Intelligent systems: towards a new synthetic agenda

D L Warnett

Ph.D.
University of Edinburgh
1999

"..there is an inherent rationality to life that makes it intelligible at a much deeper level than functional utility and historical accident." (Goodwin, 1994, p. 105)

Abstract

This thesis adopts a *complex systems stance* towards the origin and growth of intelligence in biological and artificial systems at both reactive and reflective levels of adaptation. A *synthetic* strategy is employed dually informed by characterisations of complex biological systems together with the engineering of embodied artificially intelligent systems.

Following a review of the two dominant, and often construed as mutually exclusive, "rule-based" and "behaviour-based" characterisations of systems it will be argued that neither provide a comprehensive account of the growth and organisation of complex biological systems, and that both result in brittle, hand-crafted and, critically, pre-interpreted artificial systems. Furthermore, it will be suggested that the growing consensus that this impasse can be resolved by conjoining conventional reactive and reflexive components in a single architecture should be rejected. Such considerations indicate that a new agenda is required.

A *synthetic intelligence* stance is adopted herein, motivated by convergent biological and engineering approaches within a dynamic systems framework which emphasises system *self-organisation*. In this context some new experiments on mobile robot navigation are reported which demonstrate the feasibility of the synthetic approach. The implications of this novel stance are discussed with particular emphasis on the future of self-organising artificial systems.

Acknowledgements

I would like to thank all my friends, colleagues, and family for all their support and toleration over the past four years. Specifically I would like to thank: Ben, for invaluable advice and discussions in the early days; Beth, for advice and support over the last year; Della, for conversations and coffee breaks beyond the call of duty; Esther, for her attempts to maintain my nutrient balance; John, for his positivity, support, and comradeship; Jon, for friendship and support over the past seven years; Maggie, for advice, inspiration, and conversation; and my brother, Stephen, for spell-checking chapters 3 — 6 at extremely short notice.

I am indebted to my supervisor, Brendan McGonigle, for revolutionising my view of psychology and thereby providing the inspiration for this research; for support and guidance throughout this "effort after meaning", and of course for endless pints of Guinness.

I would like to thank W. S. Stewart for his bequest to the University of Edinburgh which has made this research possible.

Finally, I would like to thank my mother for endless and invaluable support throughout; for her confidence and faith; and for teaching me always to think.

Declaration

I hereby declare that this thesis has been composed by myself, and that the work presented herein is my own.

D L Warnett
Edinburgh
April 10, 2000

Contents

List of Figures

Prologue

The concern of this thesis is the origin and growth of cognition in biological and artificial systems. Three interlinked positions are advocated: a *complex systems* stance; a *dynamical systems* perspective; and a *synthetic* strategy.

The complex systems stance eschews the traditional reductionist mode of enquiry that tends, within psychology, to focus exclusively on low-level, reactive behaviours. Instead it is argued that an understanding of systems is best achieved by investigating, and attempting to replicate, the competences of complex systems *directly*. Such a position requires accurate and comprehensive characterisations of complex biological systems. An appreciation that complex systems retain many reactive features (especially important to self-preserving behaviours), and that simple systems incorporate many precursors of more complex competences is fundamental.

In common with a growing movement (van Geert, 1994; van Gelder & Port, 1995; McGonigle & Chalmers, 1998*a*, for examples) towards a dynamical perspective in terms of both empirical observation of biological systems and associated model construction, it is argued that the fundamentally dynamic nature of cognition, and cognitive development, must be incorporated within both characterisations of biological systems, and engineering of artificial systems. The dynamical perspective holds that systems should be regarded as essentially, and influentially, *embodied* and *situated*. Furthermore, it provides a rich language with which to describe system organisation and change, and stresses that initial system conditions (de Lorenzana & Ward, 1987), in combination with behavioural mechanisms common to all systems, serve to constrain possible developmental trajectories. Critically, the dynamic stance emphasises system *self-organisation* over the whole life-span, indicating that an open-ended, life-historical approach must be taken both when characterising biological systems, and when constructing artificial robotic systems (Chalmers & McGonigle, 1997; McGonigle, 1998).

Finally, a synthetic strategy is advocated which extends Braitenburg's (1984) concept of 'uphill analysis, and downhill synthesis'. It is suggested that the most promising approach to the understanding of systems is a dialectic strategy of theory construction and refinement informed by characterisations of complex biological systems in their own right, in conjunction with reverse-engineering, and the principled construction of artificial systems.

The central theoretical arguments of, and robotic implementation described within, the thesis are not constructed to be falsifiable in the Popperian (1963) sense — recent calls for quantitative analyses of robotic architectures (Duckett & Nehmzow, 1997, for example) are suggested to be premature partly because of the currently limited range and complexity of robotic implementations available for comparison, but more importantly because the mapping of theory to model within robotics is often either implicit and therefore opaque, or lacks coherency. Furthermore, no consensus exists within the field with respect to the kinds of benchmarks of achievement which might form the bases of comparison. For these reasons the thesis should be read within the philosophical perspectives of Lakatos (1978) and Kuhn (1962)[1].

The initial theoretical portion of the research strives to demonstrate that both traditional stances within cognitive science and artificial intelligence can be regarded, in many important respects, as degenerating research programmes (Lakatos, 1978) — with the recently developed dynamical and situated stances, although providing critical theoretical advances, inspiring artificial implementations of only small improvement over the traditional approaches. Currently, following the growing prominence of dynamical and situated perspectives, the conceptual terrain is confused, both in terms of biological characterisations and, especially, the resulting prescriptions for the construction of artificial systems.

The remainder of the theoretical section of the thesis expounds a synthetic theory of intelligent systems, grounded in the reverse engineering of complex biological systems, in more detail and thus advocates a "paradigm shift" (Kuhn, 1962) towards a synthetic, complex-systems research programme. The review of studies of complex biological systems suggests a number of essential features of such systems which could be used as metrics to compare and contrast robotic architectures motivated by varying theoretical stances. The experimental section of the thesis describes an architecture which is designed to incorporate some of these features, whilst providing for the easy addition of others at a later date, and some preliminary experiments in the field of mobile robot

[1] See section 6.3.4 for further discussions of these points.

navigation which are designed to illustrate the ability of this novel stance in supporting the development of adaptive artificial agents.

The following *preçis* aims to elaborate some of the main themes which will be discussed throughout the thesis, and to trace the path that will be taken through the conceptual jungle.

Preçis

Chapter one

Classical, or "symbol-based", and "behaviour-based" characterisations of biological systems are analysed and their implications for the construction of artificial systems addressed. Classical (symbolic representational) stances might appear superficially to be promising, focusing, as they do, on complex competences as the critical indices of systems. It is argued, however, that these stances tend to lead to static characterisations of ungrounded *preadapted* cognitive agents, and brittle, semantically-blind artificial agents. Behaviour-based stances tend to adopt a reductionist strategy, focusing on observable reactive behaviours, and associationistic principles of behaviour modification. The competences of complex systems are not addressed directly which means that a plausible account of the development of complex competences from their more simple precursors must be provided. It is argued herein that behaviour-based stances fail to provide a plausible account of such development in biological systems and, furthermore, are incapable of supporting the progressive development of artificial systems. In both cases the 'scaling-up' problem remains unresolved.

The most obvious solution to this impasse is hybridisation of conventional reflective and reactive components: ideally such a system would incorporate both deliberative and reactive components thereby balancing internal goal-directedness with sensitivity to the dynamic external world. It is suggested, however, that such hybridisation is doomed to failure from the outset since although these two stances are often construed as mutually exclusive both result in artificial systems which are brittle, hand-crafted, preinterpreted, and semantically-blind. Hybridisation of the two can result only in systems which are preinterpreted and hand-crafted at *both* representational and behavioural levels.

In recent years a consensus has begun to emerge within psychology, cognitive science, and artificial intelligence, that neither the classical symbol-based nor behaviour-based approaches constitute an adequate theory of biological systems, and have been singularly ineffective within the artificial domain. Likewise, their conjunction has produced few clear benefits. It is suggested that theorising needs to progress from all mind (logic and symbols) or all behaviour (acquisition/adjustment of tight couplings) characterisations towards a more inclusive view of system competences firmly grounded in accurate characterisations of complex biological systems.

Chapter two

Dynamic and situated perspectives on cognition, both of which have been proposed as potential solutions to the impasse, are reviewed with regard to their application to cognitive systems, and their prescriptions for the implementation of artificial systems.

The major contribution of the dynamic perspective is suggested to be a rich metaphorical language for describing complex systems over time, and an emphasis on the open-ended self-organisation of systems from their lower ontological bounds over development. It is argued, however, that with regard to the construction of artificial systems the dynamic approach has been embraced only by the ALIFE community and that this represents, as currently formulated, little improvement over traditional behaviour-based approaches.

It is argued that the situated conception of intelligence, grounded in naturalistic and evolutionary epistemology, laudably stresses dynamic agent-environment interaction as the root of meaning for systems, suggesting a middle-ground between the pre-interpreted domains of symbol-based and behaviour-based stances. A brief examination of situated agent research reveals some advances over classical symbolic and behaviour-based implementations yet, it is concluded, the situated perspective, as currently interpreted by researchers within artificial intelligence, yields few coherent principles for the design of artificial intelligent systems.

Although both dynamic and situated approaches make important theoretical contributions neither, *in isolation*, have yet supported the development of increasingly adaptive artificial intelligent systems.

Chapter three

Given that both classical symbol-based and behaviour-based characterisations have failed to deliver complex artificial systems, and that dynamic and situated perspectives alone have currently failed to provide an agenda for the development of artificial systems and have motivated artificial implementations of only little improvement, it is argued that a synthetic strategy of mutually informing biological characterisations, and engineering must be adopted. The first section of this chapter discusses one such biological characterisation and the resulting prescriptions for the design of artificial systems, whilst the second describes some robotic implementations motivated by the biological characterisation.

Biological characterisation Initially it is suggested that comparative psychology has failed to adequately characterise the differences between species — largely as a result of the behaviourist domination of animal learning paradigms. This approach is contrasted with evidence from ethology and neuroscience which provides much richer characterisations of biological systems.

Next, a novel characterisation of complex biological systems is described. This characterisation, informed by dynamic and situated perspectives, has been developed by McGonigle and Chalmers (1984,1996,1998 for example), and depicts the growth of epistemic competences in complex biological systems. Central to this characterisation is the notion of the self-organisation of competences by agents on open-ended growth trajectories toward maximally economic, generative, information-handling strategies in order to deal with ever-larger search spaces. Such internal control requires *interpretation* and thus *contextual* representation, in order both to learn from error and for arbitration between competing behaviours.

Engineering It is suggested that the novel characterisation described reinforces earlier criticisms of symbol-based and behaviour-based approaches and suggests a number of important *design principles* for the construction of artificial systems. This characterisation has inspired a number of robotic implementations within our research group. These implementations, precursors of the novel architecture presented in chapter 4 are briefly described. A number of key architectural features are incorporated such as state-based (re-)interpretation of signals, and the development of a robotic task grammar.

Of the design features suggested by McGonigle and Chalmers' detailed characterisation of complex biological systems a number have already been incorporated within previous robotic architectures within our research group. The next chapter describes a new architecture designed to incorporate, and extend, design features such as state-based interpretation and the recombination of behaviours from a core repertoire, whilst providing for new architectural features such as error recovery and, critically, self-organisation in accordance with an inbuilt economy metric.

Chapter four

An architecture is presented which is inspired by the biological characterisation, and constitutes a progression from the robotic implementations, described in chapter 3. It makes no claim to embody *all* of the design features which emerge from the analysis of chapter 3 but has been constructed such that the remainder can be implemented later without requiring radical redesign of the system.

The architecture controls a mobile robot and is thus embodied and situated in a real-world environment. The system is autonomous — not merely trivially with respect to a power cable, but imbued with the potential[2] for the controller to select behaviours based on task demands. The system is inherently hierarchical, with behaviours constructed to allow serial control, a behavioural syntax, and error cognisance. The system has memory — providing the potential for progressive adaptation over time. The central feature of the architecture is the controller which maintains all global state information, implements behaviours in the task queue and logs error. Its role is easily extendible.

A pragmatic view of representation is adopted — currently state serves primarily to tag whereabouts in task-space, behaviour name and status, previous behaviours *etc.* A global world model is not maintained by the system and has not been found necessary for the competences engineered so far.

[2] Albeit at a prototype stage.

Development of the architecture is possible both horizontally (addition and modification of behaviours) and vertically (progressive addition of increasingly abstracted control and learning mechanisms) as advocated by McGonigle (1991). Although not at present inherently dynamic the system is capable of supporting self-organised competences with the further possibility of epistemic derivations.

Chapter five

A description is provided of the application of the architecture to navigational competences. Navigation was chosen both because a locative competence is absolutely critical to any locomoting system and has been argued to be ontologically and ontogenetically prior to any other (McGonigle, 1991) and because there exists a history of navigational implementations from within our research group on which to build.

Specifically two approaches were examined:

Navigation without learning A simple and reliable navigational algorithm (utilising on-board odometric data) was developed which, in conjunction with separately engineered *taxes* allowed the robot to dead-reckon 'home' from distant parts of the niche. Preliminary investigations indicated that robust navigation could be achieved in the absence of planning and that the architecture successfully supported error cognisance and recovery.

