired input-output mapping; if there are too many internal units (in relation to the number of input-output pairs in the training set) the network can *over-generalise*, essentially compiling a 'look-up table' of what output to produce for each individual input, and so generalisation to the test data is poor in such cases. In principle, the number of units can be altered 'virtually': if the learning algorithm sets the weights on all connections into and out of a unit to zero, that unit ceases to have an effect on the network and so can be considered 'deleted'; similarly if a unit with zero-valued weights on all inputs and outputs has its weight-values altered to non-zero values, it can be considered to have been 'added' to the network. Unfortunately, many learning algorithms give asymptotic weight change (i.e., the weights may approach zero but will never actually equal zero). For these reasons, some researchers developed *constructive* PDP learning algorithms that start with a small number of units (e.g., one), and gradually add extra units to achieve increases in performance (e.g., Fahlman & Lebiere, 1990).

In many applications, the network is constructed or trained to perform some task, but once in operation the learning process is suspended. There are two good reasons for this style of use. First, when in actual use, it is often no longer practicable to compare the network's actual output with the desired output (e.g. the desired output may not be specified for all possible inputs). Second, continued training of the network using its real-usage 'experiences' as training data might bias the network to learn to respond to frequently occurring inputs, thereby 'forgetting' the appropriate responses to rare but important inputs that were present in the training set.

One major problem with simple feed-forward networks is that they cannot maintain internal state: their current outputs are solely a function of their current input, and (once learning is complete) for any given input a particular network will always give the same output. Yet, in many practical applications, the same network may need to give different outputs for the same inputs, depending on the context of the network's immediate circumstances and of its past experience. For this reason, there has been increased effort to develop techniques for training and analysing *recurrent* networks (i.e., those with internal-feedback connections). Methods for employing back-propagation in simple recurrent networks were developed (e.g., Elman, 1990, 1993), as were constructive algorithms for recurrent networks (e.g., Fahlman, 1991). Of particular interest are so-called continuous-time recurrent neural networks (CTRNNs). These are more general versions of the networks studied by Hopfield (1984,1985), and can be viewed as simplified versions of the classic neuron model developed by Hodgkin & Huxley (1952). Reasons for the interest in CTRNNs is summarised well by Beer:

"In contrast to purely static feed-forward networks, which support only reactive behaviors, dynamical neural networks allow an agent to initiate action independent of its immediate situation and to organize its behavior in anticipation of future events...

"CTRNNs are an obvious choice for this type of work because (1) they are arguably the simplest nonlinear, continuous dynamical neural network model; (2) despite their simplicity they are universal dynamics approximators in the sense that, for any finite interval of time, CTRNNs can approximate the trajectories of any smooth dynamical system [of a certain large class]...; and (3) they have a plausible neurobiological interpretation, where the [internal state of a unit] is often associated with a nerve cell's mean membrane potential and the output ... is associated with its short-term average firing frequency."

(Beer, 1995b, p.470).

As Beer notes, many researchers interested in applying artificial evolution to neural networks (discussed further in Section 2.5) use CTRNNs. Because CTRNNs are capable of altering and maintaining internal state on the basis of inputs received, it is possible for them to exhibit 'learning' *without* altering the strengths of connections. Successive alterations in input value can alter the internal state of the network, tracing a trajectory of points in *state-space*. Often, the state-space trajectories traced out by the network will show regularities, such as convergence to *attractors* in state-space: either points, cyclic loops, or chaotic 'strange attractors'. Particular inputs or sequences of inputs may shift the state of the network from one attractor to another, in which

case the output behavior of the network might alter, and an external observer may want to say that the network appears to have 'learned' on the basis of the input(s) received. The roots of this approach lie in Hopfield's (1984, 1985) work on symmetrically-connected CTRNNs as associative memories, and there is now a large body of work on so-called *attractor networks*. The analytical tools and techniques of dynamical systems theory (e.g., Baker & Gollub, 1990; Abraham & Shaw, 1992) can be used to analyse such systems, to explain and predict their behavior. For discussion of (chaotic) dynamics in brain activity, see Skarda & Freeman (1987). Canavier et al. (1993) reported simulation studies in neurobiology which indicate that individual real neurons may have complex and perhaps chaotic dynamics. For demonstration of CTRNNs evolved to learn without weight change, see e.g. Yamauchi & Beer (1994), and for dynamical systems analysis of neural networks used to control robotic autonomous agents, see e.g. Husbands, Harvey & Cliff (1995) and Beer (1996). Gleick (1987) is a popular introduction to chaotic dynamics and strange attractors.

Although there had been previous collections on similar topics (e.g., Hinton & Anderson, 1989), the papers in the Rumelhart and McClelland volumes presented the PDP approach as a mature body of work, and aroused significant excitement. Dreyfus & Dreyfus (1990) state that the two-volume collection sold six thousand copies on the first day of its release in 1986, and had sold thirty thousand copies by 1988: these are outstanding sales-figures for a postgraduate-level academic text. McClelland, Rumelhart, & Hinton (1986) offer brief discussion of the historical background to the development of the PDP paradigm. The central notion of computing with neuron-style processing units can be traced back to McCulloch & Pitts's (1943) seminal paper; and to the work of Rosenberg (1962), which attracted damning criticism in Minsky & Papert's 1969 book *Perceptrons*. The idea that learning in neural networks is the product of altering connection strengths was founded by the neuroscience work of Hebb (1949), who proposed that the strength of a connection between two neurons should be increased whenever the two neurons are simultaneously excited: a concept now referred to as *Hebbian learning*. Variations of the basic CTRNN model were studied first by Grossberg (1969, 1988). For reviews of the relationship between PDP-style models and symbolic models, see Smolensky (1988), Clark (1989) and Dreyfus & Dreyfus (1990). Reprints of a set of historically important research papers were collected by Anderson & Rosenfield (1988). The textbook by Hertz, Krogh, & Palmer (1991) gives a thorough grounding with an emphasis on mathematical analysis, while a working set of programs and models was published by McClelland & Rumelhart (1988). For the most comprehensive recent collection, of almost 300 articles spanning a range of topics from detailed biological modelling to practical engineering applications, preceded by excellent background tutorials and "roadmaps", see Arbib (2003).

As interest in PDP models and neural computation increased, a large number of international journals and international conference series were established. It could be argued that not all of these were necessary: a bandwagon had been set rolling, and it seemed in the late 1980's that almost every publisher of academic journals announced at least one new PDP or neural network journal, and new conference series were being announced with a frequency that begged the question of where the international-quality research would come from to fill the pages made available for publication. By the early 1990's there were too many publishing opportunities (conferences and journals) chasing too little good research; consequently, work of questionable future relevance could be published with relative ease. Nevertheless, much good work can be found. One journal widely regarded for its editorial standards and with a strong impact rating is *Neural Computation*, published by MIT Press since 1989; the *Journal of Computational Neuroscience*

and also *Network* have established reputations as quality publications for papers dealing with computational models of biological systems. In addition to the recently established journals, more long-standing journals such as *Biological Cybernetics* and *Proceedings of The Royal Society of London Series B* have a tradition of publishing archival quality work dealing with issues in PDP and neural modelling. One very strong conference series is the annual meeting on *Neural Information Processing Systems* (NIPS), although many good papers have also been published in other conference series such as the *International Joint Conference on Neural Networks* (IJCNN) and *Computation in Neural Systems* (CNS).

At much the same time that interest in artificial neural networks, computational models of nervous systems, and other PDP systems was rapidly growing into a large international research community, orthogonal developments were taking place in research on artificial autonomous agents that exhibit *adaptive behavior*, discussed in the next section.

2.4 Artificial Autonomous Agents and Adaptive Behavior

While developments in PDP offered alternatives to symbolic models of cognition, independent questions were being posed concerning the assumptions, aims, and methods of traditional AI. PDP offered a new style of computing, closer in architecture to natural systems, but in many cases the tasks to which the early PDP systems were applied still often concentrated on high-level cognitive functions, such as language use, the application of 'expert' knowledge, or the learning of complex concepts. Without denying the importance of such topics in cognitive science, a number of researchers had turned their attention to the problem of designing or building artificial systems with a focus on lower-level cognitive functions. Although there was a long-established tradition in neuroscience and psychology of modelling low-level perceptual mechanisms, such as the abilities to detect visual motion or to distinguish between different frequencies of sound, which was also carried over to the new PDP paradigm, the new research paradigm had a key change of emphasis: the intention now was to create *complete* autonomous cognitive agents, rather than models of isolated cognitive functions.

The motivation for this style of work came from a number of directions. One was the long-standing aim to create autonomous mobile robots. Although industrial automation robots (e.g., arms used for welding or paint-spraying in automobile factories) were already in widespread use, they did not require significant intelligence or cognitive capacity to operate successfully. Such robots are rooted to the floor and the number of possible situations they can find themselves in is highly constrained in advance: the position of the part to be welded or sprayed is pre-determined, as is the sequence of movements that the robot must make to successfully complete the task. In stark contrast, if a robot is to be both *mobile* and *autonomous* (i.e., capable of self-determined movements using wheels, legs, or some other propulsion system, over extended areas of space and periods of time, without human intervention), then the number and variety of situations that can arise and problems to be solved rises very rapidly. There is a significant need for such robots, from esoteric applications such as exploring the surfaces of other planets (Brooks & Flynn, 1989); through operation in hazardous environments such as nuclear reactors, battlefields, and collapsed buildings; to more mundane tasks such as security patrols, package delivery, or cleaning operations in populated buildings such as hospitals, warehouses, or homes; along with a variety

of applications in toys and other entertainment systems. In all cases, basic ‘survival skills’ are required of the robot: the ability to wander around, without bumping into things, without getting lost, and without running out of power. It would be desirable to have robots with these and other more advanced skills pre-programmed. The more advanced skills might include the ability to learn the layout of an environment and decide sensible paths through it.

But to program such behaviors into a robot controller requires that the dynamic, unpredictable, and unforgiving nature of natural environments is taken into account. Natural environments are almost always subject to a high degree of inescapable variation: levels of temperature, ambient light, and ground traction may vary considerably; sensors deliver noisy data; motors, gear-trains, and wheels can slip; humans may be wandering around, getting in the way, or forgetting to leave doors open (or closed). Potentially, the effects of all such variation need to be accounted for. Although traditional symbolic AI had long attempted to address the issues of path-planning for mobile robots, the use of sequential programs and logic-based reasoning on real robots (e.g., Moravec, 1983; Nilsson, 1984) really had not been very successful.

Disenchantment with the state of symbolic-AI-based mobile robot research, and with AI in general, led Rodney Brooks (Professor of AI and Director of the AI Lab at MIT, the Massachusetts Institute of Technology), to write two papers (Brooks, 1985, 1986) that proved to be highly influential. Although historic precedents exist (such as the work of Walter and of Ashby, mentioned in Section 1), almost all current research in biologically-inspired autonomous robotics can be traced to these two seminal papers. In these papers, Brooks questioned the traditional focus on high-level cognitive abilities, arguing that models of cognition based on reasoning with symbolic representation of facts had serious shortcomings. Brooks’s central claims were that the basic ‘survival skills’ exhibited by all mobile creatures, even simple animals such as insects, provided a firm foundation on which more advanced cognitive functions could be based and that such systems were best studied by building mobile robots to test theories and techniques. He argued that traditional AI approaches based on separate modules for perception, reasoning, planning, and plan-execution, were inappropriate; and that instead the robots’ control architectures should be decomposed into a number of layers, each with access to sensors and motors, and each responsible for generating different levels of behavior, such as ‘wandering’, ‘collision avoidance’, and ‘map-making’. Brooks referred to this new style of controller design as ‘subsumption architecture’ (because some behaviour-generating layers could *subsume* the outputs of others, when appropriate), and demonstrated the theory in practice in a number of ‘complete-creature’ robots (reviewed in Brooks, 1990), some of which were six-legged insect-like robots while others were wheeled vehicles. His justification for studying simpler insect-like systems is summarised in the following quotation:

“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.”
(Brooks, 1986).

One of the key issues in the work of Brooks and his students is the different nature of problems faced by systems that are both *embodied* and *situated*. In stark contrast to Brooks’s robots, most

traditional symbolic AI systems were 'disembodied', existing only to reason about a given representation of knowledge about some domain: whether the domain was diagnosing blood diseases or steering a vehicle through a cluttered environment, the assumption in symbolic AI was that roughly the same set of techniques could be employed. However, experiences in constructing creature-like robots indicated that the nature of the agent's coupling to its environment, i.e., the details of its embodiment, played a key role in determining appropriate control architectures and strategies. Moreover, many issues in designing the controllers had to be resolved by considering the nature of the environment in which the agent was to operate and the types of behavior that were required of it: that is, the details of its situatedness. The issue of primary significance was the on-going *interaction* between the agent and its environment.

The concentration on interaction dynamics allowed sophisticated complex overall behavior to emerge from controllers that were relatively simple in comparison with equivalent knowledge-based AI approaches. Often, the source of the complexity was the environment with which the agent was interacting: a point first raised in connection with AI by Simon (1969). The natural human tendency to anthropomorphise apparently cognitive agents could also reinforce the false conviction that complex processing was required to achieve purposive behaviors (a point well made in Braitenberg's (1984) delightful book *Vehicles*). In many cases, traditional notions of representation and reasoning played no part in the operation of the agent, even in systems which could map their environments and determine sensible paths from their current location to a goal location (e.g., Mataric, 1990). Horswill (1993a, 1993b), one of Brooks's students, developed a rigorous methodology for analysing target environments to determine constraints that could be used to simplify the design of the controller. This approach, taken to an extreme, gave rise to one of the slogans associated with Brooks's work: *the world is its own best model*. That is, rather than try to develop sophisticated perceptual systems which could provide an internal representation of the external world, suitable for logic-based reasoning, it can be better to enter into a tightly-coupled interaction with the world relying less on complex representation or reasoning, and more on appropriately structured simple (even reflex-based) controllers. For extended discussions of Brooks's arguments and ideas, see Brooks (1991a, 1991b). Arguments such as these led to a shift in attention away from *computation* as a central concept in cognition, towards using the tools and techniques of dynamical systems theory as a more rigorous and theory-neutral approach to cognitive science. For further details, see Smithers (1992), Thelen & Smith (1994), Port & van Gelder (1995), Smithers (1995), Beer (1995a), Harvey (1996b) and Cliff & Noble (1997).

Although Brooks was initially a vocal opponent of simulation studies (Brooks, 1992), other researchers sympathetic to his arguments demonstrated important new principles and techniques using both real and simulated autonomous agents: for an early collection of papers, see Maes (1990). Autonomous agents simulated in software can act as 'virtual creatures' for scientists attempting to further the understanding of perception and cognition in animals: an issue explored briefly below. Moreover, agents with no physical realisation (i.e., that exist only in software) have significant potential for engineering and commercial applications, as was discussed in earlier in this document. Important work was done using software simulations to develop theories and models that addressed the problems of planning sequences of tasks in autonomous agents, and the *action selection problem*: how, given a set of (possibly conflicting) goals, to choose appropriate actions to perform --- "how to do the right thing", in the words of Maes (1989). For examples of such work, see Agre & Chapman (1987, 1990), Chapman (1991), Maes (1989, 1990b), Tyrrell (1993, 1994), Johnson (1991,1995) and Blumberg (1994, 1996).

While subsumption-architecture controllers were essentially PDP systems (being both parallel and distributed), there was no commitment in Brook's work to developing models of real nervous systems, or even to using biologically plausible processing strategies. Despite this, neuroscientists Altman & Kien (1989) noted the similarities between the subsumption-architecture controllers of agents such as the walking robot described by Brooks (1989) and the organisation of neural control in a variety of animals. Subsequent to this, one of Brooks's students explored a number of walking-robot controller-architectures that were more strongly inspired by biology (Ferrell, 1994).

Having spent several years researching the development of robots inspired by simple creatures, Brooks and his students at MIT turned in 1993 to a much more ambitious research program: the construction of humanoid robots. This work was initially centred on the construction and ongoing refinement of an upper-torso humanoid robot called *Cog*. Several graduate students were engaged simultaneously in the development of *Cog*, and their combined efforts led to important innovations: see Adams *et al.* (2000) for an overview. Of particular note is Williamson's (1999) work on the coupling of the dynamics of oscillations in simple "neural" circuits with the dynamics of the robot's body (its arms in particular) for smooth control. Experiences with *Cog* triggered an interest in robots capable of interacting socially with humans, and one of Brooks's students explored this in a purpose-built "emotionally expressive" robot head called *Kismet* (Breazeal, 2000), which has manifest applications in entertainment. For a populist and very readable review of Brook's contributions to behaviour-based robotics, see Brooks (2002).

An individual robot with a behaviour-based control system can be considered as a complex system, because the small-scale interactions of its component behaviour-generating modules give rise to its overall large-scale observable behavioural repertoire. Furthermore, the individual behaviour-generating modules may themselves be delivering behaviours as the large-scale consequences of small-scale interactions if, for example, a module involves the use of an artificial neural network. Moreover, at higher level of analysis, an individual autonomous robot can be considered a small-scale component in a large-scale system if it is one of a number of robots working together as some form of team. Fruitful research in so-called *collective robotics* has been underway for a little over a decade (prior to that, the material costs and high failure rates of the requisite technologies made serious research in this area prohibitively expensive). While studies of teams of humans collaborating and cooperating on the solution of tasks is one potentially valuable source of inspiration, much early work in collective robotics draws inspiration from a more lowly biological inspiration: the collective behaviour of social animals, and in particular the social insect order *hymenoptera* (which includes ants, bees, termites, and wasps). More recently, a surprisingly large amount of research in collective robotics has been directed at an even lower form of life: soccer players, in the various leagues maintained by the international "RoboCup" (www.robocup.org) organisation.

Inspired in part by Brooks's creature-robots, Beer (1990) and Cliff (1990, 1991), independently developed remarkably similar arguments for adopting a more holistic approach to PDP and artificial neural networks; an approach they both referred to as *computational neuroethology*. In doing so, they gave a name to work that had long been practised by other researchers, most notably Arbib (1987). Briefly, neuroethology (e.g., Hoyle, 1984; Simmonds & Young, 1999) is the branch of biology that deals with the neural mechanisms underlying the generation of behaviors

in animals (*ethology* is the science of animal behavior: see e.g., Alcock (1989)). Both Beer and Cliff expressed dissatisfaction at the lack of behavioral relevance of much PDP and artificial neural network research, noting that models of isolated perceptual or cognitive process often relied on questionable input and/or output representations, and that Brooks's arguments for building complete systems made it much less likely that unrealistic intermediate representations would be employed. Consequently, they both argued for the development of 'complete-creature' simulations of situated and embodied model animals: Beer (1990) developed a simulated cockroach, concentrating on issues in motor co-ordination for six-legged walking; while Cliff (1992) developed a simulated hover-fly, concentrating on issues in visual perception for flight control and target-tracking. More recent work in computational neuroethology is reviewed by Cliff (2003). Among other things, Beer's work was significant for placing new emphasis on the notion of *intelligence as adaptive behavior*.

The concept of *adaptive behavior* comes from the ethology literature. While a precise definition is still the subject of some debate, for the purposes of this discussion we'll state that a behavior is adaptive if, when exhibited by an agent, it increases the reproductive fitness of that agent. For an animal, reproductive fitness is a matter of surviving in the animal's ecological niche sufficiently long to reproduce and yield viable offspring. For an artificial autonomous agent, the reproductive fitness could perhaps be measured in terms of number of copies sold into the market, or more likely as a percentage of research and development costs recouped: McFarland (1990) introduced an economic rationalisation of adaptive behavior in artificial agents. It is important to note that adaptive behavior is different from the concept of *adaptation*, e.g. learning in animals or self-tuning control systems. For example, humans who inject heroin do so as a result of adaptation: it is an acquired behavior, rather than an instinctive one; but it is generally not an *adaptive behavior*, because it often has a severe impact on an individual's reproductive fitness by killing them. Conversely, the avoidance of stepping over the edges of steep precipices (the "visual cliff response"), widespread throughout terrestrial mammals, is adaptive but not a result of adaptation in the individual because it is an innate response. However, such innate responses could be described as a result of *evolutionary* adaptation by populations or species over time-scales measured in multiple generations. Furthermore, many adaptive behaviors can be either fully or partially a result of adaptation, either evolutionary adaptation or learning within the lifetime of an individual. These issues are returned to in Section 2.5.

Although early work in artificial autonomous agents and computational neuroethology had little or no emphasis on adaptation, it is widely acknowledged within the field that adaptation or learning is a central topic. Indeed, some of the early work on artificial autonomous agents was developed as a means of testing new approaches and techniques in machine learning. Machine learning is a diverse field, with a well-established journal (*Machine Learning*), and international conference series such as the *International Conference on Machine Learning* (ICML) and the *Annual Conference on Computational Learning Theory* (CoLT). It is beyond the scope of this document to provide a full review of techniques in machine learning: for an extensive recent textbook, see Mitchell (1997). Two early important papers describing simulated autonomous agents that learn to produce adaptive behavior are described briefly below.

As was noted in Section 2.3, the most popular supervised learning technique for artificial neural networks is back-propagation. One of the major problems with supervised learning is its requirement for a training set (and test set) of input and output vectors. There is the practical

problem of the effort required to compile these lists of data; and the methodological objection that biological systems generally do not have such training data available: an animal does not learn to walk by comparing the angles of its joints to some set of target outputs. Parisi, Cecconi, & Nolfi (1990) developed a system where back-propagation was used in a 'closed-loop' fashion by an artificial autonomous agent: because the agent was embodied and situated, it had the capability to generate its own training data. The agent existed within a simple environment, and a back-propagation network was employed to map between sensory inputs and desired outputs. Some of the network's output units were interpreted as action commands, moving the agent's body within the environment; the remainder of the output vector was a *prediction* of the sensory input vector that would result from the execution of the specified action. The actual input vector resulting from the execution of the action could be compared to the predicted values in the preceding output vector, to determine error values used to adjust the network parameters. Such work has strong connections with previous studies in using neural networks for control, where the input received by the neural network is determined at least in part by the past actions executed as a result of outputs given by the network: see e.g. the collection edited by Miller, Sutton & Werbos (1990); indeed, many of the central issues in research on controller networks were first discussed in the cybernetics literature (e.g., Weiner, 1948).

In most cases of interest, the complexity of the animat-environment interaction system to be controlled is such that supervised learning cannot be employed, and so most emphasis in learning for control applications (whether using PDP networks or other technologies) is on the development of reinforcement-based techniques. In many control-learning problems, significant difficulties are posed by the *temporal credit assignment problem*: if payoff (reward or punishment) is received, it may not be a simple consequence of the immediately preceding action; rather, actions performed at some time in the past may have led to the receipt of payoff, and it is the processes that generated those past actions that should be rewarded or punished. Deciding the relative significance of each of a sequence of actions leading to the receipt of nonzero payoff, so that the payoff can be distributed appropriately, is often a highly problematic issue (Sutton, 1984).

Not all researchers working on autonomous agents that learn to exhibit adaptive behavior employ neural networks. In a seminal early paper, Wilson (1985) argued that an understanding of the process of acquiring behaviors appropriate to a given environment, however simple, is fundamental to the understanding of higher-level cognitive functions. To this end, he reported on results from a minimal simulation model where an artificial autonomous agent existed on a toroidal world divided into a rectangular grid of cells. At each cell there could be either a blank space, a 'tree', or an item of 'food'. At any one time, the artificial agent occupied a blank-space cell, and could sense the contents of the surrounding eight cells: idealised minimal sensors returned two binary digits for each cell: one to indicate whether the contents of the cell were opaque, and one to indicate whether the contents 'smelled of food'. Thus, a blank cell was sensed as 00, a tree-cell as 01, and a food-cell as 11: the entire sensory input was sixteen bits of data formed by sampling the agent's immediate surroundings. The agent produced a two-bit output, indicating which of the eight surrounding cells it would try to move into. Moves into tree-cells were blocked, but moves into blank or food cells were allowed. The agent had initially random mappings between sensory input and motor output, and employed a type of learning algorithm known as a *learning classifier system* (LCS) to learn to take short paths to food. LCSs are similar to traditional knowledge-based systems insofar as they employ if-then 'production rules' to determine what output is appropriate for a given input: classifier system learning involves tuning the production

rules to give better behavior, on the basis of ‘reward’ or ‘punishment’ received as a result of past actions, using a combination of reinforcement learning and evolutionary techniques (described in more detail in Section 2.5). Wilson coined the word *Animat* to refer to his ‘artificial animal’, a name which has since entered the research literature as a generic term for artificial creature-like autonomous agents. For further details of variations and extensions to Wilson’s animat and the classifier system employed, see Wilson & Goldberg (1989), Cliff & Bullock (1993), Wilson (1994), Cliff & Ross (1995), and Wilson (1995).

The early 1990’s witnessed a surge of research activity on ‘autonomous agents’ and ‘agent-based’ approaches to a wide range of problems in computer science and artificial intelligence. Unfortunately a bandwagon effect took root, and the word ‘agent’ became seriously overloaded, being applied to a range of systems with widely divergent functions, architectures, and organisational principles and philosophies. There is a sizeable school of ‘agent’ researchers who are sufficiently reliant on logic-based manipulation of symbolic representations that their aims and methods are practically indistinguishable from the old-school symbolic AI of the late 1970’s. For this reason, the word ‘animat’ (Wilson, 1985) is often used to clarify that the style of agent being referred to is a biologically-inspired creature-like artificial autonomous agent, either real or virtual, responsible for co-ordinating its perception and action so as to produce adaptive behaviors for extended periods of time without human intervention. It is the animats that have been developed within the adaptive behavior research paradigm which are of most relevance to the subject of this review. In addition to the collections cited above, the journal *Adaptive Behavior* (published by MIT Press from 1992–1997; now published by Sage Publishing), is a main outlet for publishing work in this area, as are the proceedings of the *International Conference on Simulation of Adaptive Behavior* (SAB), which have been published biennially by MIT Press since 1990, under the series title *From Animals to Animats*. Clark’s (1996) book is a highly readable review of key developments over the preceding ten years in research dealing with animats, adaptive behavior, and neuroscience. For more populist accounts, see Wallich (1991) and Kelly (1994).

2.5 Evolutionary Computation

For most interesting behaviors and environments, the task of designing control systems for animats (whether real robots or software agents) is non-trivial. As was noted above, the development of learning algorithms for PDP networks was necessary because of the generally extreme difficulty of manually determining appropriate settings for the parameters of the system. In essence, this is a problem in *function optimisation*: trying to find a combination of parameter values (i.e., a point in parameter space) that minimises or maximises some function. In the case of a supervised-learning PDP network, the function to be minimised is given by the error measure. Although Darwinian evolution in biology is not actually an optimisation process, it has inspired the development of several techniques for optimisation, which will be briefly reviewed here. Evolutionary computation techniques are of particular interest for optimising the design of CTRNNs, as no traditional neural-network learning algorithms have yet been developed for training CTRNNs with the same effectiveness that gradient-descent methods such as back-propagation exhibit on feed-forward networks.

All evolutionary computation approaches to searching a given parameter space share a common core of concepts, introduced below. Evolutionary approaches have been applied to a wide variety

of optimisation problems, for a review, see the textbooks by Goldberg (1989) and Mitchell (1998). The focus here is on the use of evolutionary techniques in designing artificial neural networks for autonomous agents, because of its obvious relevance to the study of cognitive systems.

One of the most popular and deeply-studied forms of evolutionary computation is the so-called *genetic algorithm*. The basic form of genetic algorithm (GA) is relatively simple to explain:

Consider the task of finding a good or optimal set of n parameters for a neural network. That is, finding good or optimal points in the n -dimensional parameter space for the network. To use a GA, first develop a method for *encoding* the parameters as a string of characters. One simple and popular method is to encode the numeric values as strings of binary digits (bits). So for an n -dimensional parameter space, with each parameter encoded by b bits, the string is $b.n$ bits long. Start by randomly generating some number of such strings. Borrowing terminology from biology, each of the strings is referred to as the *genotype* (or *genes*) for an *individual* in the *population*. Next, enter an iterative loop of *evaluation* and *breeding*. In the evaluation phase of the loop, individuals are tested to measure their performance (e.g., on a test-set of data, in much the same way as in PDP learning algorithms), and assigned a single numeric *fitness* value, representing how good they are. When all individuals have had their fitness evaluated, *breed* a new population of individuals by *selecting* fitter members of the evaluated population as 'parents', and mixing their genes in ways inspired by the recombinant genetics of sexual reproduction, possibly also introducing random alterations or *mutations* to the genes. Finally, dispose of the old population, replacing it with the newly-bred one. This process of taking a population, evaluating each individual, selecting parents, and breeding a new population to replace the old one, is referred to as one *generation*. If the parameters governing the operation of the GA are set correctly, the best or average fitness in the population will rise over a number of generations: the GA is really not much more than a directed stochastic generate-and-test search process.

The operation of a simple GA such as this can be viewed in a similar manner to PDP learning trying to find minima on an error surface (discussed in Section 2.3). Each possible genotype represents a point in *genotype space*, and the function used for fitness evaluation defines a fitness surface or *landscape* over this space. The aim of the GA is to find points of high (or maximal) fitness in the fitness landscape. In contrast to PDP learning, where a single network traces a trajectory of points through parameter space as the learning algorithm alters settings of weights and thresholds, adaptation in a GA is a consequence of a *parallel* search of the parameter-space that is encoded in the genotype-space: each individual in the population occupies a point on the fitness landscape. Individuals at low levels in the landscape are unlikely to be selected for breeding; whereas high-fitness individuals, more likely to be selected for breeding, will combine their genes with other fit individuals. Consequently, in the next generation, there should be proportionately more individuals at high points on the fitness landscape and fewer at the low points. Over successive generations, the population is said to *converge* on peaks in the fitness landscape. In many GA applications, the GA is set to run for a pre-determined number of generations, but in some cases the system monitors convergence and terminates the evolutionary process once the population reaches some level of convergence: the 'evolutionary search' is considered to have ended when the population has converged beyond some level. Convergence can be measured by, for example, calculating summary statistics such as the mean Euclidian (straight-line) distance between individuals in genotype space. As with the problem of local minima on error surfaces, premature convergence to local peaks on fitness landscapes can cause significant problems. In

essence, the selection mechanisms help converge the population on fitness peaks, and the addition of random mutations helps 'spread' the population 'downhill', away from the peaks. In picturesque terms, mutation can be thought of as moving individuals onto lower ground, in the hope that new paths to even higher nearby peaks can be found from there. Striking an acceptable balance between selection and mutation is very important. If the mutation rate (often expressed as the probability that an individual 'child' genome will have the character at one randomly determined locus altered to a randomly determined different allele) is too low, premature convergence to local sub-optimal peaks is more likely; but with too high a mutation rate, the GA degenerates to random search. Some researchers (e.g., Fogarty, 1989) have developed techniques for dynamically varying the mutation rate as a consequence of changes in convergence.