Learning to navigate Two approaches to route learning through the niche are presented. The successful strategy was based on trial-and-error and subsequent self-selection of privileged vectors in accordance with an inbuilt economy metric. This method provided the opportunity both for progressive behavioural adaptation, and critical epistemological derivations.

Related navigational implementations are briefly evaluated with respect to the synthetic programme adopted herein specifically with regard to economy of both design and computational resource together with their potential for scaling-up, focusing on those which are pre-installed with, or acquire, map-based representations, both geometric and topological. Although successful in the navigational domain it is suggested that such implementations require a greater computational resource, and provide fewer opportunities for extendibility than the implementation described herein. It is argued that this implementation provides evidence that a synthetic analysis of intelligent systems results in novel approaches to the construction of artificial intelligent systems possessing an inherent extendibility. The self-organised navigational competence described leads, through situated activity, to representations of the environment, based on system behaviour which are, therefore, both economical and inherently meaningful to the agent.

Chapter six

The argument of the thesis is summarised, and the synthetic stance advocated herein restated. This stance is distinguished from traditional reverse engineering which tends to focus on the locomotor behaviours of simple systems. Two future research directions, error interpretation and self-organised action selection, are discussed, and potential application of the current architecture to each described, before final conclusions are drawn.

Chapter 1

Two dogmas: "symbol-based" and "behaviour-based" stances

For most of the past century characterisations of intelligent systems, within both the psychological sciences and artificial intelligence, have tended to fall in either of two camps (see table 1). One, the classical 'symbol-based' characterisation of systems, primarily adopts a top-down perspective, focusing on complex competences such as reasoning, problem-solving, and memory in predominantly static agents. The other, 'behaviour-based' characterisation, is primarily bottom-up, concentrating on the reactive behaviours of simple systems as a consequence of a quantitative view of differences in system complexity.

	Psychology	Artificial intelligence
Symbol-based	cognitive psychology (1.1.1)	classical robotics (1.1.2)
Behaviour-based	behaviourism (1.3.1)	behaviour-based robotics (1.3.2)

Table 1.1: Conflicting stances within psychology and artificial intelligence.

'Symbol-based' and 'behaviour-based' characterisations tend to be viewed as mutually exclusive stances. Indeed, the 'symbol-based' stance in psychology developed in large part as a reaction to its 'behaviour-based' predecessor. Artificial intelligence has witnessed the converse transition. Paradoxically however, as we shall see over the course of this chapter, the kinds of artificial intelligent systems inspired by these 'competing' characterisations are formally isomorphic in a number of limiting respects.

1.1 Symbol-based stances

Classical symbolic representational stances initially appear promising, focusing, as they do, on complex competences as the critical indices of systems. On closer inspection, however, it emerges that the end-state characterisation lacks psychological plausibility, and that fundamental developmental questions remain unanswered. Furthermore, the artificial systems inspired by these stances are brittle, often hand-crafted and are semantically blind — operating in highly constrained, preinterpreted domains of competence (Luger & Stubblefield, 1998*b*).

1.1.1 Cognitivism

The cognitivist stance in psychology began to cohere in the late 1950s and early 1960s largely as a consequence of the failure of behaviourism to account for phenomena such as latent, observational, and linguistic forms of learning (Chomsky, 1957), and was especially motivated by dissatisfaction with the behaviourist concentration on observable action rather than putative internal knowledge structures.

Informed by theories which characterised information in terms of internal processes of ambiguity reduction (Shannon & Weaver, 1949), and research conducted at Cambridge by Bartlett (1932, 1958) and Craik (1943, 1966) which viewed agents as *active* participants possessing unique experiential histories and thought as a collection of mechanisms similar in principle to those of artifacts, together with the discovery of 'chunking' phenomena in memory research (Miller, 1956), cognitivism re-legitimised the study of internal psychological processes.

Primitives

The mediating internal processes on which cognitivism concentrated were those regarded as the distinguishing feature of human beings, the epitome of complex systems — thought and language. The Cartesian view of the abstracted mind led to a characterisation of internal, mental processes as arbitrary and amodal (Barsalou, 1993, 1999). This characterisation immediately suggests a linguistic analogy and, indeed, the most widely accepted view of internal representation became the concept of the 'language of thought' (Fodor, 1975).

The mind was characterised, in analogy to human language, as a symbol system — a set of arbitrary tokens that can be manipulated on the basis of explicit rules which are themselves tokens, or strings of tokens. The rules by which tokens may be manipulated concern only the shapes of the tokens rather than their meanings, they are purely *syntactic*. The relevant explanatory aspect of mind became its status as a symbol system, and the internal processes of interest the logical manipulation of such internal symbols. The rules by which internal representations are manipulated were identified with those of the 'supreme' human achievement, deductive logic. The behaviours of interest to cognitive science became, therefore, those which depend on the construction and manipulation of such internal representations — language production, memory and reasoning skills. This focus on loosely-coupled, 'deliberative', abstracted cognitive behaviour, described by and attributed to logical symbol manipulation, studied predominantly in unnaturalistic laboratory settings merely served to reinforce the assumption of independence of body and mind.

The language of thought hypothesis is generally construed as an empirical theory about the nature of cognition and thus stands or falls on its ability to account for cognitive adaptation[1]. We shall see how well it fares in the task.

An abstraction mistake

Much of traditional cognitive science assumed that the mind can be safely abstracted away from its bodily instantiation and studied independently. Based on the neo-Cartesian view of the existence of the 'purely mental' (Strawson, 1994) the details of the physical embodiment of cognitive agents were ignored with respect to the explanation of cognitive phenomena (Wheeler, 1995). Theorising was purely at the algorithmic level (Putnam, 1975) and thought assumed to be analogous to language, consisting of the sequential processing of amodal symbols 'stored' in list or sentence-like internal structures (Fodor & Pylyshyn, 1988). Critically, the interrelationship of cognitive, perceptual, and bodily processes, which had formed an essential component of theories of mind from Aristotle (4th Century BC/1961) and Epicurus (4th Century BC/1994)

[1] Although some have attempted to argue for such a view on *a priori* grounds (Davies, 1989, 1991; Lycan, 1993; Rey, 1995).

onward, was abandoned with the adoption of this *disembodied* and *unsituated* view of cognition. We will see that this abstraction mistake lies at the root of the many problems which have befallen cognitivism.

Amodality

Although cognitivism has played an important role in re-legitimising internal psychological mechanisms in explanations of behaviour, what evidence is there for the type of internal processes postulated? Currently little direct empirical evidence has been reported for the existence of amodal symbols (Shannon, 1987). The correct interpretation of Moyer's finding of the symbolic distance effect (1973) remains unresolved and has prompted an ongoing debate about the nature of representation (Potts, Banks, Kosslyn, Moyer, Riley, & Smith, 1978, for a review). Debate ranges from those who adopt a strongly amodal symbolic position (Banks for example) to those who postulate the existence of both symbolic and analogue codes (Paivio, 1986). Recent reviews (Glaser, 1992; Seifert, 1997, for example) suggest that internal representations appear to have a perceptual (and therefore a situated and contextual) character.

Neuroanatomical evidence suggests that categorical knowledge is grounded in sensorimotor areas of the brain — damage to certain of these areas leads to disrupted conceptual processing of relevant categories (Damasio, 1989; Gainotti *et al.*, 1995). Furthermore, the representation of spatio-temporal forms of knowledge using purely amodal symbolic structures has been argued to be unlikely as the computational systems which would result are (expected to be) both brittle and, in some respects, intractable (Winograd & Flores, 1986; McDermott, 1987; Glasgow, 1993; Clark, 1997). Finally, when the influence of context on behaviour is universal in the remainder of the biological realm (Hinde, 1982), and a fundamental quality of human language (Wittgenstein, 1953; de Saussure, 1974) — the very phenomenon which inspired cognitivism — what would the loss of such contextual information buy a system?

Logic or rationality?

How accurate is the symbolic "rule-based" characterisation of complex systems? Although early formalisations of logic (Boole, 1854/1951, for example) were also regarded as the principles governing thought processes, the late twentieth century has rediscovered the [2] distinction between normative and descriptive views of logic (Oaksford & Chater, 1998*a*).

Over the past three decades a consensus has developed that human cognition obeys *rational* rather than strictly (deductive) logical principles (Wason, 1966; Wason & Johnson-Laird, 1972; Kahneman & Tversky, 1973; Kahneman *et al.*, 1982; McGonigle & Chalmers, 1977, 1994; Oaksford & Chater, 1998*b*). Performance on a number of laboratory tasks, construed as indices of logical ability, often falls far short of that expected of a logical agent (Wason & Johnson-Laird, 1972; Kahneman & Tversky, 1973; Kahneman *et al.*, 1982, for examples). Significantly, altering stimuli to those more often encountered in the course of everyday life results in improved performance (Wason, 1977). The Piagetian (1953, 1955) 'transitive inference' task is often regarded as a good index of the understanding of necessity through the logical coordination of internal symbolic representations (Bryant & Trabasso, 1971). However, effects of presentation mode (McGonigle & Chalmers, 1984) together with the existence of high levels of transitive choice in four year old children (Lawrenson & Bryant, 1972) and in non-human primates (McGonigle & Chalmers, 1977) cast doubt on this interpretation. Indeed the production system model of Harris & McGonigle (1994) is capable of generating 'transitive' choice in the absence of logical coordination mechanisms and, furthermore, accurately characterises performance on untrained triadic choices.

In addition to these 'transitive' abilities, hierarchical competences indexed by class inclusion (Inhelder & Piaget, 1964) and quantifier use (Johnson-Laird, 1983) are generally regarded as an outcome of logico-linguistic ability. However evidence suggests that this too is based on the internal generation of economic information management principles (McGonigle & Chalmers, 1998*a*) which presuppose no logical abilities on the part of the agent.

A consensus has developed which suggests that intelligent systems are more accurately characterised as *rational* rather than logical. Indeed systems are rule-based, or at least describable by such, yet these rules are not those of deductive logic but reflect computationally economical information-handling procedures (McGonigle & Chalmers,

[2] "Why is it that I am not compelled to do that which is entailed?" Aristotle (4th Century BC/1990) 'Essay on rational action'.

1998*a*).

Origins and growth

The origins of intelligence

As previously mentioned although one benefit of the cognitive stance is the recognition of the complexity of systems from the outset, the focus is on the *preadapted* agent — little attempt is made to connect the logico-mathematical end-state characterisation with phylogenetic development, and the progressive, life-long, adaptation of the epistemic agent over ontogeny goes largely ignored.

The most comprehensive account of the phylogenesis and ontogenesis of cognition from a classical symbolic perspective is Piaget's (1970, 1971) genetic epistemology. Although admirably aspiring to be an account of the interrelationship of 'Biology and knowledge' (1970) characterising knowledge gain as a process of regulation within the organism-environment system and thus grounding symbolic competences in earlier organic autoregulation, it is in reality adult-centred with the aim of reconstruction of the developmental mechanisms from the endpoint (Inhelder, 1990). Despite adopting an early systems stance, genetic epistemology under specifies the critical role of lower ontological bounds (Le Gare, 1987; de Lorenzana & Ward, 1987), and the mechanisms by which they constrain epigenesis.

More recently, grounded developmental accounts (Karmiloff-Smith, 1992; Halford, 1993, for example) have attempted to account for explicit knowledge through processes of reorganisation and abstraction. Karmiloff-Smith (1992) hypothesises that human cognitive development involves the complementary processes of progressive modularisation and explicitation of knowledge. 'Representational redescription' is posited to be a spontaneous reiterative process whereby domain specific (implicit) knowledge, is made progressively explicit *via* redescriptive processes both within and across domains, culminating in a publicly transmittable symbol system (Karmiloff-Smith, 1992; Clark & Karmiloff-Smith, 1993). However, the theoretical relationship between language and abstracted explicit knowledge remains unresolved: the mind's capacity to enrich "itself from within by re-representing the knowledge that is has already represented" has been argued by Karmiloff-Smith (1983, 1992) to be rarely present in other species except, possibly, the chimpanzee although here the resulting higher-level codes are suggested to be relatively impoverished.

Halford (1993) suggests that the observations which inspired developmental stage theory (Piaget, 1950) might instead reflect progressive representational differentiation over ontogenesis. Performance on Piagetian tasks within stages is interpreted as a result of an increased capacity for parallel processing of more complex relations. Here cognitive development is suggested to depend on the interaction of processing capacity and acquisition processes (Halford, 1971; Halford *et al.*, 1995, for example) such that acquired knowledge depends dually on experience and current system processing capacity. This growth of processing capacity is regarded as an important explanatory principle for cognitive development. Halford *et al.* (1998) suggest that non-human primate is capable of processing explicit representations based on the research of Premack (1983) and Holyoak & Thagard (1995), yet the relationship between pre-linguistic and linguistic relations is unclear, as is the reliance of the progressive development of relational processing on linguistic experience.

Although both of these examples strive to explain the origins of 'symbolic' competences in pre-linguistic abilities it is becoming increasingly clear that the search for the origins of explicit, transmittable knowledge in symbol-using human agents is greatly complicated by the fundamental interconnectedness of language and thought. What is required is an empirically motivated research programme which attempts to trace the origin of knowledge, and knowledge acquisition devices, over the life-history of individual non-linguistic agents (Chalmers & McGonigle, 1997; McGonigle & Chalmers, 1998*a*).

The origin of meaning

What is it to mean? Conventionally symbols are supposed to 'remind us' of their referents, not through resemblance, nor known causal relationships, but rather through arbitrary convention (Plato, 1926; Wittgenstein, 1953). By what process(es) are these relationships established? We have already seen that, for cognitivism, meaning cannot be grounded in bodily experience ruling out the logical positivist (Ayer, 1947, 1971;

Russell, 1956/88; Mach, 1959, for examples) and naturalistic (Bergson, 1944; Popper, 1985) explanations. The alternative, however, is a futile quest to trace meaning back through individual exposure until the first usage of a symbol is reached (see Kripke, 1980, for example).

So although one of the strongest assumptions for cognitive accounts is that the possession of a publicly transmittable symbol system reflecting an abstracted, internal, language-like representational system indicates a new level of system development, the origin of meaning for such symbolic systems remains unexplained. Whereas learning stances view symbols as representative of pre-encoded 'objects' (Anderson, 1990), cognitive stances do not disambiguate symbols from their semantic content. Although the meaning of such symbols can be made explicit to the agent grounding such meaning within disembodied symbol systems might be impossible[3].

The most famous critique, Searle's (1980) 'Chinese room' analogy, questioned the ability, in principle, of a symbol system to pass the Turing (1950) test and concluded that it was possible in the absence of real understanding. This led him to suggest (Searle, 1980*a*) that symbol systems lack "intrinsic meaning", or intentionality, in that a symbol system *alone* has no understanding of the symbols it manipulates. For Searle the ability of the inhabitant of the Chinese room to pass the Turing test is not sufficient, he argues that understanding can arise only in conjunction with consciousness and that consciousness arises only within systems which possess the appropriate physical structure: an evolved brain. According to Searle, the problem can be resolved by examination of the physical properties of brain which might reveal which "causal powers' might underlie intentional mental states. More recently Searle (1992) suggests that the human brain is not the only structure which can support such mental states although the appropriate causal property of brains remains unclear, as does the phylogenetic division between conscious and non-conscious brains.

More recently Harnad (1990*a*,*b*), restating Kripke's dilemma (1982), dubbed this the 'symbol grounding problem' and pointed out that the search for meaning in symbol systems results in an infinite regress through layers of meaningless syntactically related symbols. Thus we see that, once more, abstracting the mind away from the body results

[3] Unless, of course, one accepts Fodor's (1975) original claim that all concepts are innate!

in an insurmountable difficulty — how to reconnect them again!

Summary

The emphasis of cognitivism on linguistic tasks reinforces the assumption that the core complex system competence is the logical manipulation of symbolic internal representations. However, over the past two decades the cognitivist characterisation of complex systems has come under increasing attack. It increasingly appears that the end-state characterisation of purely symbolic level intelligence inadequately accounts for a range of evidence which suggests instead that human cognition is more accurately characterised as *rational* rather than logical. Furthermore, the assumption that cognition can be abstracted away from its biological instantiation has led to the *faux* problem of symbol grounding — how to *re*connect mind and body. Finally, cognitivism is unable to explain how symbol-level competence might be rooted in pre-linguistic primitives. The mischaracterisation of a logico-mathematical developmental end-point would render the attempt to trace its origin impossible, fortunately however, such attempts are rarely motivated by the formalist stance which is more concerned with the characterisation of preadapted agents than with their phylogenetic and ontogenetic origins.

The cognitive account, therefore, fails in two fundamental areas: its characterisation of intelligent systems; and their phylogenetic origins, and ontogenetic growth. The conflation of external symbolic representations with neurological and psychological structures and processes (Vera & Simon, 1993) does seem to be a category error. Such considerations do not bode well for classical AI which also shares this assumption.

1.1.2 Classical AI

The discipline of Artificial Intelligence (AI) was born in the late 1950s and early 1960s and grew out of the earlier fields of philosophical logic, computing, and cybernetics. Whitehead & Russell's (1925) proof that mathematics is substantially derivable from the predicate calculus grounded mathematics in logic resulting in the primacy of the deductive method in the field of computability. Electronic computers were first built during the 1940s as part of the Second World War effort of the United Kingdom and the United States, and were the result of previous mathematical work by Alan Turing and John von Neumann, both of whom were concerned with computability and the theory of games — critical early research areas for the newly-formed discipline of AI. Turing, particularly, published papers on the prospects of making computers intelligent (Turing, 1948/1970), and asked whether computers could think (Turing, 1950).

The symbolic representational stance within AI, derived jointly from the origins of computing in philosophical logic with consequent envy of the power of logico-mathematical symbol systems (Whitehead & Russell, 1925; Gödel, 1990; Turing, 1950), and the prevalence of cognitivism at the time of its origin, embodies a rationalist desire to apply mathematical standards of proof and analysis to the construction of intelligent agents (Luger & Stubblefield, 1998*b*). Just as cognitivism identified the rules of thought with those of deductive logic, for classical AI the primary system competence is characterised as the logical manipulation of internal, symbolic representations. An inevitable consequence of this stance was that the behaviours of interest for classical AI became, as with cognitivism, those thought to depend on reasoning and representation — problem solving and planning.

Although the primary focus of classical AI was on reasoning and problem solving implementations (Newell *et al.*, 1957, 1959, for examples) and thus the development of efficient search algorithms, given the context of this thesis the focus of this review will be classical robotics. The accepted goals of artificial intelligence are the construction of intelligent systems, and the understanding of biological intelligence (Winston, 1984; Pfeifer, 1996, for examples); what progress has been made?

Primitives

The physical symbol system (hereafter PSS) hypothesis (Newell & Simon, 1976) underlies, more or less explicitly, classical robotics. The underlying assumption is that the necessary and sufficient condition for any system to exhibit intelligent behaviour is that it is instantiated as a physical symbol system. *Only* such systems (and *any* which are appropriately organised) are capable of intelligent behaviour, where intelligent behaviour means that appropriate to the system's ends and adaptive to environmental demands, the scope of which is seen as being approximately that of human behaviour.

The PSS hypothesis commits its advocates to a number of methodological principles:

- The world should be described, as far as possible, exhaustively.
- This description should be symbolically based.
- Heuristic search should be used to explore the potential inference space.
- As all computational systems are functionally equivalent, the cognitive architecture of a system should be independent of its physical structure.

So the computational agent became characterised as a disembodied information processor, leading to the input-output model of intelligence which was to dominate AI for the first three decades of its existence. Classical robotic architectures represented the world by internal symbol systems, usually based on the predicate calculus. All aspects of the world tended to be represented, state was unpartitioned, and *a*contextual.

A classical robot

Possibly the clearest exemplar of classical robotics is the use of the Stanford Research Institute Planning System (STRIPS) and the control system PLANEX, which were used to drive Shakey, an experimental mobile robotic platform (Fikes & Nilsson, 1971). Shakey possessed on-board visual, tactile and acoustic sensors with an off-board computer used for signal processing, pattern recognition, model building and plan generation. Initially possessing two representation systems (map-like and predicate), a predicate-based system was finally chosen, incorporating predicates concerning five classes: the robot itself, objects, rooms, doors, and walls. Predicates expressed information concerning the name and coordinate location of an entity, as well as more complex relational statements such as containment of a particular object by a particular room (Nilsson, 1984).

This internal model was used to generate action sequences using STRIPS, to drive actions themselves, and to coordinate the control system (Fikes & Nilsson, 1971; Fikes *et al.*, 1972; Nilsson, 1984). PLANEX was used to execute generalised versions of a plan, having equivalent actions to all operations expressible in STRIPS. The Shakey project was designed to investigate the problems of efficiently representing the world, and implementing a planner based on heuristic search through a symbolic search space. Explicit goals included the solution of incompletely specified problems which would require the construction of intermediate goals and strategies, and to improve performance over a training period (Nilsson, 1984).

Both Shakey[4], Mark I and II exemplify classical robotics:

- The system is implemented in an inherently undynamic[5] unchanging Platonic domain.
- A global world model is constructed, maintained, and manipulated. A plan is constructed off-line, on the basis of this model, before being implemented. A sense–think–act, or sense–model–plan–act (Brooks, 1991*b*), control cycle is adopted, analogous to the TOTE units of cognitivism (Miller *et al.*, 1960).
- The domain is preinterpreted — the meaning of representations are installed within the system reflecting the ontological stance of the designer.

A second abstraction mistake — the disembodied robot

As the symbol level was regarded as a natural functional level of its own, possessing invariances independent of specific physical instantiations, intelligent behaviour was thought to arise from any PSS regardless of its particular physical implementation. Traditional AI systems are therefore only trivially embodied and situated. Arguably, applying cognitivism to robotics provided an opportunity to reintroduce embodiment and situated behaviour as fundamental constraints on the structure and processes of intelligent systems, yet the PSS hypothesis explicitly excludes this possibility thus rendering classical robotics susceptible to the fatal criticisms applied to cognitivism.

Although the aim of classical robotics was to construct artificial intelligent systems the single-minded devotion to a symbolic perspective on intelligent systems uninformed by adequate biological characterisation was to prove its undoing. Ethological and comparative research had established that the sensorimotor capacities of biological

[4] See also Hillaire (Giralt *et al.*, 1984) and CART (Moravec, 1982).

[5] Notwithstanding the Cartesian 'gremlin' which inhabited Shakey's environment.

systems critically constrain information processing[6] — as had been argued by Turing as early as 1948. Furthermore, the behaviour of biological systems was observed to be adaptive within the system's niche — behaviours contribute to evolutionary fitness thereby providing a fundamental criterion by which their success can be judged (Hinde, 1966; Bateson & Klopfer, 1982, for example). Paradoxically, by abstracting away intelligence and thus ignoring the fundamentally embodied and situated nature of cognition, classical robotics achieved the construction of 'disembodied robots'. As with cognitivism, the abstraction of mind away from body was to lie at the root of the many problems which afflicted this programme.

Modelling the world

Exhaustive symbolic representation of the sensed world rapidly led to significant difficulties — sensor inaccuracy and *environmental dynamism* meant that the formation and maintenance of a robust internal world model from logically manipulated symbols was not practicable. The environment is only partially knowable for all systems so constructing an exhaustive world model is doomed to failure from the outset in the real world. Furthermore, even if it were possible to construct such a model, maintaining it would require an unfeasibly large computational resource. Shakey, and other implementations such as the copy-demo (Winston, 1972), shared the assumption that an accurate internal world model could be constructed, maintained, and manipulated. The impossibility, in principle, of exhaustive knowledge of non-trivial domains, led to the adoption of highly constrained, "toy-world", environments.

The failure of classical AI to accurately model the world led to many criticisms of the necessity and desirability of such models (see Winograd & Flores 1986; Suchman 1987, for example) and a reaction against exhaustive, symbolic representational schemes that is the origin of all other robotic paradigms now in existence.

Platonic domains

An obvious feature of human systems is their engineering of the world around them in order to facilitate its comprehension. To what extent, however, should such engineering be employed when constructing artificial systems? Classical robotics, it is

[6] For example von Uexküll's (1921) concepts of *umwelt* and *merkwelt*.

generally agreed, engineered the environment to suit the system to the extent that the fundamental purposes of AI (Winston, 1984; Pfeifer, 1997) were betrayed: first, the environments in which systems operated were so divorced from the real world that it became extremely unlikely that such systems would be of much general use; second, the insights to be gained with respect to complex biological systems from this approach were minimal.

Winston's (1972) copy-demo, Robert's (1963) vision program (the forerunner of all modern vision programs) and Shakey, all relied on sparsely inhabited environments with controlled lighting and background — utterly unrealistic scenarios. All of these implementations, operating in toy-worlds with (near) perfect sensors, served to reinforce the assumption that an accurate symbolic world model could be constructed and used to control behaviour. In the real, dynamic world this assumption simply does not hold. One assumes that an implicit driving assumption behind the approach was that dynamism could be added later without seriously affecting performance but this was not to be.

Extreme *brittleness* became the outcome of this environmental simplification. Classical implementations were fault and noise intolerant (Pfeifer, 1996)[7] and incapable of generalising their knowledge to new domains.

Origins and growth

Development

The ability to learn must constitute part of any artificial system with aspirations towards intelligence. Learning within classical AI centres mainly on inductive generalisation within, and sometimes across, knowledge domains with 'inductive biases' providing constraints on possible generalisations.

Symbol-based machine learning is divided between supervised and unsupervised techniques. Both supervised (for example data-driven similarity based algorithms; explanation-based and analogical learning which rely on prior domain knowledge in addition to

[7] CART (Moravec, 1982), for example, existed in a static environment with the primary external source of change being angular displacement of the sun — yet even this minimal dynamism resulted in system failure Brooks (1991*c*).

training data) and unsupervised (which tend to centre around category formation and conceptual clustering) learning techniques presume some preinterpretation of domain by the designer. Notwithstanding differences between the amount of domain preinterpretation, prior knowledge, and training data required for these different techniques all model the learning process as state-space search (Luger & Stubblefield, 1998*b*, p. 606) and involve modification or construction of explicit symbolic representations.