The balance between mutation and selection is only one of many issues in constructing a working GA. Another important issue is the selection mechanism: much research has gone into exploring the effects of different styles of selection upon the evolutionary process (e.g., Whitley, 1989). One of the most popular styles is *roulette selection*, where the probability that an individual is selected to be a parent is given by calculating what proportion of the total population fitness is given by that individual's fitness: metaphorically, a roulette-wheel is used, where the size of the slice of the wheel given to each individual is proportional to its fitness. The relative differences between the probability of selection between individuals of differing fitness is referred to as the *selection pressure*: in roulette selection, the selection pressure can be varied by, for instance, making the size of each individual's slice on the wheel proportional to a power or exponential function of the individual's fitness (giving a greater bias to fitter individuals), or using a logarithmic function of fitness to reduce the advantage given to fitter individuals. However, a number of authors have argued that there are significant problems with roulette selection: in the early stages of the GA, if a fit mutant is generated, its fitness may be orders of magnitude better than the remaining (randomly generated) individuals, and so the population rapidly converges on points in genotype space near the mutant. An alternative process of *tournament selection* has been advocated, where individuals compete for the right to breed in 'tournaments': in a simple form of tournament selection, three individuals may be selected at random and the fittest two of the three are allowed to breed, with the least fit of the three being replaced by the child that results from breeding. Tournament selection gives less bias to fit mutants occurring early in the evolutionary process, and preserves diversity. Such a tournament selection process gives *rank-based* selection, where differences in the reproductive activity of a group of individuals depend on the numerical *ranking* of their fitnesses, rather than the *magnitude* of the differences between their fitnesses. For further discussion of the benefit of rank-based selection, see Whitley (1989).

A variety of other methods of selection have been developed, such as *truncation selection*, where a fixed number of high-fitness individuals (e.g., the top 10%) are all given an equal probability of selection, and the remainder of the population are excluded from the breeding process. Finally, a common technique in many GA applications is *elitism*, where a copy of the fittest member of the evaluated population (the *elite* individual) is always added into the population of the next generation without mutation, thereby ensuring that the best individual encountered so far is always preserved. Without elitism, there is an increased danger that *genetic drift* sets in, with the population never returning to previously-visited high points on the fitness landscape.

Other refinements to the GA process have been proposed for avoiding premature convergence and promoting diversity. One is the use of *spatially distributed* GAs, where the population is

arranged on a spatial grid (typically in two or three dimensions, but higher dimensions are perfectly possible). In breeding, individuals can only reproduce with 'nearby' neighbours, and the offspring are assigned grid-cells in the next generation that are close to the places occupied by their parents in the current generation. This contrasts with the conventional GA, which is *panmictic*: there is no structure to the population and any individual can, in principle, breed with any other. The intention in spatially distributed GAs is that there is convergence on a local scale, but no necessity for individuals at distant points in the grid to be genetically similar. This can be considered as a loose approximation to natural evolving populations, where geographic spread can facilitate the formation of separate *species*: members of the same species share much genetic material (are highly converged), but the presence of multiple species ensures that the overall gene-pool is diverse. The different groups can be thought of as separate local clusters or 'clouds' in genotype space. A second method is the use of a *steady-state* GA (SSGA). An SSGA differs from a conventional (or *generational* GA) in that there is no division of the evolutionary process into discrete generations: rather, at each iteration through the loop, two individuals are selected for reproduction, offspring are bred, and individuals are then selected for replacement by those offspring. Thus in a SSGA successive generations meld into one another, with smooth changes in the genetic constitution of the population, rather than the possibly discontinuous jumps of a generational GA. SSGAs can be either panmictic or spatially distributed.

The algorithmic simplicity of most GAs means that the step of evaluating the fitness scores of the individual genotypes can represent the most significant computational cost in many GA applications. In cases where the genetically encoded parameters are used to define a neural network for replicating some input-output mapping, the length of time taken to evaluate an individual is proportional to the number of input-output pairs in the training set. But when the aim is to evolve networks that produce adaptive behaviors in animats, it is generally necessary to make multiple evaluations to generate statistically reliable fitness data. Multiple evaluations are clearly necessary if the agent or environment are stochastic (e.g., if there is noise in the system), but even in purely deterministic systems it may be necessary to make multiple trials to explore the effects of changes in initial conditions, or of differing environments. Aizawa & Wah (1994) developed sophisticated statistical techniques for varying the number of trials per individual both within and between generations. Furthermore, for many desired behaviors that are easily expressed in English, formulating an appropriate mathematical evaluation function can be a surprisingly difficult task. The evolutionary process may well find peaks on the fitness landscape resulting from a given evaluation function, but even if these are global maxima they might not correspond to the behaviors desired by the designer of the system: an issue discussed further by Zaera, Cliff, & Bruten (1996), and Matarić & Cliff (1996).

The reproduction process involves the application of *genetic operators*. In addition to mutation, where characters at randomly chosen positions on the gene-string are replaced by other randomly chosen characters, sexual reproduction is modelled using *recombination*, also referred to as *crossover*. Having used the selection mechanism to identify two 'parent' genotypes, one of the two parents is randomly chosen and characters from its genotype are copied into the child's genotype. At a randomly chosen position, the copying process will 'cross-over' to the corresponding position on the genotype of the other parent, copying the remainder of that parent's genotype into the child's genome. Thus, the child inherits some of its genetic material from the first parent and the rest from the other parent. This, so-called 'one-point crossover' is the simplest form: more sophisticated forms may involve multiple crossings back and forth between the two parents; and

some researchers have experimented with systems involving more than two parents (e.g. Eiben, van Kemenade, & Kok, 1995).

In cases where the genotype represents a small number of numeric parameters, formulating an encoding scheme is not especially problematic. While the parameters can be encoded as binary strings, making mutation and crossover straightforward, it is increasingly common to eliminate the computational costs involved in converting between the binary genes and the decimal values used in evaluating the genotypes. Instead, the genotype is a string of decimal parameter values: mutation involves adding randomly generated values to one or more of the numbers on a genotype, and crossover remains much the same. Nevertheless, the ordering of the parameters on the genetic encoding can have a significant effect on the evolutionary process. If two parameters are represented by numbers or character strings that are positioned at loci close together on the genotype, there is less chance that they will be separated in the crossover process than if they are positioned far apart on the genotype. As the distance between the two loci is reduced, it is more likely that a child will inherit the values of both parameters from one of its two parents, rather than one value from the first parent and the other value from the second parent. This is important if there is some degree of *linkage* between the two parameter values: if the fitness contribution of one parameter is determined at least in part by the value of the other parameter, then once a good combination of values is present on a genotype it is desirable that these two genes are not separated in reproduction.

In most cases of interest, the fitness contribution of any one genetically specified parameter will be dependent on the values of several other parameters on the genotype. The network of cross-dependencies can be complex, and can have a radical effect on the nature of the fitness landscape. This phenomena is referred to as *epistasis*: if an evolving system is highly epistatic, there is a high degree of interdependency; while a low-epistasis system allows each of the parameter values to be altered largely independent of the others. The degree of epistasis is a consequence of the interaction between the encoding used and the fitness evaluation function, and determines the nature of the fitness landscape. A highly epistatic (or *rugged*) fitness landscape may have multiple false peaks, deep chasms, and cliffs, all of which offer potential traps for the trajectories of individuals evolving over that landscape. In general, less rugged (or more *smooth*) fitness landscapes are desirable. Unfortunately, it is frequently difficult or impossible to predict in advance the degree to which epistasis will be a problem. In applications involving complex encodings or where multiple expensive evaluations are needed to determine the fitness of each individual, it is common to try to characterise the ruggedness of the fitness landscape by empirical means. One method of doing this is to estimate or measure a metric of *fitness correlation*: an indication of the relationship between distance in genotype-space (e.g., number of single mutations) from a given genotype, and the resultant change in fitness. For smooth landscapes, moving small distances in genotype-space will result in small positive or negative effects on fitness, while in rugged landscapes the fitness changes will be less correlated, or random. Furthermore, many fitness landscapes have *plateaus*: areas of connected points in genotype-space with the same fitness value, where genetic drift can easily occur. The presence of plateaus or “neutral networks” in a genotype space may have beneficial effects on an evolutionary process operating in that space (see e.g. Huynen *et al.*, 1996; Barnett, 1997). Kauffmann (1993) developed a set of theoretical models known as *NK* landscapes, modelling arbitrary genotypes of length N with the fitness contribution of each gene being dependent on K other genes in the genotype. By varying K as a proportion of N from zero to one, fitness landscapes can be varied from perfectly

smooth to highly rugged and uncorrelated surfaces. Many researchers have used *NK* models as test-beds for developing new GA techniques, although fitness landscapes in real applications are likely to show combinations of smooth zones, rugged zones, and plateaus; rather than the uniform-epistasis landscapes generated by *NK* models.

For situations where a system to be optimised by evolution has a large number of parameters, straightforward linear encodings of the parameters on the genotype frequently become unwieldy. For example, a neural network with 20 numeric parameters, each encoded as 16-bit values, requires a genotype of $20 \times 16 = 320$ bits, thus the total number of points in the space of possible genotypes is 2^{320} , or approximately 2×10^{96} . This is significantly more points than the number of atoms in the known universe. Although not all of these points need be evaluated in the evolutionary process, it is clear that the lengths of genotypes, and the number of points in the genotype space, can rise very rapidly. Furthermore, attempting to evolve a good combination for a fixed number of parameters presupposes that the correct number of parameters is known at the outset. In evolving a neural network, this requires that the number of units and connections is specified in advance. If the number of parameters is not roughly correct, severe problems can occur: the network may be too small to perform the required task (e.g., by not having sufficient numbers of internal units); or the network may be too big, in which case the desired parameter-space is a sub-space of the evolutionary parameter space, and resources may be wasted just in getting the population into the right sub-space. For these reasons, among others, a number of researchers have developed GA principles and methods that allow *variable-length* genotypes.

If the length of the genotype can vary, and the length corresponds in some way to the number of parameters in (or the complexity of) a design, then in principle it is possible to start with a population of short genotypes encoding for simple designs. Provided that there is sufficient variation in the behaviors of the initial population for evolution to operate, longer genotypes or more complex designs will occasionally occur and will be retained in the population if they have sufficiently increased fitness. Thus, it should be possible to continue evolution, with the complexity of the designs increasing until the desired behaviors are exhibited by the evolving individuals. This is clearly a much more powerful method than evolving fixed lists of parameters. But with this added power comes additional problems. First, visualising the evolutionary process is now more difficult: as the length of the genotypes grows, the dimensionality of the search space increases. Individuals in the same population with different-length genotypes are points in parameter-spaces of different dimensions, and the search is now not just for an appropriate setting of parameters, but also for an appropriate dimension of parameter-space. Furthermore, deciding on how best to perform crossover on two parent genotypes that differ in length is problematic. These and other issues have been studied in depth by Harvey (1990, 1992a, 1992b, 1993), who has developed a set of principles referred to as *species adaptation genetic algorithms*, or SAGA, for dealing with variable-length genotypes in a variety of applications. In SAGA systems, the population is always highly converged, and can be viewed as a *species*: much more emphasis is placed on mutation (which, combined with selection, moves the population through the space of possible genotypes) than recombination (i.e., crossover, which is seen as the driving force in converging to solutions in fixed-length GA systems).

Genotype lengths can be varied by a number of means. The crossover operator can be adjusted so that, rather than jumping to the corresponding character in the other parent, the jump is to a point in the genotype either somewhere before or somewhere after the current copying location;

resulting in either a longer or a shorter child genotype, even if the two parents have identical-length genotypes. This was the method used in early tests of SAGA for evolving neural network controllers for animats (Cliff, Harvey & Husbands, 1993; Harvey, Husbands, & Cliff, 1994). Another possibility is to introduce new genetic operators such as 'deletion', which deletes a randomly selected sequence from the genotype, closing up the gap to give a shorter string, or 'duplication' which takes a sequence on the genotype with randomly chosen start and end points, and attaches a duplication of that sequence to the genotype at some other point (e.g., inserting the duplication immediately after the original sequence, or at one end of the genotype): see Miller & Cliff (1996) for discussion of the beneficial use of a duplication operator.

When genotype lengths vary, more sophisticated genetic encoding schemes are needed. The encodings must cope with variations in length while still coding for useful designs, without allowing the genetic operators used to cause severe epistatic interactions. A number of workers have addressed this problem, which continues to be an active research issue. A notable body of work has been originated by Gruau (Gruau & Whitley, 1993; Gruau, 1994), who identified a set of desirable properties for variable-length encodings of neural networks, and then developed an encoding scheme that he demonstrated in use for evolving a variety of networks, including animat controllers. Gruau uses techniques inspired by *Genetic Programming*, popularized by Koza (1992, 1994), where variable-length hierarchical tree-structured genotypes are evolved: Koza's work has concentrated on the use of these techniques to evolve expressions and programs (often using subsets of the LISP language). Gruau's work involves evolving hierarchical developmental 'programs' specifying a sequence of operations to be performed on an initial 'cell' that undergoes a sequence of 'divisions' to form the final network. This process is similar to the manner in which a single fertilised egg goes through successive divisions to produce a foetus that subsequently matures. By the same analogy, much other research on genetic encodings (for fixed or variable-length genotypes) concentrates on the genome specifying a set of parameters to a 'development' or *morphogenesis* process that determines the final agent architecture. In several cases, researchers have noted that the genes should specify not only the control architecture, but also the physical design (or 'morphology') of the agent (Brooks, 1992; Cliff et al., 1993; Sims, 1994, 1995).

When evolving neural networks, the genotype can specify the thresholds and weights of the network, which then remain fixed throughout the 'lifetime' of that network. However, a number of researchers have studied systems where the genotype specifies the initial values of the network parameters, which are subsequently adjusted by a learning algorithm during the lifetime of the network. In a landmark paper, Hinton & Nowlan (1987) provided the first demonstration of this, where the ability to learn from experience allows a population to evolve on an otherwise impossibly epistatic landscape. The significance of this paper is that, while it generally accepted that experiences acquired through the lifetime of an individual cannot be passed on in its genes (so-called *Lamarckian inheritance*), Hinton & Nowlan's abstract simulation study demonstrated that the Baldwin effect operates to the advantage of the population: in an evolving population of agents each with the ability to learn during their lifetime, the presence of learning can *guide* the population to high-fitness areas of genotype space. For further theoretical studies of the Baldwin effect see Mayley (1996a, 1996b) and Harvey (1996a); for early work on evolving neural networks that learn, see Miller, Todd, & Hegde (1989) and Belew, McInerny, & Schraudolph (1992); more recent studies include Floreano & Mondada (1996) and Nolfi & Parisi (1997).

Finally, all of the preceding discussion has assumed that the fitness landscape is essentially fixed for the duration of the evolutionary process. In real biological systems, this is rarely the case. The fitness of members of a particular evolving population may be dependent on the behaviors of members of other populations, which are themselves evolving. Thus, the fitness landscape of the first population can be altered by evolutionary changes in a second population. The clearest examples in nature of such co-evolutionary interactions are between species of parasite and host, or between predator and prey: predators may evolve more powerful bodies or behaviors, increasing their reproductive fitness by catching more prey; this alters the fitness landscapes of prey, giving a pressure to evolve better camouflage or faster and more agile escape behaviors -- when the prey have evolved to counteract the evolutionary developments of the predators, the evolutionary pressure is transferred back to the predator population, to evolve even better bodies or behaviors. These kind of co-evolutionary arms-race interactions are often referred to as examples of the *Red Queen effect* (Van Valen, 1973), after the character in Lewis Carroll's *Alice's Adventures Through the Looking Glass* who had to keep running forward in order to stay at the same place, because the landscape was moving underneath her. Evolutionary arms-races in biology have inspired workers in artificial evolution to use co-evolution as a method of encouraging continuing and open-ended evolutionary innovation. One of the first examples of applying artificial co-evolution to problems of genuine industrial interest was by Hillis (1991), who co-evolved a population of individual sorting algorithms (whose fitness was determined by their ability to sort lists of integers) against a population consisting of lists of integers (whose fitness was determined by their ability to cause problems for the sorting algorithms). For an extended discussion of the use of co-evolution in evolving animats, see Miller & Cliff (1994), and for discussion of the problems involved in monitoring progress in such systems, see Cliff & Miller (1995). Theoretical studies of co-evolution include Kaufman's (1993) extension of the *NK* models to cover a number C of interacting species (yielding *NKC* models), and the work by a number of authors studying co-evolutionary interactions in systems evolved to play the *Prisoner's Dilemma Game*: see e.g., Axelrod (1984), Lindgren (1992), Stanley, Ashlock & Tesfatsion (1993), and Batali & Kitcher (1994). Sims (1994, 1995) produced visually impressive work using variable-length genotypes for co-evolving the control-networks and body-designs of virtual animats.