Although symbol-based machine learning has been shown to be successful within circumscribed problem domains a number of issues remain unresolved. Generalisation requires rich training data and extensive preinstalled domain knowledge in addition to constraining inductive biases — learning algorithms cannot merely be supplied with examples and construct appropriate and effective generalisations. However, the relationships pertaining between these factors remain unresolved and reflect the often implicit assumptions of the designer. Unlike human cognisers, machine learning implementations are effective only within a preinterpreted world (Luger & Stubblefield, 1998*b*) and cannot yet reinterpret data or learned rules, nor move between different interpretative perspectives as the situation demands. This 'inductive bias' (Luger & Stubblefield, 1998*b*, p. 772) of symbolic learning algorithms has motivated connectionist and emergent learning approaches which conversely suffer from a lack of inductive constraint. All current machine learning approaches, whether symbolic, connectionist, or emergent, lack a fundamental feature of human learning: semantic and pragmatic constraint (Luger, 1994, p. 450) — they are predominantly neither knowledge driven nor goal oriented.

Finally, biological systems possess a range of inductive mechanisms subserved by functionally distinct neuronal circuitry (Gazzaniga *et al.*, 1998) and optimised for particular internal and external contexts (Hinde & Stevenson-Hinde, 1973; Gallistel, 1995). Classical learning techniques, in contrast, are relatively *ad hoc*, *a*contextual and imposed on top of preinterpreted, pre-installed representational schemes whose origins remain unexplained.

The origin of meaning

We have seen that the origin of meaning remains unresolved in cognitive accounts of complex systems. How did classical AI overcome this problem? Initially the system must build a model from sensed data. We have seen that this was possible only after the adoption of highly constrained environments (Moravec, 1982; Nilsson, 1984).

For problem-solving AI the problem of grounding meaning did not arise. Symbol systems are systematically *semantically interpretable* (Fodor & Pylyshyn, 1988; Harnad, 1992) so symbols can be defined syntactically in terms of other symbols and *how they are processed by some interpreter* (Quillian, 1968; Newell & Simon, 1976). For non-robotic AI, expert systems for example, the interpreter is a human operator so the mapping of sign and signified is grounded in human experience[8]. For this reason the relation of symbols to the outside world was rarely discussed explicitly (Harnad, 1993).

For classical robotics however, with the goal of the construction of artificial, *autonomous*, intelligent systems how could the meaning of symbols be made explicit to the system, or be grounded in the external world? The solution to the problem was *pre*interpretation of the world. Meaning entered the system through the translation of the predefined semantics of a given logic-based symbolic representation (Winograd & Flores, 1986, p. 18). We see, once again, that the abstraction of 'mind' from 'body' lies at the root of the difficulty.

Meaning and context

The reliance of classical systems on their designer's ontological stance lies at the root of the 'frame of reference' problem (McCarthy, 1963; McCarthy & Hayes, 1969). Although multiple interpretations of this problem have now been discussed (van Brakel 1992, 1993, discusses the 'family' of frame problems; and see also Pylyshyn, 1987) the central dilemma is the problem of handling change (Janlert, 1987). Given that the system ideally possesses an exhaustive symbolic model of the world, the number of inferences that can be drawn regarding which world states have been affected by a given action[9] results in a massive combinatorial problem.

---

[8] Attempts to remove the human operator from the loop such as the encyclopedia project (Lenat & Feigenbaum, 1991), for example, required entering hand-crafted knowledge units into the system with the vague hope that the system would be able to learn for itself by reading at some ill-defined point in the future. Smith (1991) points out that the early part of this project was devoted to finding more primitive levels of knowledge that would 'ground' the higher ones.

[9] As Harnad (1993) correctly asserts, the problem is inherent to *any* new datum rather than action *per se*.

The problem raises the importance of knowledge of the side-effects of action and, critically, the issue of the *relevancy* of knowledge. How is it possible, for a system whose ontology is preinstalled, to determine what is, and is not, relevant? Various solutions have been proposed ranging from McCarthy & Hayes' (1969) frames which link related propositions together in an attempt to reduce the search space whilst assuming that all other frames remain unaffected, to the anti-symbolism of Brooks (1991*c*). Frame relations are rules which specify which predicates remain unchanged by rule applications. However, this approach requires new frame rules to be devised whenever a new predicate is introduced with a resulting exponential increase in search time (Luger & Stubblefield, 1998*b*, p. 191).

An alternative solution, the triangle tables used to store and organise macro operations in STRIPS[10] alleviated this problem somewhat by sacrificing exhaustiveness — only postconditions of actions that are preconditions of consequent actions were retained by the system. However, this approach also suffered from the combinatorial problem — as the number of operators increases so does the degree of pattern matching required to determine whether an operation can be applied (Luger & Stubblefield, 1998*b*).

For classical systems, then, the frame problem is a consequence of the impossibility of exhaustively representing the world in non-trivial domains (Luger & Stubblefield, 1998*b*, p. 333). To compound the problem, classical representational systems are *acontextual* and therefore provide little constraint on the assessment of the relevancy of information, behaviour, or anything else. More recently circumscription (McCarthy, 1980; Lifschitz, 1984; McCarthy, 1986) has been suggested as a means of effectively delimiting problem descriptions and thereby setting an interpretative context, but classical robotic implementations could neither move between interpretative contexts as the occasion demanded nor, indeed, accurately assess that the occasion had even changed.

[10] Triangle tables (Fikes *et al.*, 1972; Nilsson, 1980) were designed as data structures for organising action sequences, including those featuring potentially conflicting subgoals. The preconditions of one action are related to the postconditions of all its preceding actions. Triangle tables were used to determine when an operator could be incorporated within a plan.

Summary

Classical AI suffers from two category errors, both of which result in insurmountable problems, just as for cognitive accounts: that the 'mind' can be abstracted away from the 'body'; and that the complexity of behaviour implies equivalent complexity of internal control structures (Simon, 1962) — in this case internal, logically manipulated symbolic representations.

The systems that result are brittle and unreactive. The assumption that an exhaustive symbolic model of the world is a necessary feature of complex systems results in the adoption of toy-worlds and *pre-interpreted* domains of competence. Even within these unnaturalistic domains the systems are hand-crafted. Little progressive adaption of systems occurs, learning is limited to refinement of hypotheses or small-scale expansion into novel domains. Finally, the ontology of the system is that of the designer; such systems are *semantically blind*.

1.2 Interim summary

Over the past two decades 'symbol-based' characterisations of intelligent systems have been subject to increasing attack both within psychology where evidence accumulated over the past three decades suggests that systems are more accurately characterised as 'rational' rather than logical, and within classical AI which has failed to synthesise systems capable of operating robustly in non-trivial environments. Instead such systems are restricted to highly-constrained, preinterpreted domains of competence.

Although the symbolic representational stances appreciate system complexity, they ultimately fail to deliver a comprehensive account of intelligent systems. The characterisation falls on two counts: (1) *Embodiment* as a critical constraint on system competences and future development, and the fundamentally *situated* nature of systems are ignored, as are the *dynamics* of system behaviour. Furthermore, the classical preoccupation with representations *qua* logically-manipulable symbols fails to take account of the adaptive function of representation in contextually marking action, perception and their (re)interpretation which by no means necessitates all representations to be symbolic. Rather it seems much more likely that cognitive systems make use of many forms of representation both internal and external[11], symbolic, distributed, and hormonal; and (2) classical symbolic stances fail to deal adequately with development — the mechanisms by which complex system competences might be rooted in pre-linguistic behavioural system primitives are unaddressed and the stances therefore fail to answer the fundamental question — what are the origins of intelligent systems?

The symbolic stance, then, provides neither a plausible account of the development of complex intelligent systems from their more simple precursors over phylo- and ontogenesis, nor an adequate characterisation of such systems at later developmental stages, nor does it support the construction of robust artificial intelligent systems.

1.3 Behaviour-based stances

The major alternative characterisation of systems, within both psychology and artificial intelligence, regards (predominantly reactive) behaviour as the indexical system competence. These stances developed as a result of dissatisfaction with introspective, or symbolic characterisations and, therefore, eschew the explanatory role of internal representational mechanisms. Complex system competences tend not to be examined which means that a plausible account of the development of complex systems from their more simple precursors must be provided. The behaviour-based stances posit *associationistic* behaviour modification and acquisition mechanisms as the processes underlying such development.

We shall see that the rejection of the explanatory role of internal mechanisms leads to problems with the characterisation of the behaviour of complex systems and, indeed, is in some important respects inadequate for descriptions of simple systems. Furthermore, behaviour-based stances fail to provide a plausible account of the progressive development of biological systems and are incapable of supporting such development in artificial systems. In both cases the 'scaling up' problem remains unresolved.

[11] Much has been made of the use of the environment to externalise computation recently with situated (Suchman, 1987, for example) and extended-mind (Clark & Chalmers, 1998) theories of cognition and action gaining prominence.

1.3.1 Behaviourism

> [Psychology] "as the behaviourist views it, is a purely objective branch of natural science. Its theoretical goal is the control and prediction of behaviour, [it] recognises no dividing line between man and brute." (Watson, 1913, abstract)

Behaviourism dominated anglo-american psychology from the 1920s to the mid-1950s. At its inception it was an attempt to objectivise psychology as a science by discarding the traditional subjective, introspective modes of investigation: rather than consciously acting, agents react. From the outset behaviourist research assumed a neo-positivist stance — as mental (cognitive) processes are unobservable, *scientific* investigation must be confined to the external manifestations of cognition — behaviours (Watson, 1913).

Primitives

Behaviourism, unsurprisingly, began with the assumption that the appropriate level of investigation for all systems was observable, and measurable, behaviour (Watson, 1913, 1925, 1929; Skinner, 1953, for examples). Since 'no dividing line' was drawn between man and brute, and the generalised laws of learning applicable to all species were the goal of the project, simple reflexive systems, or reflexive subsystems of more complex (epistemic) systems were the experimental targets.

All learning and behaviour were thought to be characterisable in terms of stimulus-response or operant associations. The relationship between stimulus and response over short spatio-temporal intervals and how such relationships change with varying forms of local environmental contingencies were investigated. The larger the number of these pairings, and the greater the capacity for their acquisition, the more complex the organism. System differences were therefore characterised *quantitatively*.

One benefit of the behaviourist focus on primarily reflexive behaviours in simple systems was that the tight couplings between inputs and outputs made such systems transparent to traditional forms of reductionist analysis. Moreover, theories could be constructed about the behaviour of simple systems (or subsystems) which were themselves (relatively) simple and easily communicable.

The task of behaviour-based approaches, though, is to account for the behaviour of complex as well as simple systems, and to elucidate mechanisms by which systems might progress from simple to complex over the course of ontogenetic development. Critically, for the behaviourist approach to be viable, an account must be provided of the internal reorganisation of system structure and function over ontogenesis which underlies the progressive *epistemic* adaptation of systems, allowing them to become ever-more inductively powerful over the life-span. To what extent does the behaviourist programme accomplish this task?

A double edged sword

The *anti-representationalism* of behaviourism allowed the paradigm to make great advances in the years after its inception. By focusing purely on simple behaviours with clearly defined operational measures whilst simultaneously denying the import of internal structures, it began to briefly seem as though objective descriptions of system behaviour might be achievable and, furthermore, that changes in system behaviour might be describable by an elegantly small set of principles.

However, notwithstanding this initial benefit, the laudable attempt of the forefathers of behaviourism to distance scientific psychology from the previously all-pervasive introspective methodology through the rejection of the explanatory value of all forms of internal representation went too far — dooming behaviourism to the characterisation of simple systems alone.

Trivialising induction

Whereas early behaviourist research examined relatively complex behavioural competences[12] later research focussed almost exclusively on the performance of simplified tasks[13] culminating in Skinner box experiments featuring a binary choice between pressing a lever to obtain food, or not. Such procedures served to maximise the rate of learning at the expense of severe restriction of the inductive space. In addition to

[12] For example maze learning using the Hampton Court Palace maze (Small, 1901).

[13] Largely as a result of the underlying assumption that the same learning mechanism underlies all forms of learning at all levels of complexity in all systems.

this inductive trivialisation, measures of learning became purely operational: *how* a response was made became irrelevant. Furthermore the puzzle box paradigm, especially, forced the agent to find the solution accidentally as a consequence of the fundamental impossibility of perceiving the functions of, and therefore the material relationships between, differing parts of the apparatus.

As there was little to learn within the restricted inductive space of classic behaviourist experiments it is unsurprising that research was unable to fractionate simple and complex systems — but such findings merely served to reinforce the fundamental assumptions of the behaviourist programme (Macphail, 1982, for example).

Origins and growth

What is learned, and how?

Behaviourism viewed adaptation as modification of the relationship between stimulus and response. All adaptation was considered to be of the same form, all tasks were thought to be subject to the same laws of learning regardless of the species studied or the complexity of the task. No account was taken of the vast range of types of learning or differences in competence across species. Furthermore it is unclear *what* is learned in many traditional behaviourist tasks; the stimulus–response model of learning has been subject to much criticism. Watson (1913) defined a response as a particular motor or activity pattern yet evidence suggests that it is not particular motor patterns that are learned but rather some 'purpose' or 'goal awareness' (Macfarlane, 1930; Tolman, 1932, 1948; Lashley, 1951). Yet if learning is not the modification of sensorimotor links then, in the absence of representational systems of any kind, what can it be?

To compound this problem, the role of reinforcement in animal learning is not fully understood. Notwithstanding the difficulties inherent in determining the nature of reinforcement required for a particular task, learning can occur both before (Tolman & Horcik, 1930; Tolman, 1932, latent learning phenomena), synchronously with (Guthrie, 1952), and subsequent to (Hull, 1952) the provision of reinforcement. Furthermore, modification of reinforcer *value* modulates response frequency (Adams & Dickinson, 1981) indicating that conditioning processes recruit epistemic/representational, in addition to purely behavioural, adaptive mechanisms.

How general are the 'laws' of learning?

One goal of the behaviourist research programme was determination of the laws of learning, those principles which might generate all forms of behaviour through learning processes possessed by, and fundamental to, all biological systems at every level of complexity. However, the *assumption of equivalence of association* underlying behaviourist research has been conclusively demonstrated to be false. Associations which are in accord with species-specific defence mechanisms, or relate to innate fixed action patterns are much easier to learn that those which are not (Garcia & Koelling, 1966; Bolles, 1970; Seligman, 1970; Hinde, 1982).

What of context?

In the biological realm stimuli do not elicit identical system responses at all times but rather invoke varying responses modulated by internal factors such as endocrine state, motivation, and memory. Within ethology, for example, the concept of innate releasing mechanisms as under-specified forms of input eliciting adaptive behavioural responses[14] has been abandoned with the recognition that there are many internal and external constraining factors influencing both learning and the control of behaviour (Hinde & Stevenson-Hinde, 1973). Furthermore, there is now much evidence to suggest that the capacity of biological systems to learn and the nature of what is learnt is, in important respects, biologically determined (Lorenz, 1935; Bateson & Klopfer, 1982; Hinde, 1982).

The complexity of behaviour

Lashley (1951) suggested that the problem of *syntax* is not exclusively linguistic but rather pertains to the organisation of all forms of activity. Movement, for example, presents the problem of sequences of actions that cannot be explained simply in terms of chains of motor responses. How does behaviourism explain the complex serial behaviours of natural systems which are pervasive within the biological world?

In keeping with the reductionist strategy of behaviourism it was argued (Skinner, 1938, 1953; Hull, 1934, for example) that all complex behaviours could, in principle,

[14] For example Tinbergen's (1952) discovery that a painted lollipop stick can release attack behaviours in the male stickleback.

be explained as chains of responses each of which provides feedback (either internal or external) which then becomes the discriminative stimulus for the consequent response. Although the concept of chaining of responses became widespread within the behaviourist paradigm, and more recent behaviourist-inspired research, many forms of possible chaining mechanism have been suggested and it remains unclear which, if any, of these occurs. A competence such as maze learning which necessitates sequential goal-approach behaviour might potentially be based upon many forms of chaining mechanism indicating that a multitude of mechanisms[15] are recruited by systems *as required.* No single underlying mechanism underlies sequential behaviour.

Yet if it is the case that systems, at all levels, pragmatically recruit multiple mechanisms what can the (non-representational) mechanism of this selection be? In laboratory situations where stimulation is held constant, behaviour is not always stereotyped, but permits of great variation. What mediates between stimulus and response in these cases?

Scaling what?

A further problem associated with the minimised induction of classic behaviourist experimentation concerns the nature of the acquired associations. The environment within which systems were tested was far removed from their natural habitats with the effect that the stimuli to which subjects were exposed, and the responses expected of them, were fundamentally *unsituated* and inherently arbitrary. Behaviourism focussed on an associative inductive machinery which forms only part of the complement of adaptive mechanisms which intelligent biological systems possess (Gazzaniga *et al.*, 1998). The associations constructed by this mechanism within the behaviourist paradigm, formed purely on the basis of repeated connections of particular objects and events (McGonigle & Chalmers, 1998*a*), leave little possibility of the acquisition of systematic knowledge (Fodor & Pylyshyn, 1988) or necessity (Smith, 1993) deriving from experienced material relationships between causally linked phenomena, and thus little basis on which to scaffold progressive epistemic system competences.

[15] Behaviourist researchers have identified a number of chaining mechanisms underlying serial behaviours including (but not limited to) response chains (Carpenter, 1984), stimulus-approach chains, place (stimuli-configuration) chains, and S-R-S* (stimulus-response-stimulus) chains (Tolman, 1932, argued that such chains might support economising inferences) with inference supporting models constructed by Gallistel (1980) and Schmajuk & Thierne (1992).

Summary

> "Science today stands on something of a divide. For two centuries it has been exploring systems that are either intrinsically simple or that are capable of being analysed into simple components. The fact that such a dogma as 'vary the factors one at a time' could be accepted for a century shows that scientists were largely concerned in investigating such systems as *allowed* this method; for this method is often fundamentally impossible in complex systems." (Ashby, 1956, p. 5, italics retained)

Ashby wrote this at a time of nascent paradigm shift — when the traditional reductionist methodology of western science was beginning to come under threat from the emerging sciences of complexity. His words remain equally valid today: almost a century's investigation of simple reflexive systems has failed to elucidate the fundamental characteristics of complex intelligent systems.

We have seen that the experimental methodology of the behaviourist paradigm was fundamentally unsuited to studying the behaviours of complex systems. Mishkin & Petri (1984) suggest that behaviourist tasks recruit phylogenetically older areas of the nervous system such as the *corpus striatum*, resulting in a levelling effect (Vygotsky, 1934/1961). There is no convincing evidence that those reflexive learning competences, which are dependent on more primitive structures within the nervous system, are extendible to more complex, especially epistemic, forms of adaptation (McGonigle & Chalmers, 1998*a*). Unfortunately, misunderstanding of the weaknesses of the paradigm has led to statements such as Macphails's (1982) null hypothesis: no qualitative changes in the adaptive competences of systems appear until humanity.

Furthermore, the behaviourist characterisation of even simple systems lacks credibility. What is learned remains unclear — it is not purely sensorimotor couplings but, without representations of any kind, what else can it be? The role of reinforcement is similarly confused: what is being reinforced, and when? Finally, by what mechanisms are these poorly characterised behaviours transformed into the richness of those possessed by complex systems? Over the past century no indications that systems can make the transition from simple to complex, inductively weak to strong, through associationistic mechanisms of behaviour modification alone have been reported (Fodor, 1983; McFarland, 1999; McGonigle & Chalmers, In press*a*).

1.3.2 Behaviour-based robotics

The reactive behaviour-based (hereafter BB) approach within AI, exemplified by subsumption architectures (Brooks, 1985), arose as an explicit rejection of the symbolic approach of classical AI (Brooks, 1991*c*, for example). Inspired by a view of intelligence developed by Minsky (1986) where multiple simple elements operate concurrently within limited problem domains, subsumption approaches adopt an explicit *anti-symbolic, situated agent* stance. Intelligent behaviour is viewed as the outcome of the internal interaction of multiple concurrent units together with interaction between the agent and its environment and is, therefore, characterised as *emergent.*

Behaviour-based subsumption architectures support system behaviour which is an immediate, and generally stereotypical, response to a given sensed situation. No representational system is instantiated[16] thus no planning stage intervenes between sensing and acting. Behaviours are engineered as sensor(s)–actuator(s) links, and are therefore best described as *tightly-coupled.* All behaviours operate concurrently with inhibitory connections between them allowing higher level behaviours to 'subsume' (or, take control of the system).

Primitives

Behaviour-based systems are constructed in accordance with the following guidelines (Malcolm *et al.*, 1989; Brooks, 1991*c*). Systems should be embodied and situated in the real world. System competence should be decomposed by activity rather than function. Intelligence should arise through emergence. Furthermore, systems should be autonomous, and tested in real-world, dynamic environments. These principles embody an explicit rejection of classical AI.

Critically, BB systems should scale-up: becoming increasingly more complex over time. Behaviour-based roboticists would follow the 'oxen-trail' of evolution, ultimately at-

[16] At least in first generation implementations. See below.

taining human-like levels of intelligent behaviour. With this end in sight, Brooks (1986*a*,*b*, 1989, 1991*c*) developed a distributed architecture for intelligence. The 'subsumption' architecture features individual 'layers' each one supporting a single task-achieving behaviour in order to optimise computation. Within a layer individual finite-state-mechanisms are often hardwired to sensors and actuators. Scaling of the system is envisaged to occur through the addition of new layers, which can suppress or inhibit their predecessors. Each layer of the architecture is designed to be continuously, and asynchronously active. In the absence of any form of central control the behaviour of the system is determined by the sensed environment. Initially, at least, no explicit representational system was incorporated within the system, although later implementations feature non-manipulable, deictic forms of representation (Matarić, 1990, 1991, 1992) or low-level ('endocrine') state (Brooks, 1991*a*).

The design principles described above were initially hailed as revolutionary. What is the current status of BB robotics[17]?; and how does it compare to its classical predecessor?

Embodiment

Brooks (1991*b*) provides two reasons why embodiment is important. The first is for validation purposes — an architecture for intelligence should be capable of supporting an embodied agent. In recent years this point has become widely accepted. This latter consideration makes it doubly strange that out of the 42 implementations discussed by Brooks (1997), only four were embodied.

The second reason, to provide a solution to the symbol-grounding problem (Harnad, 1990*b*) of the classical stances, is more problematic. Certainly "the world grounds regress" (Brooks, 1991*b*, p. 16) but for at least the first generation of *a*representational systems and, arguably, the second generation also, regress of what? Subsumption systems, like their classical predecessors yet for different reasons, are *semantically-blind.*

The tyranny of the flex[18]

[17] A review of current BB research by Brooks (1997) of articles published in *'Adaptive Behaviour'* (MIT Press), "the journal most closely aligned with the behaviour-based approach" (p. 294) will be referred to a number of times.

[18] Or, power cable.

The explicit aim of BB robotics is to

> "build completely autonomous mobile agents that co-exist in the world with humans, and are seen by those humans as intelligent beings in their own right [that are] able to maintain multiple goals [and] have some purpose in being" (Brooks, 1991*c*, p. 145).

Despite these laudable aims, and the purposive language used to describe layers within the subsumption architecture[19] autonomy, for BB robotics, is simply autonomy from the flex. BB systems lack that most vital ingredient of complex biological systems — internal control. It is claimed that subsumption architectures are "not constrained to be purely reactive" (Brooks, 1991*b*, p. 17) yet an example of a non-reactive layer of a subsumption architecture, the second layer of 'Allen' (Brooks, 1986*b*) which 'visits distant visible places', simply instantiates the rule "if space then move into it".

Subsumption architectures are not autonomous in any non-trivial sense. The state of the world determines the actions of the system at any given time[20] Furthermore, of 31 implementations discussed by Brooks (1997), 29 were capable of only *one* competence. It is unsurprising, therefore, that this review could only identify two papers (Tyrrell, 1993, 1994) which strove to deal with action selection — a fundamental control feature of all complex biological systems (see Lashley, 1951, on the problem of the syntax of behaviour).

This confusion of autonomy of power and autonomy of control is clearly illustrated by the *apologia* of Brooks (1991*c*) that the research programme suffers delay due to engineering difficulties[21].

Origins and growth

Emergent functionality

---

[19] For example, "layers can decide on the appropriateness of their goals" (Brooks, 1991*b*).

[20] 'Pick up soda can', one high level 'non-reactive' competence of 'Herbert' (Connell, 1989), a first generation BB robot *par excellence*, was in reality a hand-crafted condition of the absence of body movement.

[21] Specifically those of miniaturisation.

We saw previously that classical systems were reduced to operating in undynamic 'toy-worlds'. Situating systems in the real world from the outset ensures that dynamism is not crippling, allows the system to take advantage of environmental constraints, and to 'offload' computation. Brooks (1991*b*) also states that intelligence should emerge from the situated interaction of a system and its environment. This 'emergent functionality' (Steels, 1991) helps the designer to avoid one of the category errors of classical AI — complexity of behaviour does not directly imply complexity of internal control structures (Simon, 1962) hence: "intelligence is in the eye of the observer" (Brooks, 1991*b*, p. 16).

For BB robotics two levels of interaction, inter-layer and system-environment, are supposed to give rise to this 'emergent' functionality. However, in reality the behaviour of such systems does not truly emerge through system-environment interaction but is rather *hand-crafted* by the designer *at both levels*:

Inter-layer interactions "We advocate careful engineering of all the interactions within the system" (Brooks, 1991*b*).

System-environment interactions "as each layer is built it must be tested extensively in the real world. The system must interact with the real world over extended periods. Its behaviour must be observed *and be carefully and thoroughly debugged*" (Brooks, 1991*c*, italics added)

So we see that the functionality of BB systems lies not in emergence through interaction but is rather *installed* in the system by the designer. The final outcome is a system that exhibits behaviour which is as equally hand-crafted as that of classical robotics.

Scaling-up

The scaling of BB systems to ever more complex levels of adaptation is the mission of the paradigm. Complex biological systems become progressively more adaptive over both phylogeny and ontogeny. What kind of scaling does the BB paradigm support?