One novel issue in co-evolution, explored by Todd & Miller (1993), is the use of *sexual selection*, where the reproductive fitness of individuals is dependent on their 'desirability' to members of the opposite sex. Sexual selection as a force in evolution was long overlooked but has received increasing attention in recent years: see e.g. Cronin (1993) and Ridley (1993). Classic examples of the effects of sexual selection are physical and behavioral 'ornaments' such as the brightly coloured tails of peacocks or elaborate patterns of 'ritual' courting behavior. Todd & Miller (1993) demonstrated sexual selection in action using a simple simulation, where breeding-pairs were chosen on the basis of the mate-preferences of individuals. They showed that the effects of sexual selection can drive populations of individuals away from local fitness peaks, possibly resulting in the discovery of higher peaks. Sexual selection can also be a cause of *speciation*, where a single-species population evolves into two or more distinct species that do not interbreed or are not capable of doing so, as a result of mate-choice. Miller (2000) argued persuasively that runaway sexual selection was a primary factor driving the evolution of those cognitive and intellectual abilities in humans that distinguish us from other primates.

As with artificial neural networks, the fundamental ideas in artificial evolution have been present in the literature since the 1960's, but have only attracted significant interest in the past two

decades or so. John Holland of Michigan University is widely acknowledged as the founder of genetic algorithms: see Holland (1975, 1992), and Holland et al. (1986). Holland (1975) is credited with the development of the *Schema Theorem*, a mathematical argument to the effect that useful ‘building blocks’ on the individuals’ genotypes will be combined by the crossover operator, giving efficient search despite the essentially random nature of the GA.

There are a number of respectable conference and workshop series addressing issues in evolutionary computation. Two of the earliest and longest-established conference series are the biennial International Conference on Genetic Algorithms (ICGA: established 1985) and the biennial conference on Parallel Problem Solving from Nature (PPSN: established 1990). A number of smaller and more specialised conference series ran through the 1990’s, but recently a series of mergers have led to the two main conferences in the field being the annual Congress on Evolutionary Computation (CEC), and the annual Genetic and Evolutionary Computation Conference (GECCO), within which the ICGA series is now subsumed. The premier journals dedicated to artificial evolution are *Evolutionary Computation*, published since 1992 by MIT Press; and the *IEEE Transactions on Evolutionary Computing*, established 1996. In addition to the regular conferences proceedings, a number of edited collections resulting from workshops and other meetings have been published: the Foundations of Genetic Algorithms (FOGA) series is notable for the high technical standard of its contributions. Papers involving the application of artificial evolution are also frequently published in related journals such as *Artificial Life*, *Adaptive Behavior*, *Machine Learning*, and *Complex Systems*; and related conferences such as the *ICML*, *CoLT*, *SAB*, *ECAL* and *A-Life* series introduced in previous sections. Readers interested in the underlying biological issues are referred to Sigmund (1993) and Dawkins (1982, 1986, 1989) for popular accounts; to Griffiths et al. (1993) for detailed discussion of genetics; and Gilbert (1992) or Edelman (1988) for detailed discussion of morphogenesis and other issues in developmental biology. For an extended review of the application of evolutionary techniques to robotic autonomous agents, and a discussion of some of the significant problems and challenges, see Matarić & Cliff (1996).

2.6 Artificial Biochemistries

Artificial biochemistries have been a constant theme of research within the A-Life literature since the first meetings in the late 1980s. Most of the published research has concentrated on modelling the origin of life, i.e. the emergence of early self-replicating molecules (such as RNA) from a ‘prebiotic soup’. In particular, attention has been focused on studying abstract computational models of *autocatalytic sets*.

In brief, an autocatalytic set is a group of reactive molecules that each act as *catalysts*. A catalyst is a chemical that can influence (e.g., enable or accelerate) a chemical reaction without being permanently changed by its involvement in that reaction. The catalysts in an autocatalytic set are involved in the production of a cyclic chain-reaction leading to the generation of more of the same catalysts. For example, if a catalytic molecule *A* is added to a ‘soup’ of chemicals, and it catalyses a reaction that produces a new molecule *B*, and *B* catalyses a reaction yielding another molecule *C*, and *C* is a catalyst for a reaction that produces another molecule *D*, then a chain of reactions yielding *D* can be initiated by adding a quantity of *A* to an appropriate ‘soup’. However, if *D* itself catalyses a reaction that yields some quantity of molecule *A*, then the chain becomes a loop or

cycle. If the conditions are right, such a cycle might give rise to a self-sustaining chemical reaction: the addition (or random synthesis) of a small quantity of any one of the catalysts in the autocatalytic set establishes the cycle of reactions which then continues as long as resources allow. In real autocatalytic sets, there may be multiple loops with cross-connections between them. Such cyclic reactions have been argued to be the basis for the emergence of complex biochemical reactions that sustain life (Kauffman, 1993), and also to be a useful metaphor for economic systems (Waldrop, 1993, pp.120–125). The concept of autocatalytic sets is attributed to Kauffman. Almost all the early relevant papers on this and related topics in the A-Life literature are covered by the following references: Tamayo & Hartman (1989); Zeleny, Klir, & Hufford (1989); Lugowski (1989); Rasmussen (1989); Rasmussen et al. (1989); Kauffman (1991, 1993); Bagley & Farmer (1992); Bagely, Farmer, & Fontana (1992); Fontana (1992); Schuster (1992); Morowitz (1993); Kitano (1994); & Banzhaf (1994).

The degree of research activity in such artificial biochemistries is not yet sufficiently large to warrant establishing an international learned journal dedicated to the topic, but journals such as *Physica D*, *Complex Systems*, and *The Journal of Theoretical Biology* are possible publishing outlets in addition to the *Artificial Life* journal and the A-Life conferences.

In addition to abstract theoretical models, computer simulations are increasingly being used to model real biochemical systems for research in molecular biology and the pharmaceuticals industry. A number of companies have been established in the last two decade, dedicated to the use of artificial evolution techniques for drug design. However, rather than use genetic algorithms or other *in silico* computational approaches, real molecules are generated at random, evaluated, and selected *in vitro*.⁴ For further details, see e.g. Joyce (1992) and Schuster (1993, 1995).

Farmer (1991) draws explicit mathematical links between connectionist networks and autocatalytic sets, offering the possibility of a unified treatment of the two topics. Subsequent to the publication of Farmer's paper, new developments in neuroscience research have opened up the possibility of creating systems that incorporate significant interactions between artificial neural networks and artificial biochemistry. While much of that literature deals with issues in the origins of complex organic compounds and self-sustaining cyclic chain reactions, there is a small but growing body of work that addresses the interaction of biochemical activity with the operation of a neural network.

Although the effects of psychoactive hormones and drugs have long been studied in neurobiology, the assumption was long held that all signalling from one neuron to another was via *direct* connections at the synapses (the point where an output terminal or *axon* of a signalling neuron connects with the membrane of receiving neuron). Most such connections are chemical: electrical activity on the membrane of the signalling neuron initiates the release of small packets of *neurotransmitter* chemical from the axon onto the membrane of the receiving neuron, a process which alters the level of electrical activity on the receiving neuron's membrane. Much less common are electrical connections, where electrical activity on the signalling neuron's membrane directly affects that of the receiving neuron, without intervening neurotransmitters being released.

⁴ The prospect of genetically modifying simple life-forms such as bacteria or plant-life to create new drugs means that *in vivo* techniques are now frequent complements or rivals to *in vitro* procedures. When using *in silico* computational simulations there is a clear need for verification and validation by cross-checking the output of computational models against data from similar experiments with real data: *in vivo veritas*, one might say.

However, recent developments in neuroscience have identified the presence of *gaseous* neurotransmitters: nitric oxide in particular. This discovery indicates that neurons may be capable of signalling in a *diffuse* manner, by release of gases to nearby neurons. For further details, see e.g. Elphick et al. (1995, 1996). One of the first artificial neural network models that interacted with an associated computational biochemistry was described by Grand, Cliff & Malhotra (1996) and Grand & Cliff (1998); see also Grand (2000), but this was closer in spirit to the way in which hormones affect neural activity. Recent research in artificial neural networks that incorporate models of gaseous neuro-modulation in addition to direct connections between neurons includes the work by Phillippides *et al.* (2000) and Husbands *et al.* (2001).

3 Applications of Artificial Autonomous Cognitive Agents

“Real-world” applications of artificial autonomous agents with animal-like capabilities have attracted increasing attention over the last two decades, and a primary influence is the desire to build truly autonomous mobile robots. Nevertheless, autonomous agents that exist purely as software entities (with no physical realisation) also have a number of important commercial and scientific applications. This section briefly reviews the state of the art in Section 3.1 and then discusses what role biological metaphors might play in future developments in Section 3.2.

3.1 State of the Art

The word “agent” has in the past decade come to mean many things to many people, and there is a tension between some of those meanings. Some perfectly well-meaning and respected researchers consider the use of an agent metaphor to be a natural next step for use in computer programming in general. Such a transition to “agent-oriented programming” would be similar to the spread and adoption of object-oriented programming techniques (as embodied in programming languages such as C++ or Java) that occurred during the 1990’s. Under such a view, any procedure, method, or function performed by a computer program can in principle be referred to as an agent if the programmer chooses to do so, and indeed non-agent legacy software systems can also be “wrapped” in an agent interface, thereby hiding their non-agent origins. While this approach has some appeals viewed within the context of the history of programming language design, it also has the effect of widening the definition of the word “agent” almost to the point of vacuity. It is beyond the scope of this document to provide a review of the entire field of research in artificial autonomous agents. The interested reader is referred to the *International Journal of Autonomous Agents and Multi-Agent Systems* (published by Kluwer since 1997) and the associated international conference series that has been running since 1998 (see www.aamas-conference.org), which are noted both for their high editorial standards and also for their very broad interpretation of what counts as an “agent”. Here we focus our attention on those strands of autonomous agent research that have a strong complex adaptive systems focus, and where ideas or metaphors from biological systems have influenced the development of new techniques or technologies; that is, we focus here on BICAS treatments of agents.

One obvious distinction within BICAS agent research is between real physical artificial autonomous agents (i.e., robots), and agents with no physical embodiment, i.e. software agents that exist purely in virtual environments. We discuss the state of the art in these two areas separately in Sections 3.1.1 and 3.1.2 respectively.

3.1.1 Autonomous Robots

As was discussed above, the mid-1980’s work of Brooks and his students at MIT argued forcefully for a behaviour-based or “bottom-up” approach to cognition, which presumes that displays of intelligence are the product of complex interactions between the behavioural repertoire of an agent and its environment, where that agent’s behavioural repertoire is itself the product of the non-linear system formed from multiple interacting behaviour-generating modules within the

agent. Although the practicality of this approach was demonstrated in a series of “insect-like” autonomous mobile robots developed by Brooks and his students at MIT during the late 1980’s, subsequent attempts to apply behaviour-based engineering techniques in creating more cognitively complex behaviours proved difficult. In the early 1990s, Brooks’ MIT group shifted its attention to the construction of the humanoid robot *Cog* (see www.ai.mit.edu/projects/cog), again using behaviour-based control techniques. While the mechanical engineering and low-level sensory-motor coordination and control aspects of *Cog* were widely agreed to be novel innovations (e.g., Williamson, 1999), successful demonstration of higher-level cognitive functions was somewhat more elusive. One insight of the *Cog* project was the significance of social interactions between humans and humanoid robots, an issue subsequently studied in more depth using the *Kismet* “socially expressive” robot head developed by one of Brooks’ PhD students (Breazeal, 2002). The *lack* of social ability in *Cog* led to the initiation of research on behaviour-based humanoid robots as models for the diagnosis and quantification of social development disorders such as autism (Scassellati, 2000). Furthermore, although *Kismet* was developed for principled scientific study of human-robot social interaction, it clearly has potential applications in creating animatronic automata for entertainment purposes.

Behavior-based physical robots with animal-like capabilities for autonomous action and survival have many obvious applications in areas such as hazardous environments (including battlefields), industrial automation, domestic cleaning and security, and in the entertainment and leisure industries. Behavior-based robots for all these applications are available commercially from Brooks’s company *iRobot* (www.irobot.com).

Robot models may also act as physical simulations of real creatures, used to test scientific hypotheses concerning the organisation of a real animal’s sensory-motor control system (i.e. its nervous system), as in the work of Franceschini *et al.* (1992), Srinivisan *et al.* (1997, 1998), or Webb (2000, 2002, 2003). In addition to using humanoid socially expressive robots for the exploration of social development disorders such as autism, autonomous biologically-inspired non-humanoid robots for the remedial therapeutic *treatment* of autistic children are currently under development (Dautenhahn *et al.*, 2002). Furthermore, biologically-inspired control systems have recently been developed for “intelligent” prosthetic limbs and other assistive robotic technologies; most notably at the MIT Leg Lab (www.ai.mit.edu/projects/leglab).

The journals *Autonomous Robots* (published by Kluwer since 1994) and *Robotics and Autonomous Systems* (published by Elsevier/North-Holland since 1985) are good sources for current work in all approaches to the development of autonomous robots. The recent collection edited by Webb & Consi (2001) focuses on the role of biologically-inspired robots used as scientific tools for biologists, while Arkin’s (1998) text provides a thorough overview of biological inspiration in robotics for engineering purposes.

Many academic biologically-inspired robot researchers in the UK and in mainland Europe perform their research using the miniature *Khepera* robot platform, manufactured in Switzerland by K-Team (www.k-team.com). This wheeled robot has two independent drive motors, on-board batteries, an array of eight paired infra-red reflectance emitter-detector sensors, and a powerful single-chip processor with associated memory to which control programs can be downloaded for autonomous running, all in a compact package that is approximately the size of a small tin of tuna-fish. Although a real physical entity, there is a view among autonomous robotics researchers

(especially in the USA) that *Khepera* is best considered as a *physical* simulation of a real robot, as its miniaturized format severely limits the range of environments in which it can operate.

3.1.2 Software Agents

Research exploring BICAS approaches to the creation of autonomous agents that exist purely in software has the advantage of requiring no expenditure on the construction and maintenance of real physical robots. If the interactions of a real robot and its environment can be accurately simulated in software, this is an obvious advantage. However, there is a constant (and often seriously under-appreciated) danger of working with simulations: if the simulation has been poorly verified (or not verified at all) then there is the possibility that the simulation does not faithfully model the real world system that it is intended to represent, and hence that the results from the simulation study will not be replicable in the real world. This danger is heightened when the agent is adaptive, as the adaptation mechanisms (e.g., learning in a neural network, or the use of a genetic algorithm to tune the design of the agent) may exploit flaws in the simulation, and this may go undetected. One of the biggest computational costs in creating accurate simulations of real-world robots (and their real-world environments) are those costs associated with simulating the physics of mechanics and kinematics and dynamics in sensing and in acting (i.e. simulating the sampling of the ambient optic array by a video camera, or simulating the results of torque being applied by a motor). A number of independent third-party “middleware” software suppliers have developed general-purpose “physics engine” software libraries that can save much time and money in development of accurate simulations. Leading suppliers in this field are MathEngine (www.mathengine.com), Havok (www.havok.com), and Criterion Software (www.csl.com).

Despite the potential heavy computational cost of simulating phenomena that “come for free” when working with real robots, studies of simulated agents allow for rich streams of data to be collected, data that it may be impracticable or impossible to gather from a physical robot. Robot simulations also allow studies of failure modes that could be prohibitively expensive when working with real robots. For example, when developing flying-robot unmanned air vehicles (UAVs), many real-world failure modes will involve the loss (destruction) of the UAV, and a simulated crash is a much less costly event. A number of research teams have worked with well-validated simulations of real robots, where the lessons learnt in simulation have been demonstrated to be transferable to the real system. One notable body of work in this area was conducted by Jakobi (1997) for his PhD studies at Sussex. Jakobi developed a principled methodology for radically simplifying the computational cost of simulating agent-environment interactions, albeit one that is primarily applicable where those interactions are themselves simple: for realistically complex or dynamically varying interactions, significantly more work is required to establish how best to make computational savings in the simulation. K-Team offer a popular software simulator for evaluating *Khepera* control programs, and there exist a number of independent *Khepera* simulators (e.g. Jakobi *et al.*, 1995).

However, not all software agents are accurate models of physical robots: many software agents of interest have no need to accurately model real-world robots, or even real-world physics. It is reasonable to talk of non-physically-accurate software agents in two broad classes: abstract scientific agent-based models, and commercial engineering applications.

Agent-based models for scientific purposes are no less rigorous than simulated robot models, but they eliminate major computational costs by abstracting away from levels of analysis where detailed and accurate models of physical interaction are at all relevant. This is often allowable in cases where collective behaviour is the primary object of study. For example, in collective robotics, it is an item of faith that some form of inter-agent communication is useful for coordination among the group of agents. Many interesting but different forms of communication, or different constraints on the space of communicative behaviours, can in principle be explored in simulations where the software agents inhabit a world with minimally simple “laws of physics”. For example, a limited vocabulary of communicative utterances (grunts) could be modelled as simply emitting one of a small number of grunt-types, which are heard instantly by all nearby agents, without modelling any of the details of sound production, or sound-wave propagation in air, or auditory sensing of sound waves. Exactly this approach has proved very successful in the simulation-based scientific study of the development or evolution of a number of communication systems, including the evolution of human language use (McLennan & Burghardt, 1994; Noble, 2000; Noble *et al.*, 2001; Kirby, 2001). Two notable UK research clusters in this area are the BioSystems group at the Informatics Research Institute at the University of Leeds (www.scs.leeds.ac.uk/research/inf); and The Language Evolution and Computation group at Edinburgh (www.ling.ed.ac.uk/lec).

One very constrained form of inter-agent communication occurs in microeconomics, where traders interact within auction markets: i.e., buyers and sellers communicate by signalling prices of bids and offers. These and other abstract artificial economic systems have also been studied with some success using minimal simulation techniques (e.g., Epstein & Axtell, 1996); in the UK the use of agent-based simulations in economics and the social sciences has been pioneered by Gilbert’s team at Surrey (www.soc.surrey.ac.uk/research/cress).

Returning to the issue of abstracting away from accurate simulation of real-world physics, there are sound scientific models of agents moving over some area of space that pay little or no attention to modelling the physics of movement. For instance, a country-scale model of traffic flows across a highway network gains nothing by accurately simulating the physics of each car’s individual movement, provided that the abstractions in the model preserve a representation of phenomena important at a higher level, such as the fact that if one car hits another, both are likely to stop and (at least partially) block the road. Agent-based simulations of human activity have found increasing use over the last decade in health informatics applications (such as epidemiology and the associated prediction of healthcare demand for planning purposes) and in geographic information systems applications (such as those used to predict the spatial growth and spread of a city, and the effect that growth has on natural resources). The use of complex systems thinking in academic schools of business and management has also grown steadily over the past decade, and in America both the Santa Fe Institute (www.santafe.edu) and the New England Complex Systems Institute (www.necsi.org) appear to generate sizeable revenue from their offerings of business seminars and consultancies. In the UK, complex adaptive systems research applied to the sphere of business and management is well represented by the Complex Adaptive Systems Group at Oxford University’s Saïd Business School, by the Complexity Research Programme at the London School of Economics, and by NEXSUS, the Complex Systems Management Centre at Cranfield University School of Management.

There are many *potential* commercial or applied-engineering uses of autonomous software agents, but those *actually* deployed and making money are more rare. Currently one of the most lucrative markets is computer-based entertainment, where BICAS-type software agents have been used in computer games (e.g. Maes, 1995; Cliff & Grand, 1999) and in the animation of computer-generated characters for Hollywood movies (e.g. Reynolds' (1987) *Boids* algorithm was used for animating stampeding herds of animals in Disney's feature *The Lion King*). In recent years, global revenues from computer games have consistently exceeded global revenues from Hollywood movies, and the production costs of main-title computer games now routinely match those of medium-budget movies. A recent huge commercial success involving a computer-game application of simulated human agents is the *The Sims* series of games produced by Maxis (<http://thesims.ea.com/>), in which the user creates a number of human-like agents, designs their home, and then guides their relationships and careers. More abstract computer games, such as the perennially popular *SimCity* series of computer games (also produced by Maxis) in which the player takes the role of town planner and mayor for an abstract simulation model of a city, can also be classed as BICAS entertainment applications. In a style similar to these models developed for entertainment, commercial scientific modelling of real-world systems has recently turned to using autonomous software-agent techniques, for the predictive simulation of real-world events or scenarios. Examples include training police strategists in the prevention, containment, and control of crowds of rioters; or using demographic data and spatial geographic information models to predict the effects on revenue stream of relocating a factory or choosing a specific site for a new superstore. Prominent commercial companies in this space include the UK's GMAP (www.gmap.co.uk) and the US's Bios Group (www.biosgroup.com), who both offer, at various levels of abstraction, agent-based models of humans interacting in some space that represents a real-world geography or corporate organization, that can be used in management planning and training applications. Again, in such simulations, the small-scale interactions compound to give large-scale overall activity that is not readily predictable in advance.

However, not all software autonomous agents are designed to interact with a simulated environment that is intended to represent some real-world situation (or a realistically plausible but imaginary world, as is more often used in entertainment applications). Many researchers have studied the development of autonomous software agents intended to coordinate their perception and action in environments that are abstract "cyber-spaces", typically formed from a number of dynamic data-streams. One potential application area is for individual personalisable software agents that are instructed by a user to do that user's business on e-commerce sites such as online exchanges or auctions, simultaneously monitoring the bids and offers in multiple auctions so as to get the best deal; or possibly also so as to arbitrage across those auctions. Such an agent could potentially be simultaneously active in tens or hundreds of different auctions, where those auctions do not necessarily all operate according to the same rules and protocols. This is an application area recently explored by Byde, Preist, & Jennings (2002), albeit not using BICAS techniques. For several years, developing software agents for autonomous automated trading on the international financial markets has been a research topic pursued by a number of academic researchers, but the take-up of such technology for live applications by investment banks and financial exchanges appears to be very poor. Solid data on successfully fielded applications in financial trading are notoriously sparse, because the developer of any consistently profitable automated-trading method (agent-based or otherwise) has a manifest vested interest in keeping very quiet about that success, at least until they have banked enough money to retire comfortably. For reviews of BICAS-oriented approaches to the engineering design of trading agents, see (Cliff &

Bruten, 1999; and Tesfatsion, 2002). Prominent international research groups with a strong interest in artificial autonomous agents for business and e-commerce include academic groups at the MIT Media Lab and at Michigan University in the US, at Liverpool and Southampton Universities in the UK; at major industrial research labs such as IBM in New York State, and Hewlett-Packard Labs in Bristol UK; along with smaller commercial enterprises such as Frictionless Commerce (www.frictionless.com) in the US and LostWax (www.lostwax.com) in the UK. It is perhaps worth noting that in many of these groups the desire for raw profit typically takes a much higher precedence than any considerations of biological verisimilitude.

It has long been known from studies in experimental economics (e.g. Smith, 1962) that when groups of human traders come together in an appropriate free-market environment, the transaction-prices in the market can rapidly and reliably converge on the market's theoretical equilibrium price. This is the price at which the quantity supplied by the population of sellers best matches the quantity demanded by the population of buyers, and so represents an optimal allocation of those scarce resources that are supplied by the sellers and that are demanded by the buyers. This view of real-world free-market economies as resource allocation mechanisms is appealing because they are typically asynchronous and decentralised (in particular, they do not require a centralised auctioneer to orchestrate proceedings). Hence, they offer another metaphor from the natural world that can be used to influence the engineering design of distributed and decentralised systems where scarce resources are demanded by some population of consumers. For instance, in a networked computer facility, the scarce resources demanded by the users are likely to include processor time, disk space, and network bandwidth. If autonomous software agents are attached to each network resource, acting as sellers of the resource, and if autonomous software agents are also associated with each user's request for a job to be processed, then the agents can negotiate prices by, for instance, engaging in an auction. The intent is that at times of high demand the price of some of the facility's resources will rise, making them less attractive to some users, who hold off from consuming those resources until a period of reduced demand, during which the price lowers. This dynamic and decentralised market-based approach to computer load-balancing is one instance of a new approach to robustly solving dynamic resource-allocation problems, an approach known as *Market-Based Control* (MBC). Much of the groundwork for MBC was laid in the collection of papers edited by Huberman (1988), who pioneered MBC approaches while a researcher at Xerox PARC. More recently, Clearwater's (1996) collection includes accounts of a number of successful MBC systems, including distributed computer system load-balancing; industrial job-shop scheduling; and office-block air-conditioning management. Research groups with significant activities or investment in MBC include Southampton University and Hewlett-Packard Labs Bristol in the UK, and groups at the University of Michigan, University of Southern California, IBM T.J. Watson Research Labs, and Hewlett-Packard Labs Palo Alto in the USA.

3.2 Future Research Issues and Application Areas

Experiences to date indicate that for any artificial autonomous agent, robotic or virtual, parallel distributed processing architectures (such as artificial neural networks) offer many advantages over centralised sequential control programs. Experiences to date also demonstrate that purely manual design of such processing architectures is extremely difficult, because traditional

engineering design methodologies are not well-suited to the creation of asynchronous distributed networks of processors intended to operate without central control. Thus, the use of automated adaptation techniques, both within the lifetime of an agent and also over successive evolving generations of agents, remains the most promising approach to creating processing architectures. For this reason, biological metaphors such as adaptive artificial neural networks, and evolutionary computation techniques such as genetic algorithms are likely to remain strong influences in future BICAS research.

As the number of individual processing units (e.g. artificial neurons, or behaviour-generating modules) in an agent increases, it becomes increasingly difficult to specify an appropriate connectivity between the components in advance, and also to reconfigure the connectivity to account for component failure or malfunction. For this reason, ideas from developmental biology are likely to become more influential, as artificial autonomous agents undergo some kind of embryological morphogenesis process. The *Amorphous Computing* team at MIT produced a review of their pioneering work in this area (Abelson *et al.*, 2001).

Similar issues arise when using evolutionary optimisation techniques such as genetic algorithms (GAs) for semi-automated design of autonomous agent architectures. The space of possible genotypes in the system defines a space of possible designs to be explored. This definition is often made implicitly, via the specification of how the “agent genotype” genetic encodings operated on by the GA are interpreted as “agent phenotypes” in the evaluation of the genotype’s “fitness” value. Although past research has demonstrated many successful applications of GAs in the design both of robotic and software autonomous agents, the design of appropriate genetic encodings (and their associated mappings onto agent phenotypes via a morphogenesis process) and of productive fitness evaluation functions remains an *ad hoc* art, rather than an operationalised engineering discipline. This has long been recognised by GA-agent practitioners, but no clear solutions are yet in sight.

One final issue that is starting to cause some concern among practitioners and sympathetic observers of BICAS autonomous agent research is the relatively slow rate of increase in the desired or intended cognitive complexity of the autonomous agents studied; whether GA-evolved or hand-designed. As it is now over 15 years since the publication of Brooks’ papers that established the field of biologically-inspired behaviour-based systems, convincing excuses for not tackling more cognitively challenging problems than navigating an environment while avoiding collisions can no longer be based on appeals to the relative youth of the approach. The fear is that the BICAS approach is reaching an impasse similar to the one that occurred in traditional logic-based top-down AI around the time that Brooks wrote his seminal papers. One response to this is that it is a fear based on impatience and ignorance, symptomatic of failing to appreciate the inherent difficulty of creating artificial systems that can attain the cognitive complexity needed for even simple, restricted, task domains and environments. The counter to this response is that it is exactly the excuse made by practitioners of logic-based top-down AI.

While it is possible that economic constraints force researchers to employ robot platforms that are ill-suited to (or simply incapable of) use in studying more cognitively complex behaviours, experiences in humanoid robotics research seem to indicate that possession of sophisticated robot hardware is no panacea. Support for this hypothesis comes from the observation that the research team most actively trying to accelerate the cognitive complexity of tasks studied within BICAS-

agent research are using abstract idealised simulation studies (Beer, 1996, 2000), despite having previously worked with advanced robotic hardware.

Turning to potential lucrative applications of BICAS cognitive systems research, the continued development of behavior-based robots for niche applications, such as those pioneered by Brooks's company *iRobot*, seem set to continue, as does the use of robots as models for furthering our scientific understanding of cognition in animals (including us human animals). In the next decade or so, scientific applications of software-agent techniques may offer genuinely predictive computer-simulation models of a simple invertebrates (perhaps *C. elegans*; less possibly *D. melanogaster*, but see Hamahashi & Kitano, 1998), allowing accurate studies *in silico* of morphogenesis and development processes, and lifetime adaptation/habituation, and the interaction between complete sensory and motor systems.

However, in the immediate future the most promising applications for biologically-inspired software agents appear to remain in entertainment and leisure software. It does not require a crystal ball to predict that, very soon, "live" but entirely computer-generated versions of some sports will become available over internet broadband and/or broadcast TV and/or as a mobile phone handset content source; where the participants in the sports (e.g. soccer players, race-horses, car drivers, or robot warriors) are synthetic agents, possibly with BICAS architectures. These agents are trained and/or bred (evolved) by individual users/players/viewers (or networked syndicates of users/players/viewers) on their desktop PCs. The funding model would be based on income from online gambling; from provision of cheap "filler" content to TV broadcasters; from virtual track-side/pitch-side advertising hoardings; and from premium-rate phone lines used to provide a back-channel from viewers/participants into the broadcast. A clear precursor to this development is the *iRace* virtual horse-racing system planned as a joint venture in the UK by Telewest and VIS Entertainment, to be broadcast on Sky Digital TV: see www.irace.com.

Taking a slightly longer view, it seems plausible that stable market-based control systems, populated entirely by artificial agents, could be used for resource allocation in clustered compute facilities and potentially also in national GRID computer networks. It seems perfectly plausible that as federated networks of warehoused central computing facilities, housing many tens of thousands of server machines (all connected on an ultra-high-bandwidth network — providing what is referred to by IBM as "computing on demand" and by Hewlett-Packard as "utility data centres") come on-stream, they will incorporate BICAS-style "autonomic" or "adaptive infrastructure" technologies that provide self-healing resilience to load fluctuations, to component failures, and to attack by computer viruses and worms; and possibly also with market-based control for load-balancing and thermal resource management.