Ontogenesis BB systems' learning occurs intra-run. No reports of learning across runs have been provided. Notwithstanding the lack of scope for progressive adaptation over the life span of individual agents, what kind of learning can occur in BB systems? Currently, learning in first-generation systems is limited to recalibration of sensor-actuator links (Viola, 1990) or of locomotor interactions (Maes & Brooks, 1990). The second generation system reported by Matarić (1990, 1991) builds a deictic map of its surroundings based on local landmarks. This implementation, however, is an exception for BB systems and it is unclear whether any further (epistemic) derivations are obtainable from this representational scheme.

The latter example is instructive. The initial mission statement of the BB approach was that:

> "When we examine a *very simple level intelligence* we find that explicit representations and models of the world simply get in the way. It turns out to be better to use the world as its own model." (Brooks, 1991*c*, p. 140, italics added)

Since this seminal article, there has been much confusion over exactly what form of anti-representationalism BB robotics should adopt, with Brooks stating that his objection was to explicit and symbolic forms only (1991, 1997). Meanwhile representation, albeit implicit, or deictic (Agre, 1988; Chapman, 1990), has crept back in — presumably in recognition that complex behaviour requires internal control, and internal control, state. BB robotics could not even come close to achieving[22] its initial target — 'insect level' intelligence — without adopting some form of internal state. This is unsurprising: 'even' insects possess chemically-based state and modulate their behaviour according to *internal* as well as external criteria — identical stimuli do not always elicit identical responses. Indeed, Brooks' own extension to the basic subsumption architecture (Brooks, 1991*a*), is based on a characterisation of the hormonal system of the lobster (Kravitz, 1988). Furthermore, Brooks (1991*b*) concedes that a system "may sometimes need to build and maintain a map" (p. 19).

The review article of Brooks (1997), a long awaited response to the scepticism of Kirsh (1991), mentions 'self-adaptation' as an important area of future research. Even though at this stage Brooks himself, the founding father of BB robotics,

[22] No implication that this level has now been attained is intended!

accepts that it has not achieved the scaling of systems desired and therefore advocates a hybrid 'cognitive robotics' (p. 296), the 'self-adaptation' advanced is restricted to sensor-actuator (re)calibration, habituation, and sensitisation (p. 298). The lack of appreciation of the importance of internal control of behaviour, based on contextual interpretation of explicit knowledge strongly suggests that subsumption-based approaches, whether redescribed as 'cognitive' or not, are fundamentally unsuited to the construction of adaptive intelligent systems.

'Phylogenesis' We have seen that BB robotics has not developed systems capable of scaling-up over ontogenesis, what of incrementing the system through the addition of layers, the original path toward complexity proposed by Brooks (1991*c*)? A serious constraint on such scaling is that, as we have seen, emergence of behaviour really occurs through . Interactions between layers are unpredictable making it very difficult to scale the system (Brooks, 1991*b*).

Maybe this is the reason why, after a decade and a half of research, BB systems still exhibit only very simple behaviours. Brooks' (1997) review of 31 robotic implementations[23] *none* exhibit behaviour more complex than 'chasing prey' or 'avoiding predators'. Furthermore, only two papers (Tyrrell, 1993, 1994) even begin to address the issue of action selection. BB robotics still clearly has a long way to travel on the oxen-trail of evolution.

Summary

Behaviour-based robotics attempted to steer away from the pre-interpreted symbolic domains characteristic of classical AI by grounding the behaviour of agents directly in the world, yet the result became equally pre-interpreted behavioural systems. Focusing on behaviour and eschewing the role of internal control mechanisms and representational schemes has meant that little remains to be grounded — these systems are semantically blind.

The hope was for the emergence of high-level competences through the interaction of lower-level behavioural modules, yet despite the hand-crafting of such 'emergence'

[23] Twenty-seven of which were, in fact, simulated.

only primitive behaviours such as wall following, pursuit, and avoidance have been engineered. Higher level behaviours, such as map-building, have been demonstrated only after hybridisation with some form of representational system. It seems unlikely that such implementations, which form the minority of BB research will become much more advanced — partly due to the problems of hand-crafting emergence but also because the paradigm lacks both an accurate characterisation of the target systems and a principled path along which to travel toward them.

1.4 An impasse

Classical systems, whilst focussed on goal-directed, internally-driven behaviours are limited to highly constrained pre-interpreted domains of competence. Even within these domains, the systems are brittle and unreactive, their behaviour hand-crafted, and adaptation limited. The meaning of representation is ungrounded, and reflects the ontology of the designer rather than that of the system. Reactive behaviour-based systems feature tightly-coupled, hand-crafted behaviours. But internal control is largely absent; systems are not goal-directed but rather respond to the local environment. Whilst the behaviour of such systems is grounded in their environments, they are semantically-blind.

How should such an impasse be resolved? The most obvious solution to the problem is *hybridisation* of the two classes of system within a single architecture. What are the logical features of such architectures; and do they resolve, in practice, the problems inherent in either stance alone?

1.5 Hybrid systems

Superficially a hybrid approach appears promising. Classical systems are designed to mimic the competences of complex biological systems, specifically humans, and are therefore focussed on issues of representation and reasoning. However such systems are inherently unreactive and therefore lack robustness. Behaviour-based systems, in contrast, eschew representation, reasoning, and internal control in favour of tightly-coupled behaviours which are robust in the face of a noisy dynamic world. Surely hybridisation of these two approaches would lead to systems capable of both representing and reasoning about the world and behaving robustly within that world[24]?

What are the logical features of such a classical/behaviour-based hybrid architecture? The system must have a representational component. This system will be symbolic and manipulation of symbols will underlie reasoning and planning processes. The semantics of the representational system will be pre-installed by the designer yet 'grounded' in the world through the reactive behaviour of atomic units within the architecture. These atomic units will be hand-crafted[25] by the designer in order to achieve optimised reactive coupling with the world. Such systems should utilise their world model in order to plan actions within the world; the behaviour-based atoms should overcome noise and dynamism in order to reliably and robustly carry out the plan.

All well and good. Over the past twenty years a multitude of 'hybrid', or reactive planning, architectures have been developed. However, such architectures are generally not pure classical/behaviour-based hybrids but rather incorporate many new design features — such novel features were introduced in recognition that the problems pertaining to classical and reactive behaviour-based systems were not resolved by their simple hybridisation.

1.5.1 World models and grounding meaning

We have seen previously that the classical stance encountered innumerable difficulties through attempting to construct and maintain an exhaustive symbolic model of the world. Conversely, BB robotics eschewed such representational schemes. An obvious design feature of a hybrid architecture is a reduced world model, albeit symbol-based. Such a model should be easier to maintain, and help to overcome the combinatorial difficulties of the frame problem. Reducing state size in this way, however, raises the upper limit of behavioural complexity but does not provide a solution to the fundamental problem (Horswill, 1997).

One aim of the designers of hybrid systems is to ground symbolic representational sys-

[24] As advocated by Sloman (1998) for example.

[25] We have seen previously that the 'emergent' functionality of reactive behaviour-based systems is really a consequence of careful design.

tems in the world through reasoning about lower-level behavioural modules (Malcolm, 1997, p. 33). However, this tactic does not overcome the problem of domain preinterpretation. For classical systems the meaning of symbols is apparent only through human interpretation (Quillian, 1968; Winograd & Flores, 1986). Similarly for BB systems behaviours are hand-crafted to 'mean' something by the designer.

Hybrid systems feature low-level reactive behavioural units which are hand-crafted, and representations about the world[26] and their behaviours (Malcolm, 1997). The meaning of these representations of behaviours, however, remains pre-installed within the system. Such representations are no more grounded than those of traditional AI.

1.5.2 Robust execution of plans

So hybrid classical/behaviour-based systems can be no more grounded than their classical predecessors — their domain of competence remains pre-interpreted: representations are of pre-determined external features of the world and of reactive behaviours pre-engineered to fit a target niche. But what of robustness? The most promising outcome of hybridisation seemed to be augmenting planning with robust plan execution. Plans would be constructed on the basis of an internal symbolic model of the world and then implemented by behaviour-based modules which would be able to carry out their functions reliably, reacting to changing external situations.

The strategy certainly appears to be an improvement over either planning and failing, or blindly reacting yet, on *a priori* grounds alone, it is insufficient — context is ignored. The robust, goal-directed behaviour of biological systems relies on the ability to distinguish situations and determine appropriate action. How does the strategy work in practice?

Although pure classical/BB hybrid systems have met some success in the field of assembly robotics[27] (Malcolm, 1987, 1998, for example) autonomous agent implementations have adopted an alternative strategy. Indeed the majority of hybrids do not rely on autonomous reactive behaviour-based modules to provide robustness but rather in-

[26] Susceptible to all of the criticisms levelled against classical AI systems.

[27] Albeit in domains which are accepted to be of reduced complexity and dynamism (Malcolm, 1997, 1998).

terleave planning and acting in an attempt to reactively respond to changing events whilst retaining internal control. Unfortunately the lack of success achieved by systems such as Schoppers (1987) universal plans[28], and McDermott's (1978) NASL[29] have led to the development of *situated* planning implementations (Linney, 1991; Levison, 1996). Such systems are *contextually-informed*[30]. Furthermore, the difficultly of robustly implementing plans in a dynamic world, despite hybridisation, has led to implementations which stress rational behaviour sequencing (Firby, 1989; Gat, 1991*b*; Simmons, 1991; Gat & Dorais, 1994), replanning (Gat, 1991*a*) and error cognisance and recovery (Simmons, 1991; Gat, 1991*b*).

1.5.3 Summary

Hybridisation of conventional classical and behaviour-based architectural components cannot resolve the impasse. Such systems must remain preinterpreted — now at both representational and behavioural levels. The meaning of representations is preinstalled within such systems. The problem of grounding meaning has not been resolved — representations now denote elements of a preinterpreted external world together with hand-crafted behaviours.

Furthermore such systems are brittle. The most successful implementations of pure classical/BB hybrids operate in constrained domains such as assembly robotics (Malcolm, 1997, 1998). Indeed, hybrid autonomous agents operating in dynamic domains remain susceptible to plan failure and consequent system collapse.

Architectures which support more robust behaviour, although described as 'hybrid', move beyond pure classical/BB elements, incorporating contextual knowledge, behavioural sequencing, and error recovery.

---

[28] The strategy adopted is to construct a plan for every possible world state at compilation and select between them at run-time. Plans perform conditional tests in an attempt to predict all run-time contingencies.

[29] A characteristic reactive planning implementation. Planning and action are interleaved by a cycle of sensing, plan generation and execution of one step of the plan. Unfortunately, although robust in the face of environmental changes, this approach was found to generate many errors due to unforeseen interactions between steps.

[30] Firby's (1987) RAPS, PENGI (Agre & Chapman, 1990*a*), ItPlanS (Geib, 1992; Geib *et al.*, 1994) and the situated automata of (Kaelbling & Rosenschein, 1990), for example.

1.6 Summary: two dogmas

Neither symbol-based nor behaviour-based approaches adequately characterise intelligent systems. The classical mischaracterisation of disembodied and unsituated symbol using systems leads to failure both to root the competences of complex systems in earlier adaptations, and to account for the origin of meaning for such systems. The focus of behaviourism on tightly-coupled responses to local events cannot explain the richness of behaviour of complex systems, nor the processes by which simple systems become complex over phylogeny and ontogeny.

The consequences for the construction of artificial agents are devastating. Classical AI has succeeded in producing only brittle and unreactive systems operating in pre-interpreted toy worlds. The systems are hand-crafted, non-adaptive and semantically blind. Similarly, reactive BB systems, although robust, are capable only of tightly-coupled, hand-crafted behaviours. Control of the system resides with the local environment.

The most obvious solution, hybridisation of conventional classical and BB components, does not resolve the impasse. Such hybrid systems would remain hand-crafted. Their domains would be preinterpreted at both representational and behavioural levels. Finally, meaning would still be ungrounded. Indeed, such pure classical/BB systems are rare and successful only within limited domains. The majority of hybrid systems are fundamentally *situated* — incorporating elements of contextually-marked non-symbolic representation. Internal control is achieved through rational sequencing of behaviours together, in some implementations, with the ability to recognise, and attempt to recover from, error.

Neither of the traditional stances, nor their conjunction, have supported the development of robust and adaptive artificial intelligent systems. However, their joint failure has stressed the importance of a characterisation of systems which recognises the constraints provided by embodiment and the situated nature of action and cognition. Yet without an accurate characterisation of complex *biological* systems, the targets of design, what hope can there be for progress?

The next chapter discusses two frameworks which seek to replace the traditional characterisations of complex systems. Both are inspired by characterisations of biological systems and have inspired the construction of artificial systems. Do either of these stances resolve the problems associated with symbolic and behaviour-based positions; and what are their implications for the construction of adaptive artificial systems?

Chapter 2

Two contenders: dynamic and situated perspectives

We have seen that neither symbol-based nor behaviour-based stances adequately characterise intelligent biological systems. Nor does either stance support the development of robust adaptive artificial intelligent systems. The disembodied reasoning systems of the symbolic stances are contrasted with the undeliberative reactive systems of the behaviour-based approaches. We have seen that purely reactive systems are forever at the mercy of their environments, lacking internal control. Deliberative systems, in contrast, lack the reactivity required to successfully achieve their internally motivated goals. Hybridisation assumes that cognition serves to supplant and extend purely reactive behaviour yet in the attempt to recombine action, perception and cognition it is unclear how to make low-level reflexes sensitive to cognitive goals (Laird & Rosenbloom, 1990; Lewis *et al.*, 1990).

This chapter introduces the 'emergent' stances, dynamicism and situated theories, which seek to overcome this impasse by assuming coordination of perception and (situated) action, within a given niche, as the basis of higher-level, cognitive competences. Such higher-level competences are thought to emerge through inter-system and system-environment[1] interactions.

[1] Precursors of this position are clear within: biology and ethology where organisms are conceptualised with respect to their adaptive environment (Haldane, 1917; von Uexküll, 1928); neurophysiology where the nervous system is construed as such only when tied to an environment (Bethe, 1931; Anokhin, 1978); and psychology (Bretano, 1874; Dewey, 1896; Köhler, 1927; Merleau-Ponty, 1962; Gibson, 1979).

The implications of these stances for the construction of artificial systems are discussed later in the chapter. Initially situated robotic implementations are described: where classical robots centralised control but were unreactive and behaviour-based robots relied on control exerted from without in the designer's attempt to distance them from over-centralisation, 'situated' robotics strives to utilise constraints present in environment and agent, whilst retaining contextually-informed internal representations. The designers of such systems strive for both reactivity and appropriate adaptive behaviour, under some internal control.

Next, 'emergent' computational approaches, connectionism and ALIFE, are briefly reviewed. Finally their conjunction 'evolutionary robotics', which strives to develop adaptive systems through evolutionary and self-organisational mechanisms, is examined.

2.1 The dynamic perspective

> "The problem of origins requires an understanding of how new levels of order emerge from complex patterns of interaction and what the properties of these emergent structures are in terms of their robustness to perturbation and their capacity for self-maintenance. Then all levels of order and organisation are recognised as being of equal importance in understanding the behaviour of living systems, and the reductionist insistence on some basic material level of cause and explanation such as molecules and genes can be recognised for what it is, an unfortunate fashion or prejudice that is actually bad science" (Goodwin, 1994, p. 168).

Throughout western science and philosophy for more than two thousand years the tension between a focus on form or process has remained unresolved. First the Eleatics, and then rationalism, and finally cognitivism concern themselves with *being*. Conversely, beginning with Heraclitus and, much later, empiricism and behaviourism the emphasis is on *becoming*. Neither stance alone is sufficient to deal with the problems of complexity. Dynamic systems theory offers itself as the solution by emphasising the interrelationship of both: becoming through being, and being through becoming

(Prigogine, 1980).

The empirical success of dynamic systems theory, now the most powerful explanatory framework within physical science, has naturally led to a desire to apply these explanatory tools to cognition (Turvey & Carello, 1981; Swenson & Turvey, 1991). Indeed, within cognitive science the dynamic hypothesis has become increasingly popular over the past decade (Globus, 1992; Robertson *et al.*, 1993; Thelen & Smith, 1994; van Gelder, 1995; van Gelder & Port, 1995, for examples), arising in part from the unresolved problems of, and general dissatisfaction with, both symbolic representationalist and associationist behaviour-based stances (Bechtel & Abrahmsen, 1991).

2.1.1 What is a dynamic system?

A system is defined as a set of interdependent variables, with the state of the system equating to the value of all variables at a given time, and its behaviour corresponding to state transitions (von Bertalanffy, 1967). Initially the field of complex dynamic systems dealt only with closed systems — ones in which there is no transfer of matter or energy between the system and its environment. Such systems exhibit reversible processes, in accordance with the second law of thermodynamics, ensuring that the system, over time, reaches a state of maximum entropy — one of maximal homogeneity and minimal organisation (von Bertalanffy, 1967; Prigogine, 1980). Since closed systems are subject to the second law of thermodynamics, perturbations of the system tend to result in self-regulations towards the inevitable maximum entropy state.

Open systems, in contrast, exchange matter or energy with their environment. Such systems exhibit irreversible processes, and therefore never attain final equilibrium states, but instead fall into steady-states where the system's functional organisation is maintained with respect to its environment. Open systems are capable of self-regulation following disturbances of system processes, this fact explains why similar systems tend to converge on similar goal states. However, changes in system conditions cannot be regulated: the set of possible trajectories is altered and the system may settle into final states very different from those which were initially possible. Final system state can therefore be seen to be the result of constraints imposed by the initial system conditions together with environmental influences (Köhler, 1927).

Organisms are examples of open systems, in continual interaction with their environments. The irreversible processes occurring in open systems underly the unidirectionality of growth in organisms. The somewhat paradoxical growth of structure in biological systems, as opposed to increasing entropy as is predicted by the second law of thermodynamics, is explained by the fact that open systems and their environments as a whole tend towards maximum entropy. Such a situation allows for local growth of structure, simultaneous with a global trend towards homogeneity. The cognitive system can be regarded as an abstract system implemented by a concrete system at a lower level (van Gelder, 1998) and can therefore be subjected to a dynamic systems analysis.

As the universe can be conceived as a hierarchy of dynamic systems the demarcation point between systems must be *conventional* and is therefore chosen with a theoretical question in mind — the level of analysis appropriate to the question determines the demarcation point. A system should therefore be conceptualised as an arbitrary restriction of the total environment (de Lorenzana & Ward, 1987). For many questions, including analysis of intelligent systems, a number of interlinked systems may have to be considered simultaneously. All levels are explanatory to varying degrees in complex hierarchical systems:

> "The theory of boundary conditions recognises the higher levels of life as forming a hierarchy, each level of which relies for its workings on the principles of the levels below it, even while it itself is irreducible to these lower principles" (Polyani, 1976, p. 136).

Intelligent biological systems consist of many cognitively-relevant dynamical systems and might also be partially located outside the body (van Gelder, 1998; Clark & Chalmers, 1998).

2.1.2 The dynamic perspective in cognitive science

A dynamic approach within cognitive science implies that: (1) cognitive systems should be regarded as nested complex dynamical systems; and (2) cognition should be understood dynamically, using dynamical models and tools (van Geert, 1994; van Gelder, 1998).

Dynamic systems theories shift concentration from consideration of structure towards process. The important property of complex dynamic systems is not their physical substrate but rather their *relational order*. Whereas traditional cognitive science focuses on state, static internal structure, an input-output metaphor, and a serial view of what states developing systems pass through using the language of symbolic representation and logical coordination, the dynamical perspective emphasises state changes, continuous internal processes leading to self-organisation, and coupling between agent and environment employing the terminology of system theory — states, parameters, attractors, trajectories *etc.*

Where traditional symbolic approaches tend to disregard the temporal dynamics of cognition, and behaviour-based stances are limited to change over time in simple systems, through descriptions in terms of the continuous self-organisation of systems in interaction with the environment and each other, the dynamic perspective reintroduces time. The characterisation of disembodied, abstracted, symbol grinding cognitive systems is replaced by an embodied, coupled, dynamic view of intelligent biological systems. This characterisation seeks to explain the features of cognition ignored by traditional accounts: its embodiment in nervous system, body, and environment; its continuous dynamic nature; and its emergence and stability over time.

The dynamic perspective has been recently criticised for being *only* metaphorical and thus lacking explanatory value (Robertson *et al.*, 1993; Barton, 1994, for example). The central issue is the *predictive* ability of dynamic models — it has been argued that the application of rigorous terminology from nonlinear dynamics to often imprecise psychological variables, together with the difficulty of generating simulations of cognitive systems, is a critical failure for the dynamic model of cognition (Eliasmith, 1996). It is still early days, however, and there have been few attempts to construct truly dynamic models of cognitive and developmental processes[2]. Furthermore, this lack of rigour in newly emerging disciplines is common within science: as Hesse (1988) suggests, the value of a new model, or a novel application of a trusted model, lies in the

[2] See the olfactory bulb model of Skarda & Freeman (1987) and the more recent work of van Geert (1994) for examples.

exploitation of a successful and familiar explanatory framework from one domain to describe less well characterised phenomena in another. Only time will tell if the dynamic perspective proves to provide predictive models of (aspects of) the cognitive system, but at the current time it brings a rich descriptive language to the characterisation of intelligent systems.

The *state space* of a system refers to the possible range of variation in system behaviour and structure over its existence. State space is dually constrained by the *lower ontological bounds* of the system[3] and its environment[4]. Viewing the cognitive system as one example of such a system leads to an emphasis on the constraints imposed by 'wetware' (sensory organs, their structure, effectors *etc.*), 'software' (cognitive and perceptual primitives), and environment on developmental trajectories. Interactive system-environment processes determine, over the course of epigenesis, the *trajectory* taken through state-space. *Attractors* in state space lead to dynamic stabilisation of the system — the emergence of a temporary stable state.

System theory provides many possible mechanisms of change over both phylogeny, and ontogeny. Lorenzana & Ward (1987), for example, postulate two mechanisms which might underlie the trend toward complexity in evolutionary systems. *Combinatorial expansion* is the process by which a core system converts its unrealised potential (given by its global properties) through a process of progressive differentiation. The ontological lower bounds of the system set a limit to the expansion possible hence *generative condensation* occurs at the limit stage: the expanded system 'condenses' to form the ontological lower bound of a new system.

The emergence of detailed structure and refined behaviour through the *self-organisation* of systems over ontogenesis can occur through processes of spontaneous pattern generation and *bifurcation*, or symmetry-breaking[5], which leads to progressive hierarchical differentiation. A critical consequence of the dynamic stance is that no end point to development is envisaged. Systems follow open-ended trajectories, becoming increasingly

[3] For example, see Goodwin's (1994) discussion of the actions of genes over morphogenesis to select and stabilise the alternative forms which are available to an organism (p. 16).

[4] The generative layer of a system, its ontological lower bounds together with the richness of the environment, constrains development of evolutionary systems (de Lorenzana & Ward, 1987).

[5] Bifurcation from spatial uniformity to pattern occurs in non-linear systems when energy flow through the system displaces it from equilibrium (Goodwin, 1994).

differentiated over time.

2.1.3 Implications for the construction of artificial systems

Currently the dynamic perspective has been embraced most wholeheartedly by the ALIFE, and evolutionary robotics communities. Here, the focus of dynamicism on coupled systems, changes in system structure over time, and environmental influences on system development has been translated into a manifesto for synthesising intelligent artificial systems through the evolution of connectionist networks. Notwithstanding the paradox of the adoption of a stance which emphasises internal *self-organisation* as the key to progressive system development by one which concentrates on blind selective processes, ALIFE focuses on behavioural adaptation over evolving populations and is not, it is suggested[6], the way forward.

What else does the dynamic perspective suggest? Obviously a key feature of artificial systems must be their coupling to the environment, we shall see that this plays a key role in many situated theories (section 2.2) and implementations (section 2.3.1). It is important, however, that the dynamic perspective is not hijacked as a means of justifying reactivity in the absence of internal control. The rejection of the explanatory value of representation in the characterisation of cognition by many dynamic theorists (Beer, 1995*a*,*b*; Freeman & Skarda, 1990; Harvey, 1992; Husbands *et al.*, 1995, for example) is a worrying example[7]. Its postulated replacements, shared parameters in coupled differential equations (van Gelder, 1995) for example, are unlikely to lead to a clearer understanding of the behaviour of cognitive agents, and provide no clear guidelines for the construction of artificial systems. Critically the assertion of many dynamic theorists that representation is not needed to explain cognitive phenomena is alarmingly reminiscent of behaviour-based stances. Not only do these fail to account for the behaviour of complex systems but were also inadequate for characterising even basic animal learning phenomena (Thagard, 1992).

[6] See section 2.3.2.

[7] "In realistic nets [...] it is not the representations that are changed; it is the self-organising process that changes via chemical modulation. Indeed it no longer makes sense to talk of 'representations'," (Globus, 1992, p. 302); "it is the concept of representation which is insufficiently sophisticated" (van Gelder, 1993, p. 6); "We are not building representations at all!" (Thelen & Smith, 1994, p.338).

The most promising contribution of the dynamic perspective lies in the ideas of constraint, self-organisation, and unbounded development. It suggests that we should strive to build systems whose lower ontological bounds (system primitives), and environment of application, are sufficiently rich to enable, through interaction and consequent self-organisation, the development of progressively more adaptive competences. Critically, it suggests the adoption of a life-historical approach which appreciates the influence of past experience on present and future competence, and implies that systems should be constructed with no pre-determined end-point in sight.

2.1.4 Summary

The dynamic perspective, although still in its infancy within cognitive science, provides a rich metaphorical language for the characterisation of complex intelligent systems. Emphasis is shifted from the disembodied reasoning of symbolic stances towards the dynamic interaction of the coupled system-environment meta-system — intelligent systems are embedded within the real world. Initial system conditions and environmental potential dually constrain the developmental trajectories open to systems. Self-organisation occurs over ontogenesis leading to more highly differentiated systems. No developmental end-point is specified.

For the construction of artificial systems, the main contribution of the dynamic perspective lies in appreciation of the constraints imposed by system primitives and environment on the range of possible self-organised developmental trajectories not in a rejection of representation and consequent justification of pure reactivity.