It also seems possible that, within the next decade, small-scale live trials of the first online international financial markets to be populated (at the point of execution) entirely by artificial autonomous trader-agents, might be operated by smaller "boutique" exchanges in major financial centres such as London and New York.

4 References

- Aarts, E., & Korst, J. 1988. *Simulated Annealing and Boltzmann Machines: A Stochastic Approach to Combinatorial Optimization and Neural Computing*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley and Sons Ltd.
- Abelson, H., Allen, D., Coore, D., Hanson, C., Rauch, E., Sussman, G.J., & Weiss, R., 2001. Amorphous Computing. *Communications of the ACM*, 43(5):74-82.
- Abraham, R. H., & Shaw, C. D. 1992. *Dynamics: The Geometry of Behavior* (Second Ed). Addison Wesley, Redwood City, CA.
- Adams, B., Breazeal, C., Brooks, R., & Scassellati, B. "Humanoid Robots: A New Kind of Tool", *IEEE Intelligent Systems*, July-August 2000.
- Agre, P. E., & Chapman, D. 1987. Pengi: An implementation of a theory of activity. In: *Proceedings AAAI-8*. San Mateo, CA. Morgan Kaufmann.
- Agre, P. E., & Chapman, D. 1990. What are plans for? In: Maes, P. (Ed.), *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 17--34. MIT/Elsevier, Cambridge MA.
- Aizawa, A. N., & Wah, B. W. 1994. Scheduling of genetic algorithms in a noisy environment. *Evolutionary Computation*, 2 (2), 97--122.
- Alcock, J. 1989. *Animal Behavior: An evolutionary approach*. Sinauer Associates.
- Altman, J. S., & Kien, J. 1989. New models for motor control. *Neural Computation*, 1, 173--183.
- Anderson, J. A., & Rosenfield, E. (Eds.). 1988. *Neurocomputing: Foundations of research*. MIT Press, Cambridge MA.
- Arbib, M. A. 1987. Levels of modelling of mechanisms of visually guided behaviour. *The Behavioral and Brain Sciences*, 10, 407--465.
- Arbib, M. A. (Ed.). 2003. *The Handbook of Brain Theory and Neural Networks*. Second edition. MIT Press Bradford Books, Cambridge MA.
- Ashby, W. R. (1952) *Design for a Brain*. London: Chapman & Hall.
- Axelrod, R. 1984. *The Evolution of Cooperation*. Basic Books, New York.
- Bagley, R. J., & Farmer, J. D. 1992. Spontaneous emergence of a metabolism. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Volume X of the series *Santa Fe Institute Studies in the Sciences of Complexity*, pages 93--140 Redwood City, CA. Addison-Wesley.
- Bagley, R. J., Farmer, J. D., & Fontana, W. 1992. Evolution of a metabolism In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Volume X of the series *Santa Fe Institute Studies in the Sciences of Complexity*, pages 141--158 Redwood City, CA. Addison-Wesley.
- Baker, G. L., & Gollub, J. P. 1990. *Chaotic dynamics: an introduction*. Cambridge University Press, Cambridge.
- Banzhaf, W. 1994. Self-organisation in a system of binary strings In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV*, pages 109--118 Cambridge MA. MIT Press Bradford Books.

- Barlow, H. 1989. Unsupervised learning. *Neural Computation*, 1, 295--311.
- Barnett, L. 1997. Tangled Webs: Evolutionary Dynamics on Fitness Landscapes with Neutrality. MSc Thesis, School of Cognitive and Computing Science, University of Sussex.
- Batali, J. 1994. Innate biases and critical periods: Combining evolution and learning in the acquisition of syntax. In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV*, pages 160--171 Cambridge, MA. MIT Press Bradford Books.
- Batali, J., & Kitcher, P. 1994. Evolutionary dynamics of altruistic behavior in optional and compulsory versions of the iterated prisoner's dilemma. In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV*, pages 343--348 Cambridge, MA. MIT Press Bradford Books.
- Beer, R. D. 1990. *Intelligence as Adaptive Behaviour: An Experiment in Computational Neuroethology*. Academic Press.
- Beer, R. D. 1995a. A dynamical systems perspective on agent-environment interaction *Artificial Intelligence*, 72, 173--215.
- Beer, R. D. 1995b. On the dynamics of small continuous-time recurrent neural networks *Adaptive Behavior*, 3(4), 471--511.
- Beer, R. D. 1996. Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 421--429. MIT Press Bradford Books, Cambridge MA.
- Belew, R. K., McInerney, J., & Schraudolph, N. N. 1992. Evolving networks: Using the genetic algorithm with connectionist learning. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Volume X of the series *Santa Fe Institute Studies in the Sciences of Complexity*, pages 511--548 Redwood City, CA. Addison-Wesley.
- Blumberg, B. 1994. Action-selection in hamsterdam: Lessons from ethology. In Cliff, D., Husband, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals To Animats 3, Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pages 108--117 Cambridge MA. MIT Press Bradford Books.
- Blumberg, B. 1996. *Old Tricks, New Dogs: Ethology and Interactive Creatures*. Ph.D. thesis, MIT Media Lab. Unpublished.
- Boden, M. A. 1996. *The Philosophy of Artificial Life*. Oxford Readings in Philosophy. Oxford University Press, Oxford.
- Booker, L. B., & Belew, R. K. (Eds.). 1991. *Proceedings of the 4th International Conference on Genetic Algorithms (ICGA91)*, San Mateo, CA. Morgan Kaufmann.
- Braitenberg, V. 1984. *Vehicles: Experiments in Synthetic Psychology*. MIT Press Bradford Books, Cambridge MA.
- Breazeal, C. 2002. *Designing Sociable Robots*. The MIT Press.
- Brooks, R. A. 1985. *A robust layered control system for a mobile robot*. AI Memo 864, MIT AI Lab.
- Brooks, R. A. 1986. *Achieving artificial intelligence through building robots*. AI Memo 899, MIT AI Lab.
- Brooks, R. A. 1989. A robot that walks: Emergent behaviours from a carefully evolved network. *Neural Computation*, 1, 253--262.
- Brooks, R. A. 1990. Challenges for complete creature architectures. In Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90)*, pages 434--433. MIT Press Bradford Books, Cambridge MA.
- Brooks, R. A. 1991a. Intelligence without reason. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pages 139--159 San Mateo, CA: Morgan Kaufmann.
- Brooks, R. A. 1991b. Intelligence without representation. *Artificial Intelligence*, 47, 139--159.

- Brooks, R. A. 1992. Artificial life and real robots. In Varela, F. J., & Bourgine, P. (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, pages 3--10 Cambridge MA. MIT Press Bradford Books.
- Brooks, R. A. 2002. *Robot: The Future of Flesh and Machines*. Allen Lane The Penguin Press.
- Brooks, R. A., & Maes, P. (Eds.). 1994. *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press Bradford Books, Cambridge MA.
- Brooks, R. A., & Flynn, A. M. 1989. Fast, cheap and out of control. *Journal of the British Interplanetary Society*, October, 1989. Also appears as *MIT AI Laboratory Memo 1186*, November, 1989. Artificial Intelligence Laboratory, Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts.
- Byde, A., Priest, C. & Jennings, N.R., 2002. Decision procedures for multiple auctions. *Proc. 1st Int Joint Conf. on Autonomous Agents and Multi-Agent Systems*, Bolgna, Italy, 613-620.
- Canavier, C. C., Baxter, D. A., Clark, J. W., & Byrne, J. H. 1993. Nonlinear dynamics in a model neuron provide a novel mechanism for transient synaptic inputs to produce long-term alterations of postsynaptic activity. *Journal of Neurophysiology*, 69(6), 2252--2257.
- Chapman, D. 1991. *Vision, Instruction, and Action*. MIT Press Bradford Books, Cambridge, MA.
- Charniak, E., & McDermott, D. 1985. *Introduction to Artificial Intelligence*. Addison Wesley.
- Chatfield, C., & Collins, A. J. 1980. *Introduction to Multivariate Statistics*. Chapman and Hall, London and New York.
- Clark, A. J. 1989. *Microcognition: Philosophy, Cognitive Science, and Parallel Distributed Processing*. MIT Press Bradford Books, Cambridge MA.
- Clark, A. J. 1996. *Being There*. MIT Press Bradford Books, Cambridge MA.
- Clearwater, S.H. 1996. *Market-Based Control: A Paradigm for Distributed Resource Allocation*. World Scientific.
- Cliff, D. 1990. *Computational neuroethology: A provisional manifesto*. Technical Report CSRP162, University of Sussex School of Cognitive and Computing Sciences.
- Cliff, D. 1991. Computational neuroethology: A provisional manifesto. In Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90)*, pages 29--39 Cambridge MA. MIT Press Bradford Books.
- Cliff, D. 1992. Neural networks for visual tracking in an artificial fly. In Varela, F. J., & Bourgine, P. (Eds.), *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.
- Cliff, D. 2003. Neuroethology, computational. In Arbib, M. A. (Ed.), *The Handbook of Brain Theory and Neural Networks*, Second edition, pages 737-741. MIT Press Bradford Books, Cambridge MA.
- Cliff, D., & Bullock, S. G. 1993. Adding 'foveal vision' to Wilson's animat. *Adaptive Behavior*, 2 (1), 47--70. Also available as University of Sussex School of Cognitive and Computing Sciences Technical Report CSRP263.
- Cliff, D., Harvey, I., & Husbands, P. 1993. Explorations in evolutionary robotics. *Adaptive Behavior*, 2 (1), 71--108.
- Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. (Eds.). 1994. *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*, Cambridge MA. MIT Press Bradford Books.

- Cliff, D., & Miller, G. F. 1995. Tracking the Red Queen: Measurements of adaptive progress in co-evolutionary simulations. In Morán, F., Moreno, A., Merelo, J. J., & P.Chacón (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)*, pages 200--218 Berlin. Springer-Verlag.
- Cliff, D., & Miller, G. F. 1996. Coevolution of pursuit and evasion II: Simulation methods and results. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 506--515. MIT Press Bradford Books, Cambridge MA.
- Cliff, D., & Ross, S. 1995. Adding temporary memory to ZCS. *Adaptive Behavior*, 3 (2), 101--150.
- Cliff, D., & Noble, J. 1997. Knowledge-based vision and simple visual machines. *Philosophical Transactions of the Royal Society of London: B*, **352**:1165-1175.
- Cliff, D. & Grand, S. 1999. The Creatures Global Digital Ecosystem. *Artificial Life* 5(1): 77-93.
- Cliff, D. & Bruten, J. 1999. Animat Market-Trading Interactions as Collective Social Adaptive Behavior. *Adaptive Behavior* 7(3&4):385—414.
- Clocksink, W. F., & Mellish, C. S. 1984. *Programming in Prolog* (Second (Ed.)). Springer-Verlag.
- Coderre, B. 1988. Modeling behavior in petworld. In Langton, C. (Ed.), *Artificial Life*, SFI Studies in the Sciences of Complexity, pages 407--420. Addison-Wesley Publishing Co. Inc.
- Coveney, P., & Highfield, R. 1995. *Frontiers of Complexity : the search for order in a chaotic world*. Faber and Faber, London.
- Cronin, H. 1993. *The Ant and the Peacock: Altruism and sexual selection from Darwin to today*. Cambridge University Press, Cambridge.
- Dautenhahn, K., Werry, I., Rae, J., Dickerson, P., Stribling, P., & Ogden, B., 2002. "Robotic Playmates: Analysing Interactive Competencies of Children with Autism Playing with a Mobile Robot" In: K Dautenhahn, A Bond, L Canamero, B Edmonds (eds.): *Socially Intelligent Agents - Creating Relationships with Computers and Robots*, Kluwer Academic Publishers.
- Davis, L. (Ed.). 1991. *Handbook of Genetic Algorithms*. Van Nostrand Reinhold, New York.
- Dawkins, R. 1982. *The Extended Phenotype*. Oxford University Press, Oxford.
- Dawkins, R. 1986. *The Blind Watchmaker*. Longman, Essex.
- Dawkins, R. 1989. *The Selfish Gene* (second edition). Oxford University Press, Oxford.
- Dreyfus, H. L. 1979. *What Computers Can't Do* (second edition). Harper and Row, New York.
- Dreyfus, H. L. 1981. From micro-worlds to knowledge representation: AI at an impasse. In Haugeland, J. (Ed.), *Mind Design: Philosophy, Psychology, Artificial Intelligence*, pages 161--204. MIT Press Bradford Books, Cambridge MA.
- Dreyfus, H. L., & Dreyfus, S. E. 1990. Making a mind versus modelling the brain: Artificial intelligence back at a branch-point. In Boden, M. A. (Ed.), *The Philosophy of Artificial Intelligence*, pages 309--333. Oxford University Press, Oxford, UK. First published in 1988.
- Edelman, G. M. 1988. *Topobiology: An Introduction to Molecular Embryology*. Basic Books, New York.
- Edelman, G. M. 1989. *Neural Darwinism: The Theory of Neuronal Group Selection*. Oxford University Press, Oxford.
- Eiben, A. E., van Kemenade, C. H. M., & Kok, J. N. 1995. Orgy in the computer: Multi-parent reproduction in genetic Algorithms. In Morán, F., Moreno, A., Merelo, J. J., & P.Chacón (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)*, pages 934--945 Berlin. Springer-Verlag.

- Elman, J. L. 1990. Finding structure in time. *Cognitive Science*, 14 (2), 179--211.
- Elman, J. L. 1993. Learning and development in neural networks: The importance of starting small. *Cognition*, 48 (1), 71-79.
- Elphick, M. R., Kemenes, G., Staras, K., & O'Shea, M. 1995. Behavioural role for nitric oxide in chemosensory activation of feeding in a mollusc. *Journal of Neuroscience*, 15 (11), 7653--7664.
- Elphick, M. R., Williams, L., & O'Shea, M. 1996. New features of the locust optic lobe: evidence of a role for nitric oxide in insect vision. *Journal of Experimental Biology*, 199, 2395--2407.
- Epstein, J.M., & Axtell, R.L., 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press.
- Fahlman, S. E. 1991. The recurrent cascade-correlation architecture In Lippmann, R. P., Moody, J. E., & Touretzky, D. S. (Eds.), *Advances in Neural Information Processing Systems 3*, pages 190--196. Morgan Kaufmann, Denver, CO.
- Fahlman, S. E., & Lebiere, C. 1990. The cascade-correlation learning architecture. In Touretzky, D. S. (Ed.), *Advances in Neural Information Processing Systems 2. Proceedings of the 1989 Conference*, pages 524--532. San Mateo, CA. Morgan Kaufmann.
- Farmer, J. D. 1991. A rosetta stone for connectionism. In Forrest, S. (Ed.), *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*, pages 153--187. MIT Press Bradford Books, Cambridge MA. Reprinted from *Physica D*, 42.
- Farmer, J. D., & Belin, A. 1992. Artificial life: The coming evolution. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 815--840 Redwood City, CA. Addison-Wesley.
- Ferrell, C. 1994a. Robust and adaptive locomotion of an autonomous hexapod. In Gaussier, P., & Nicoud, J. D. (Eds.), *Proceedings: From Perception to Action Conference (PerAc)*, pages 66--77. IEEE Computer Society Press, Los Alamitos, California.
- Ferrell, C. 1994b. Failure Recognition and Fault Tolerance of an Autonomous Robot. *Adaptive Behavior*, 2(4):375-398.
- Floreano, D., & Mondada, F. 1996. Evolution of plastic neurocontrollers for situated agents. In Maes, P., Matarić, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 402--410 Cambridge MA. MIT Press Bradford Books.
- Fogarty, T. C. 1989. Varying the probability of mutation in the genetic algorithm. In Schaffer, J. (Ed.), *Proceedings of the 3rd International Conference on Genetic Algorithms (ICGA89)*, pages 104--109 San Mateo, CA. Morgan Kaufmann.
- Fontana, W. 1992. Algorithmic chemistry. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 159--210 Redwood City, CA. Addison-Wesley.
- Forrest, S. 1991. *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*. MIT Press Bradford Books, Cambridge MA.
- Forrest, S. (Ed.). 1993. *Proceedings of the 5th International Conference on Genetic Algorithms (ICGA93)*. Morgan Kaufmann, San Mateo, CA.
- 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(1281):283--294.
- Gardner, H. 1985. *The Mind's New Science*. Basic Books, New York.

- Gilbert, S. F. 1991. *Developmental Biology* (Third edition). Sinauer Associates, Sunderland, MA.
- Gleick, J. 1987. *Chaos: Making a new science*. Viking, New York.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison Wesley.
- Gonzalez, A. J., & Dankel, D. D. 1993. *The Engineering of Knowledge-Based Systems*. Prentice-Hall International, Englewood Cliff, New Jersey.
- Grand, S., Cliff, D., & Malhotra, A. 1997. Creatures: Artificial life autonomous software agents for home entertainment. In Johnson, W. L. (Ed.), *Proceedings of the First International Conference on Autonomous Agents*, pages 22--29 New York. ACM Press.
- Grand, S., & Cliff, D. 1998 "Creatures: Entertainment Software Agents with Artificial Life". *Autonomous Agents and Multi-Agent Systems* 1(1):39-57.
- Grand, S. 2000. *Creation: Life and How to Make It*. Weidenfeld & Nicholson.
- Griffiths, A. J., Miller, J. H., Suzuki, D. T., Lewontin, R. C., & Gelbart, M. W. 1993. *An Introduction to Genetic Analysis* (Fifth edition). Freeman, New York.
- Grossberg, S. 1969. On learning and energy-entropy dependence in recurrent and nonrecurrent signed networks. *Journal of Statistical Physics*, 1, 319--350.
- Grossberg, S. 1988. Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks*, 1, 17--61.
- Gruau, F. 1994. Automatic definition of modular neural networks. *Adaptive Behavior*, 3 (2), 151--183.
- Gruau, F., & Whitley, D. 1993. Adding learning to the cellular development of neural networks: Evolution and the Baldwin effect. *Evolutionary Computation*, 1 (3), 213--234.
- Hamahashi, S. & Kitano, H., 1998. Simulation of **Drosophila** Embryogenesis. In *Proceedings of the Sixth International Conference on Artificial Life (AlifeVI)*, pp151-160, MIT Press.
- Hanson, S. J. 1991. A stochastic version of the delta rule. In Forrest, S. (Ed.), *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*, pages 265--272. MIT Press Bradford Books, Cambridge MA. Reprinted from *Physica D*, 42.
- Harvey, I. 1990. The artificial evolution of behaviour. In Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior (SAB90)*. Cambridge MA. MIT Press Bradford Books.
- Harvey, I. 1992a. The SAGA cross: the mechanics of crossover for variable-length genetic algorithms. In Männer, R., & Manderick, B. (Eds.), *Parallel Problem Solving from Nature*, 2, pages 269--278. North-Holland.
- Harvey, I. 1992b. Species adaptation genetic algorithms: A basis for a continuing SAGA. In Varela, F., & Bourgine, P. (Eds.), *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, pages 346--354. MIT Press Bradford Books, Cambridge MA.
- Harvey, I. 1993. Evolutionary robotics and SAGA: the case for hill crawling and tournament selection. In Langton, C. (Ed.), *Artificial Life 3 Proceedings*, pages 299--326. Santa Fe Institute Studies in the Sciences of Complexity, Proc. Vol. XVI, Addison Wesley.
- Harvey, I. 1996a. Is there another new factor in evolution? *Evolutionary Computation*, 4 (3), 313--329.
- Harvey, I. 1996b. Untimed and misrepresented: connectionism and the computer metaphor. *AISB Quarterly*, 96, 20--27.

- Harvey, I., Husbands, P., & Cliff, D. 1994. Seeing the light: Artificial evolution; real vision. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*, pages 392--401. MIT Press Bradford Books.
- Hebb, D. O. 1949. *The Organization of Behavior*. Wiley, New York.
- Hertz, J., Krogh, A., & Palmer, R. G. 1991. *An Introduction to the Theory of Neural Computation*. Lecture Notes Volume I. Addison Wesley.
- Hillis, W. D. 1985. *The Connection Machine*. MIT Press, Cambridge MA.
- Hillis, W. D. 1992. Co-evolving parasites improves simulated evolution as an optimization procedure. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 313--324. Addison-Wesley, Redwood City, CA.
- Hinton, G., & Nowlan, S. 1987. How learning can guide evolution. *Complex Systems*, 1, 495--502.
- Hinton, G. E. 1992. How neural networks learn from experience. *Scientific American*, 267 (3), 105--109.
- Hinton, G. E., & Anderson, J. A. (Eds.). 1989. *Parallel Models of Associative Memory*. Lawrence Erlbaum Associates. Updated edition of the original 1981 volume.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. 1986. Distributed Representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel Distributed Processing -- Explorations in the Microstructure of Cognition, Volume 1 -- Foundations*. A Bradford Book, MIT Press, 1986, pages 77--109. MIT Press, Cambridge, MA.
- Hodges, W. 1977. *Logic*. Penguin, London.
- Hodgkin, A. L., & Huxley, A. F. 1952. A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117, 500--544.
- Holland, J. H. 1975. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor. Since reprinted by MIT Press.
- Holland, J. H. 1992. Genetic algorithms. *Scientific American*, 267 (1), 66.
- Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. R. 1986. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge MA.
- Hopfield, J. J. 1984. Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences (USA)*, 81, 3088--3092.
- Hopfield, J. J. 1985. Neural computation on decisions in optimization problems. *Biological Cybernetics*, 52, 141--152.
- Horgan, J. 1995. From complexity to perplexity. *Scientific American*, 272 (6), 74--79.
- Horswill, I. D. 1993a. A simple, cheap, and robust visual navigation system. In Meyer, J.-A., Roitblat, H., & Wilson, S. (Eds.), *Proceedings of the Second International Conference on Simulation of Adaptive Behaviour (SAB92)*, pages 129--136. MIT Press Bradford Books, Cambridge, MA.
- Horswill, I. D. 1993b. *Specialization of Perceptual Processes*. Ph.D. thesis, MIT AI Lab.
- Hoyle, G. 1984. The scope of neuroethology. *The Behavioral and Brain Sciences*, 7, 367--412.
- Hubel, D. H. 1982. Explorations of the primary visual cortex, 1955-1978. *Nature*, 299, 515--524.

- Hubel, D. H. 1988. *Eye, Brain, and Vision*. W. H. Freeman, New York.
- Huberman, B.A., 1988. *The Ecology of Computation*. Elsevier/North-Holland.
- Husbands, P., Harvey, I., & Cliff, D. 1995. Circle in the round: State space attractors for evolved sighted robots. *Robotics and Autonomous Systems*, 15, 83--106.
- Husbands, P., Philippides, A., Smith, T.M.C. and O'Shea, M., 2001. Volume Signalling in Real and Robot Nervous Systems. *Theory in Biosciences*, 120:253-269.
- Huynen, M.A., Stadler, P.F., & Fontana, W. 1996. Smoothness Within Ruggedness: The Role of Neutrality in Adaptation. *Proc. Natl. Acad. Sci. (USA)* 93:397-401
- Hyvarinen, A., Karhunen, J., & Oja, E., 2001, *Independent Component Analysis*. John Wiley & Sons Inc.
- Jakobi, N., Husbands, P. & Harvey, I., 1995. Noise and the Reality Gap: The use of Simulation in Evolutionary Robotics. In *Advances in Artificial Life: Proc. 3rd European Conference on Artificial Life*, Moran, F., Moreno, A., Merelo, J., Chacon, P. (eds.) Springer-Verlag, Lecture Notes in Artificial Intelligence 929 (1995) pp. 704--720.
- Jakobi, N. 1997. Evolutionary Robotics and the Radical Envelope-of-Noise Hypothesis. *Adaptive Behavior*, 6 (2), 325-368.
- Johnson, M. B. 1991. *Build-a-dude: Action selection networks for computational autonomous agents*. Master's thesis, MIT Media Laboratory, Cambridge, MA. Unpublished.
- Johnson, M. B. 1995. *WavesWorld: A Testbed for Three Dimensional Semi-Autonomous Animated Characters*. Ph.D. thesis, MIT Media Laboratory, Cambridge, MA. Unpublished.
- Joyce, G. 1992. Directed molecular evolution. *Scientific American*, 267 (6), 90.
- Kauffman, S. 1991. Requirements for evolvability in complex systems: orderly dynamics and frozen components. In Forrest, S. (Ed.), *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*, pages 135--152. MIT Press Bradford Books, Cambridge MA. Reprinted from *Physica D*, 42.
- Kauffman, S. 1993. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.
- Kelly, K. 1994. *Out of Control: The New Biology of Machines*. Fourth Estate, London.
- Kirby, S. (2001). Spontaneous evolution of linguistic structure: an iterated learning model of the emergence of regularity and irregularity. *IEEE Transactions on Evolutionary Computation*, 5(2):102--110.
- Kitano, H. 1994. Evolution of metabolism for morphogenesis. In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV*, pages 49--58 Cambridge MA. MIT Press Bradford Books.
- Koch, C., & Segev, I. (Eds.). 1989. *Methods in Neuronal Modeling: From Synapses to Networks*. MIT Press Bradford Books, Cambridge MA.
- Koza, J. R. 1992. *Genetic Programming: On the programming of computers by means of natural selection*. MIT Press.
- Koza, J. R. 1994. *Genetic Programming II: Automatic Discovery of Reusable Programs*. MIT Press.
- Langton, C. G. 1989a. Artificial life. In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 1--47 Redwood City, CA. Addison-Wesley.
- Langton, C. G. (Ed.). 1989b. *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*.

- Langton, C. G. 1991. Computation at the edge of chaos: phase transitions and emergent computation. In Forrest, S. (Ed.), *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*, pages 12--37. MIT Press Bradford Books, Cambridge MA. Reprinted from *Physica D*, 42.
- Langton, C. G. 1993b. Editor's Introduction. *Artificial Life*, 1 (1), v-- viii.
- Langton, C. G. (Ed.). 1995. *Artificial Life: An Overview*. MIT Press Bradford Books, Cambridge MA.
- Lehky, S. R., & Sejnowski, T. J. 1988. Network model of shape from shading: neural function arises from both receptive and projective fields. *Nature*, 333, 452--454.
- Leith, P. 1990. *Formalism in AI and Computer Science*. Ellis Horwood, Chichester.
- Levy, S. 1993. *Artificial Life: The quest for a new creation*. Penguin.
- Lindgren, K. 1992. Evolutionary phenomena in simple dynamics In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 295--312 Redwood City, CA. Addison-Wesley.
- Lugowski, M. W. 1989. Computational metabolism: Towards biological geometries for computing. In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 341--368 Redwood City, CA. Addison-Wesley.
- MacGregor, R. J. 1987. *Neural and Brain Modeling*. Academic Press.
- MacLennan, B. J. 1992. Synthetic ethology: An approach to the study of communication. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 631--658 Redwood City, CA. Addison-Wesley.
- MacLennan, B. J., & Burghardt, G. M. 1994. Synthetic ethology and the evolution of cooperative communication. *Adaptive Behavior*, 2 (2), 161--188.
- Maes, P. 1989. How to do the right thing. *Connection Science*, 1 (3), 291--323.
- Maes, P. (Ed.). 1990a. *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. MIT/Elsevier, Cambridge MA.
- Maes, P. 1990b. Situated agents can have goals . In Maes, P. (Ed.), *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pages 49--70. MIT/Elsevier, Cambridge MA.
- Maes, P., 1995. Artificial Life Meets Entertainment: Lifelike Autonomous Agents. *Communications of the Association for Computing Machinery*, 38(11):108--114.
- Marr, D. 1982. *Vision*. Freeman.
- Matarić, M. J. 1990. Navigating with a rat brain. In Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90)*, pages 169--175 Cambridge MA. MIT Press Bradford Books.
- Matarić, M. J., & Cliff, D. 1996. Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19 (1), 67--83.

- Mayley, G. 1996a. The evolutionary cost of learning. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 458--467 Cambridge MA. MIT Press Bradford Books.
- Mayley, G. 1996b. Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4 (3), 213--234.
- McCarthy, J., Minsky, M. L., Rochester, N., and Shannon, C. E. 1955. *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. Copy available at: <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>
- McClelland, J. L., & Rumelhart, D. E. (Eds.). 1986. *Parallel Distributed Processing. Volume 2: Psychological and biological models*. MIT Press Bradford Books, Cambridge MA.
- McClelland, J. L., & Rumelhart, D. E. 1988. *Explorations in Parallel Distributed Processing*. MIT Press Bradford Books, Cambridge MA.
- McClelland, J. L., Rumelhart, D. E., & Hinton, G. E. 1986. The appeal of parallel distributed processing. In Rumelhart, D. E., & McClelland, J. L. (Eds.), *Parallel Distributed Processing, Volume 1: Foundations*, pages 3--44. MIT Press Bradford Books, Cambridge MA.
- McCulloch, W. S., & Pitts, W. 1943. A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115--133.
- McDermott, D. 1981. Artificial intelligence meets natural stupidity. In Haugeland, J. (Ed.), *Mind Design: Philosophy, Psychology, Artificial Intelligence*, pages 143--160. MIT Press Bradford Books, Cambridge MA. (Reprinted from SIGART Newsletter 57, April 1976).
- McDermott, D. 1987. A critique of pure reason. *Computational Intelligence*, 3 (3), 151--237.
- McFarland, D. 1990. What it means for robot behaviours to be adaptive. In Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (SAB90)*, pages 22--28 Cambridge MA. MIT Press Bradford Books.
- Miller, G. F. 1995. *Artificial life as theoretical biology: How to do real science with computer simulation*. Technical report CSRP378, University of Sussex School of Cognitive and Computing Sciences.
- Miller, G. F. 2000. *The Mating Mind: How Sexual Choice Shaped the Evolution of Human Nature*. Heinemann.
- Miller, G. F., & Cliff, D. 1994. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB94)*, pages 411--420. MIT Press Bradford Books, Cambridge MA.
- Miller, G. F., Todd, P. M., & Hegde, S. U. 1989. Designing neural networks using genetic algorithms In Schaffer, J. D. (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*, pages 379--384 San Mateo, California. Morgan Kaufmann.
- Miller, W. T., Sutton, R. S., & Werbos, P. J. (Eds.). 1990. *Neural Networks for Control*. MIT Press.
- Minsky, M., & Papert, S. 1969. *Perceptrons : an introduction to computational geometry*. MIT Press, Cambridge MA.
- Mitchell, M. 1998. *An Introduction to Genetic Algorithms*. MIT Press Bradford Books.
- Mitchell, T. 1997. *Machine Learning*. McGraw-Hill.
- Moravec, H. P. 1983. The Stanford Cart and the CMU Rover. *Proceedings of the IEEE*, 71 (7), 872--884.

- Moravec, H. P. 1989. Human culture: A genetic takeover underway. In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 167--200 Redwood City, CA. Addison-Wesley.
- Morowitz, H. J. 1993. Artificial biochemistry: Life before enzymes. In Langton, C. G. (Ed.), *Artificial Life III: Proceedings of the workshop on Artificial Life*, June 1992, volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 381--388. Reading, MA, USA. Addison-Wesley.
- Newell, A., & Simon, H. A. 1976. Computer science as empirical enquiry: Symbols and search. *Communications of the Association for Computing Machinery*, 19 (3), 113--126.
- Nilsson, N. J. 1982. *Principles of Artificial Intelligence*. Springer Verlag.
- Nilsson, N. J. 1984. Shakey the robot. Technical note 323, SRI International, Menlo Park, California.
- Noble, J., & Cliff, D. 1996. On simulating the evolution of communication. In Maes, P., Matarić, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 608--617. MIT Press Bradford Books, Cambridge MA.
- Noble, J. 2000. Co-operation, competition and the evolution of pre-linguistic communication. In C. Knight, J.R. Hurford & M. Studdert-Kennedy, editors, *The Evolutionary Emergence of Language: Social Function and the Origins of Linguistic Form*. Cambridge: Cambridge University Press.
- Noble, J.J; Di Paolo, E A; Bullock, S. 2001. Adaptive factors in the evolution of signalling systems. In: Cangelosi, A & Parisi, D (editors) *Simulating the Evolution of Language*, pp. 53-78 Springer-Verlag.
- Nolfi, S., & Parisi, D. 1995a. Evolving artificial neural networks that develop in time. In Moran, F., Moreno, A., Merelo, J. J., & P. Chacon (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)*, pages 351--367 Berlin. Springer-Verlag.
- Nolfi, S., & Parisi, D. 1995b. 'Genotypes' for neural networks. In Arbib, M. A. (Ed.), *The Handbook of Brain Theory and Neural Networks*, pages 431--434. MIT Press Bradford Books.
- Nolfi, S., & Parisi, D. 1997. Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, 5 (1), 75--98.
- Packard, N. H. 1989. Evolving bugs in a simulated ecosystem. In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 141--156 Redwood City, CA. Addison-Wesley.
- Parisi, D., Cecconi, F., & Nolfi, S. 1990. Econets: Neural networks that learn in an environment. *Network*, 1, 149--168.
- Philippides, A.O., Husbands, P. and O'Shea, M. 2000. Four-dimensional neuronal signaling by nitric oxide: A computational analysis. *Journal of Neuroscience*, 20(3):1199-1207.
- Port, R., & van Gelder, T. (Eds.). 1995. *Mind as Motion: Explorations in the Dynamics of Cognition*. MIT Press Bradford Books, Cambridge MA.
- Rasmussen, S. 1989. Toward a quantitative theory of the origin of life In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 79--104 Redwood City, CA. Addison-Wesley.
- Rasmussen, S., Knudsen, C., Feldberg, R., & Hindsholm, M. 1991. The coreworld: emergence and evolution of cooperative structures in a computational chemistry. In Forrest, S. (Ed.), *Emergent Computation: Self-Organizing, Collective, and Cooperative Phenomena in Natural and Artificial Computing Networks*, pages 111--134. MIT Press Bradford Books, Cambridge MA. Reprinted from *Physica D*, 42.

- Ray, T. S. 1992. An approach to the synthesis of life. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 371--408 Redwood City, CA. Addison-Wesley.
- Ray, T. S. 1994. An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artificial Life*, 1 (1/2), 179--210.
- Reynolds, C. 1987. Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, 21 (4), 25--34.
- Ridley, M. 1993. *The Red Queen: Sex and the evolution of human nature*. Viking, London.
- Rosenblatt, F. 1962. *Principles of Neurodynamics*. Spartan, New York.
- Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. 1986a. A general framework for parallel distributed processing. In Rumelhart, D. E., & McClelland, J. L. (Eds.), *Parallel Distributed Processing, Volume 1: Foundations*, pages 45--76. MIT Press Bradford Books, Cambridge MA.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. 1986b. Learning internal representations by error propagation. In Rumelhart, D. E., & McClelland, J. L. (Eds.), *Parallel Distributed Processing, Volume 1: Foundations*, pages 318--362. MIT Press Bradford Books, Cambridge MA.
- Rumelhart, D. E., & McClelland, J. L. (Eds.). 1986. *Parallel Distributed Processing, Volume 1: Foundations*. MIT Press Bradford Books, Cambridge MA.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. 1986. Schemata and sequential thought processes in PDP models. In McClelland, J. L., & Rumelhart, D. E. (Eds.), *Parallel Distributed Processing, Volume 2: Psychological and Biological Models*, pages 7--57. MIT Press Bradford Books, Cambridge MA.
- Rumelhart, D. E., & Zipser, D. 1986. Feature discovery by competitive learning. In Rumelhart, D. E., & McClelland, J. L. (Eds.), *Parallel Distributed Processing, Volume 1: Foundations*, pages 151--193. MIT Press Bradford Books, Cambridge MA.
- Scassellati, B., 2001. "Investigating Models of Social Development Using a Humanoid Robot," in Barbara Webb & Thomas Consi, eds., *Biorobotics*, MIT Press.
- Schuster, P. 1992. Complex optimization in an artificial RNA world. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 277--292 Redwood City, CA. Addison-Wesley.
- Schuster, P. 1993. Extended molecular evolutionary biology: Artificial life bridging the gap between chemistry and biology. *Artificial Life*, 1(1), 39--60.
- Schuster, P. 1995. Artificial life and molecular evolutionary biology. In Moran, F., Moreno, A., Merelo, J. J., & P. Chacon (Eds.), *Advances in Artificial Life: Proceedings of the Third European Conference on Artificial Life (ECAL95)*, pages 3--19 Berlin. Springer-Verlag.
- Searle, J. R. 1980. Minds, brains and programs. *The Behavioral and Brain Sciences*, 3, 417--424.
- Sejnowski, T. J., Koch, C., & Churchland, P. S. 1988. Computational neuroscience. *Science*, 241, 1299--1306.
- Sharkey, A. (Ed.). 1997. *Connection Science* 9 (1): special issue on Combining Artificial Neural Nets: Modular Approaches.
- Sigmund, K. 1993. *Games of Life: Explorations in Ecology, Evolution, and Behavior*. Oxford University Press, Oxford.
- Simmons, P.J. & Young, D. 1999. *Nerve Cells and Animal Behaviour*. Cambridge University Press, Cambridge.

- Simon, H. A. 1969. *Sciences of the Artificial*. MIT Press, Cambridge MA.
- Sims, K. 1994. Evolving 3D morphology and behavior by competition. In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV*, pages 28–39 Cambridge MA. MIT Press Bradford Books.
- Sims, K. 1995. Evolving 3D morphology and behavior by competition. *Artificial Life*, 1 (4), 353–372.
- Skarda, C. A., & Freeman, W. J. 1987. How brains make chaos in order to make sense of the world. *The Behavioral and Brain Sciences*, 10, 161–195.
- Slocum, A.C., Downey, D.C. and Beer, R.D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In J. Meyer, A. Berthoz, D. Floreano, H. Roitblat and S. Wilson (Eds.), *From Animals to Animats 6: Proceedings of the Sixth International Conference on Simulation of Adaptive Behavior* (pp. 430-439). MIT Press.
- Sloman, A. 1990. Motives, mechanisms, emotions. In Boden, M. A. (Ed.), *The Philosophy of Artificial Intelligence*, Oxford Readings in Philosophy, pages 231–247. Oxford University Press, Oxford.
- Sloman, A., & Croucher, M. 1981. Why robots will have emotions. In Hayes, P. J. (Ed.), *Proceedings of the 7th International Joint Conference on Artificial Intelligence (IJCAI '81)*, pages 197–202 Los Altos, CA. Morgan Kaufmann.
- Smith, V.L., 1962. An Experimental Study of Competitive Market Behavior. *Journal of Political Economy*. 70:111 – 137.
- Smithers, T. 1992. Taking eliminative materialism seriously: A methodology for autonomous systems research. In Varela, F. J., & Bourgine, P. (Eds.), *Towards a Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life (ECAL91)*, pages 31–40. MIT Press Bradford Books, Cambridge, MA.
- Smithers, T. 1995. Are autonomous agents information processing systems? In Steels, L., & Brooks, R. (Eds.), *The Artificial Life Route to Artificial Intelligence*. Lawrence Erlbaum, Hillsdale, New Jersey.
- Smolensky, P. 1988. On the proper treatment of connectionism. *The Behavioral and Brain Sciences*, 11, 1–74.
- Srinivasan, M. V., Chahl, J.S., Nagle, M.G., & Zhang, S.W. 1997. Embodying natural vision into machines. In: *From Living Eyes to Seeing Machines*, M.V. Srinivasan and S. Venkatesh (eds), Oxford University Press, U.K., pp. 249-265.
- Srinivasan, M. V., Chahl, J.S., Weber, K., Venkatesh, S., Nagle, M.G., & Zhang, S.W., 1998. Robot navigation inspired by principles of insect vision. In: *Field and Service Robotics*, A. Zelinsky (ed), Springer Verlag, Berlin, New York, 12-16.
- Stanley, E. A., Ashlock, D., & Tesfatsion, L. 1993. Iterated prisoner's dilemma with choice and refusal of partners. In Langton, C. G. (Ed.), *Artificial Life III: Proceedings of the workshop on Artificial Life*, June 1992, volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 131 – 176. Reading, MA, USA. Addison-Wesley.
- Steels, L. 1993. The artificial life roots of artificial intelligence. *Artificial Life*, 1 (1), 75–110.
- Stone, J.V., 2002, Independent Component Analysis: An Introduction. *Trends in Cognitive Sciences*, 6(2), pp59-64, Feb 2002.
- Stork, D., Jackson, B., & Walker, S. 1992. 'Non-Optimality' via pre-adaptation in simple neural systems. In Langton, C., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 409–429. Addison-Wesley, Redwood City, CA.
- Sutton, R. S. 1984. *Temporal credit assignment in reinforcement learning*. Ph.D. thesis, University of Massachusetts, Amherst, MA. (COINS Technical Report 84-02).
- Sutton, R. S. 1988. Learning to predict by the methods of temporal difference. *Machine Learning*, 3, 9–44.

- Sutton, R. S., & Barto, A. G. 1987. A temporal-difference model of classical conditioning. In *Proceedings of the ninth Annual Conference of the Cognitive Science Society*. Erlbaum.
- Tamayo, P., & Hartman, H. 1989. Cellular automata, reaction-diffusion systems, and the origin of life. In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 105--124 Redwood City, CA. Addison-Wesley.
- Tesfatsion, L., 2002, Agent-Based Computational Economics: Growing Economies from the Bottom Up, *Artificial Life*, 8(1):55-82.
- Thelen, E., & Smith, L. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press Bradford Books.
- Todd, P. M., & Miller, G. F. 1993. Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior*, 2 (1), 5--47.
- Travers, M. 1988. Animal construction kits. In Langton, C. (Ed.), *Artificial Life*, SFI Studies in the Sciences of Complexity, pages 421--442. Addison-Wesley Publishing Co. Inc.
- Tyrrell, T. 1993. The use of hierarchies for action selection. *Adaptive Behavior*, 1 (4), 387--420.
- Tyrrell, T. 1994. An evaluation of Maes's bottom-up mechanism for behavior selection. *Adaptive Behavior*, 2 (4), 307--348.
- Van Valen, L. 1973. A new evolutionary law. *Evolutionary Theory*, 1, 1--30.
- Waldrop, W. M. 1993. *Complexity: the emerging science at the edge of order and chaos*. Viking.
- Wallich, P. 1991. Silicon babies. *Scientific American*, 265 (6), 82--91.
- Walter, W.G. 1950. An imitation of life. *Scientific American* 182(5): 42--45.
- Walter, W.G. 1951. A machine that learns. *Scientific American* 185(2): 60--63.
- Watkins, C. J. C. H. 1989. *Learning from Delayed Rewards*. Ph.D. thesis, University of Cambridge. Unpublished.
- Watkins, C. J. C. H., & Dyan, P. 1992. Technical note: Q-learning. *Machine Learning*, 8, 279--292.
- Webb, B. 2000. What does robotics offer animal behaviour? *Animal Behaviour*, 60:545-558.
- Webb, B. 2002. Robots in invertebrate neuroscience. *Nature* 417:359-363
- Webb, B. 2003. Can robots make good models of biological behaviour? *The Behavioural and Brain Sciences* 24 (6).
- Weiner, N. 1948. *Cybernetics: or Control and Communication in the Animal and the Machine* (second edition). MIT Press, Cambridge MA.
- Werbos, P. J. 1994. *The roots of backpropagation: from ordered derivatives to neural networks and political forecasting*. Wiley, New York.
- Werner, G. M., & Dyer, M. G. 1992. Evolution of communication in artificial organisms. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, volume X of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 659--688 Redwood City, CA. Addison-Wesley.
- White, H. 1989. Learning in artificial networks: A statistical perspective. *Neural Computation*, 1, 415--424.

- Whitley, L. D. 1989. The GENITOR algorithm and selection pressure: Why rank-based allocation of reproductive trials is best. In Schaffer, J. D. (Ed.), *Proceedings of the 3rd International Conference on Genetic Algorithms (ICGA89)*, pages 116--123. Morgan Kaufmann.
- Williamson, M. M., 1999. *Robot Arm Control Exploiting Natural Dynamics*, Ph.D. Thesis, MIT Department of Electrical Engineering and Computer Science.
- Wilson, S. W. 1985. Knowledge growth in an artificial animal. In Grefenstette, J. J. (Ed.), *Proceedings of an International Conference on Genetic Algorithms and their Applications (ICGA85)*, pages 16--23 Hillsdale: New Jersey. Lawrence Erlbaum Associates.
- Wilson, S. W. 1994. ZCS: A zeroth level classifier system. *Evolutionary Computation*, 2 (1), 1--18.
- Wilson, S. W. 1995. Classifier fitness based on accuracy. *Evolutionary Computation*, 3 (2), 149--175.
- Wilson, S. W., & Goldberg, D. E. 1989. A critical review of classifier systems. In Schaffer, J. D. (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*, pages 244--255. Morgan Kaufmann.
- Winograd, T., & Flores, F. 1986. *Understanding Computers and Cognition: A New Foundation for Design*. Addison-Wesley Publishing Co. Inc.
- Winston, P. H. 1984. *Artificial Intelligence* (Second (Ed.)). Addison-Wesley Publishing Co. Inc.
- Winston, P. H., & Horn, B. 1980. *LISP*. Addison-Wesley Publishing Co. Inc.
- Yaeger, L. 1993. Computational genetics, physiology, metabolism, neural systems, learning, vision, and behavior or PolyWorld: Life in a new context. In Langton, C. G. (Ed.), *Artificial Life III: Proceedings of the workshop on Artificial Life*, June 1992, volume XVII of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 263--298 Reading, MA, USA. Addison-Wesley.
- Yamauchi, B. M., & Beer, R. D. 1994. Sequential behavior and learning in evolved dynamical neural Networks. *Adaptive Behavior*, 2 (3), 219--246.
- Zaera, N., Cliff, D., & Bruten, J. L. 1996. (Not) evolving collective behaviors in synthetic fish. In Maes, P., Matarić, M. J., Meyer, J.-A., Pollack, J., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior (SAB94)*, pages 635--644 Cambridge MA. MIT Press Bradford Books.
- Zeleny, M., Klir, G. J., & Hufford, K. D. 1989. Precipitation membranes, osmotic growths, and synthetic biology. In Langton, C. G. (Ed.), *Artificial Life: Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Los Alamos, NM, Sept, 1987, volume 6 of *Santa Fe Institute Studies in the Sciences of Complexity*, pages 125--140 Redwood City, CA. Addison-Wesley.