2.2 Situated cognition

> "Expert-knowledge is comprised of context-dependent, personally constructed, highly functional but fallible abstractions." (Agnew *et al.*, 1993)

The situated perspective on cognition aims to rectify the abstraction mistake of the classical symbolic stance. Systems are viewed, as from the dynamic perspective, as fundamentally and influentially *embodied* and *situated* within an adaptive niche. The cognitive system is not abstracted away from its physical roots but rather these roots constrain the nature of intelligent systems, and provide a method of grounding meaning in experienced reality.

The conflation of external symbolic representations and internal mechanisms characteristic of traditional cognitive science is eschewed as a category error (Winograd & Flores, 1986; Clancey, 1991*a*). Inspired by ideas such as relativity, Heisenburg's uncertainty principle, quantum indeterminacy, and Gödel's theorem which emphasise the contingency of 'truth', together with the Kuhnian (1962) view of the contextual interpretation of 'objective' science, situated stances deny that human cognition can be accurately modelled by 'objective' context-independent symbolic descriptions since the validity of inferences derived from such symbolic knowledge varies in differing interpretative contexts (Dreyfus, 1972, 1992; Menzies, 1996).

Since the influence of the environment is so great, we must characterise systems in terms of their direct interaction with the environment rather than postulating intermediary symbolic deliberative devices. Control of systems is passed, therefore, from logical coordination of internal symbols to direct coordination of perception and action, modulated (in some characterisations) by indexical, contextually-informed representations. According to this view, symbolic descriptions are less important[8] than the dynamic internal deictic representations that are constructed through direct experience of the world[9].

Thus situated stances adopt a constructivist epistemology, where meaning is derived from system-environment interaction and is valid only within limited domains. Furthermore, progressive adaption of systems is characterised as 'emerging' through similar dialectic interactions within a given niche. The implications for the construction of artificial systems constitute a progression from both classical AI and behaviour-based robotics. Pure reactivity is rejected, as is internal deliberation, in favour of 'coordination' mechanisms which, it is hoped, will produce adaptive behaviour. Representation

[8] In fact, post-hoc justifications of a non-symbolic internal coordinative mechanism according to Clancey (1993).

[9] "The neural structures and processes that coordinate perception and action are created during activity, not retrieved and rotely applied, merely reconstructed, or calculated via stored rules and pattern descriptions." (Clancey, 1993, p. 94).

is either absent (implicit in the coordination mechanisms) or indexical-functional (Agre & Chapman, 1990*b*) — agent-centred.

2.2.1 Embodiment and its constraints

Perception and motor skills, for situated stances, are the 'hard' problems solved by intelligent systems; their solutions impose constraints on the remaining aspects of natural intelligence. Indeed, for situated stances all thoughts are acts, rooted in previous bodily acts — often cited as a possible reason why so much of language consists of bodily metaphors and metonymies (Lakoff & Johnson, 1980, for example). Where behaviour-based stances were mind-less, and symbolic stances disembodied, situated stances emphasise the bodily roots of cognition.

Recognition of the fundamental importance of embodiment is not new to situated stances of course: although much of western philosophy equated the laws of deductive logic with those of thought, the naturalistic epistemological tradition, a clear philosophical predecessor of situated cognition, holds that the cognitive system is a natural phenomena interacting with other natural phenomena, and that epistemology must be founded on empirical studies of the human cognitive system (Shimony, 1987, p. 1).

Spencer's *'Principles of psychology'* (1855), published only shortly before the *'Origin of species'* (Darwin, 1859), stressed that human mentality should be studied from an evolutionary perspective. However Spencer, and other early naturalistic epistemologists[10], regarded the process of cognitive adaptation as closed (Capek, 1987, p. 95). The environment to which the cognitive system was thought to be adapted was the total 'cosmic' environment and the fundamental concepts of classical science were regarded as accurate descriptions of the objective features of reality. For these early naturalistic epistemologists, evolution was complete — a view which paradoxically resulted in similar conclusions to Kant and other *apriorists*: a logico-symbolic characterisation of the cognitive system.

A second school of thought developed through the ideas of Bergson (1944), Piaget (1971), Popper (1972), and Campbell (1974). For this second group evolutionary

[10] Mach (1959), Avenarius (1890), von Helmholtz (1873), Poincaré (1946) for examples.

adaptation was still ongoing with both physical and cognitive structure dependent on this adaptive process. Popper (1985) makes the distinction between explicit knowledge, available to the system in terms of thought, language *etc.*, and implicit knowledge, embedded in the biological machinery of a system (Michalski, 1986; Longo, 1994), such as homeostatic mechanisms (Belew, 1992). The adaptive process by which both explicit and implicit forms of knowledge are constructed was thought to be (neo-)Darwinian evolution.

Analogously, within ethology (Uexküll, 1864/1944), and psychology (Miller, 1966), theorists had been considering the possibility of a disembodied mind and reflecting on the constraints provided by embodiment on cognitive structure. Early appreciation of how the *umwelt*, or perceptual world, of a system is highly constrained by architectural sensorimotor features of the system has been recently supplemented by the work of Ballard and colleagues (1993, 1997) who emphasise the interaction of the constraints imposed by the physical structure of a system with cognition.

Computational and power resources, for example, provide a fundamental overarching constraint on natural cognitive architectures. Systems must be able to process information rapidly in the face of a dynamic world. Computation must be fast, efficient, and *economical*. Systems might need to minimise the 'quantity' of working memory at any given time by utilising constraints provided by their own physical structure (Ballard *et al.*, 1997) and by 'extending' computation, *via* natural and engineered niche constraints, into the environment (Clark & Chalmers, 1998).

Resource constraints suggest that features such as the following might be *necessary* features of intelligent systems:

Hierarchy and modularity The physical structure of the brain reflects (at least a partial) modular structure. It has been suggested that this is an outcome of the time penalty of conducting neuronal signals across long distances, which would make local computation, and thus a modular, hierarchical organisation most efficient (Newell, 1990; Ballard *et al.*, 1997). Furthermore this kind of organisation would provide robustness (Simon, 1962; McGonigle, 1991; Gat, 1991*a*).

A need for both rapid processing at low levels in order to cope with a dynamic world, in combination with slower, more 'deliberative' processes suggests a hierarchy of control structures operating at increasingly slower, and more abstracted levels (Newell, 1990; Ballard *et al.*, 1997). Furthermore, a modular construction can be utilised to constrain the potential inductive space, thereby supporting a 'divide and conquer' strategy, which enables systems to solve problems which might otherwise by inductively intractable (McGonigle, 1991; Luger & Stubblefield, 1998*b*).

Sequential processing The work of Ballard *et al.* (1997) on saccadic eye movements suggests that serial order is a low-level property of (at least the the primate) visual systems, inherent in perceptual processing rather than a feature of high-level competences alone. As Lashley (1951) stated, serial competence is a fundamental feature of the nervous system rather than a novel development consequent upon a linguistic competence as implicitly assumed by traditional cognitivism.

System-centred representation Computational resource limitations indicate that an internal rather than external frame of reference is more likely to be utilised by systems (Ballard, 1991). Representing features of the world dynamically, from an egocentric viewpoint, rather than in a global reference system reduces memory load and yields simplifications of algorithmic complexity (Ullman & Richards, 1984; Agre & Chapman, 1987; Jeannerod, 1988; Milner & Goodale, 1995).

2.2.2 The ontogenesis of situated systems

For situated stances, as with the dynamic perspective, interaction between system and environment is the principle method by which development is thought to occur. Situated theories are constructivist: dialectic processes are regarded as the key to system scaffolding, and to grounding systems in their environments. Action is central to development (Bertenthal & Pinto, 1993; Rutkowska, 1993; Costall, 1994).

Although a multitude of mechanisms have now been proposed for this interactive process, *emergent functionality* and *enaction*, representative of the two extremes of situated approaches, will be briefly described in order to present the flavour of situated theories.

Emergent functionality is currently the major theoretical principle underlying situated agent approaches within robotics as it was previously for behaviour-based robotics. The complex abilities of situated systems are assumed to emerge indirectly from the operation of independent, simpler components in the absence of centralised control and (traditional symbolic) representational systems through inter-system interactions and transactions with the environment (Maes, 1990; Steels, 1991, for example). It has been argued that emergent functionality is most suited to situations:

> "when there is a lot of dependence on the environment, and it is difficult to foresee all possible circumstances in advance" (Steels, 1991, p. 459).

Rutkowska (1994, 1997) has applied this framework to human development arguing that the range of possible motor activities might be selected on the basis of environmental interaction rather than hardwired into the system, with novel coordinations established through environmental change. Temporarily engineered competences might later become permanent adaptive changes.

Emergent functionality is generally applied to coordination of responses to environmental stimuli, and motor coordination rather than more complex competences: it is hoped that these will emerge over time. It remains unclear, however, exactly how emergent functionality might operate, and especially how it might be incorporated in artificial systems. We saw previously that the reactive behaviour-based paradigm, although stressing emergent functionality as an important design principle, resorted to the hand-crafting of behaviour. Furthermore, this perspective emphasises reactivity through coordinative mechanisms, what are the roles envisaged, if any, for internal control and representation?

Enaction "The basic notion is that cognitive capacities are inextricably linked to a history that is lived, much like a path that does not exist but is laid down in walking. Consequently the view of cognition is not that of solving problems through representations, but as a creative bringing forth of a world where the only required condition is that it is effective action: it permits the continued integrity of the system involved." (Varela, 1989, pp. 59–60)

> Enaction is the 'bringing forth of the world' through activity (Varela, 1979; Maturana & Varela, 1980; Varela *et al.*, 1993). But whereas for the arch-constructivist Piaget (1970) action gives rise to explicit representations that model selected aspects of an objective reality from which systems construct ever more abstracted knowledge structures, from this perspective development is a widening of constraints on action. Increased complexity of behaviour relies not on logical co-ordination of knowledge structures, but on action-centred representations of the preconditions for successful behaviour (Willatts, 1989; Rutkowska, 1993).
>
> Enaction, therefore, emphasises contextual representation of the potential of actions as a key factor underlying progressive system development.

Situated stances although clearly rejecting *symbolic* representations in an attempt to avoid the associated combinatorial and grounding problems lack unity on its replacement within a theory of cognition. At one extreme are those who eschew explicit representation of any kind, and on the other those who push representation out into the world and, between these two, those who accept some form of egocentric, non-symbolic *pragmatic* representation.

The non-representational side of situated theories stresses direct coordination of behaviour — maintenance of an adaptive dynamic relationship between system and environment is critical. Merleau-Ponty's (1962) and Dreyfus' (1996) 'homeostatic' models fall into this category. The focus on coordination mechanisms leads to a view of systems untroubled by representational systems. This view could inspire artificial systems of only limited improvement over those of reactive BB robotics.

The majority of situated cognition theorists suggest that system-centred, contextually-informed representations underlie coordination. Motivated by Dretske's (1988) analysis of natural systems of representation whose meaning is intrinsic and dependent on construction rather than objective, these representations establish selective correspondence with the world; their function is not to express states of (external) affairs but rather to allow a system to maintain a functional relationship with its external environment.

Suggestions for representational systems which would support the maintenance of such a relationship include:

- Implicit representation
  - Clancey (1995) stresses that situated cognition researchers claim no internal representation in the first-person sense — as structures created, interpreted, and manipulated by a system. Rather representation/symbolising occurs through behaviour rather than between sensing and acting.
  - Rosenfield (1988) discusses 'process memory' which coordinates behaviours in space and time as they have been previously coordinated.
  - Roitblat (1991) maintains that any change internal to a system as a result of experience which later influences behaviour can be considered to be a representation.

  These stances, therefore, are neutral with respect to the construction of representational schemes in artificial systems.

- Pragmatic representation
  - Thornton (1997) suggests that replication of external 'triggers' for behaviours within a system is critical. Salient structural features of the environment should be replicated by any number of alternative forms of representation.
  - Prem (1997) argues that representation is based upon the affordances (Gibson, 1979) of perceived objects. Representations are inherently functional and meaningful.
  - Chapman & Agre (1987); Agre & Chapman (1990*b*) suggest that representations must be both functional and indexical — system centred. Representational systems need not be global nor objective as long as they permit contextual features to be encoded and allow for adaptive behaviour.
  - Ballard (1989, 1993) stresses the importance of deictic representation. These are inherently egocentric and support contextual interpretation of perception and action.

  For these stances meaningful structures are not pre-given nor static but rather arise through interactive processes, and are continually reinterpreted over the lifetime of a knowing subject. Meaning enters a system through contextually-marked activity (Suchman, 1987; Agre, 1988; Clancey, 1991*b*). Continual lifelong construction, interpretation, and reinterpretation of meaningful contextual representations also avoids both the frame problem, and the combinatorial problems associated with search through a global symbolic model of the world (Bickhard & Richie, 1983). The implicit idealism of symbolic stances is also rejected: information is not filtered from a pre-given ontology but rather constructed through perceptual processes (Maturana, 1983; Reeke & Edelman, 1988). Representation is best characterised, therefore, as a method of controlling action and perception in the environment (Clark, 1994). Situated inference (Barwise, 1987), for example, involves system exploitation of environmental regularities. Reasoning in this way depends on the embedding circumstances of the system: if the relevant external regularities break down any inferences dependent upon them will become invalid. Dretske (1988) has suggested that infants' avoidance of looming objects can be interpreted as an example of situated inference — the soundness of the 'inference' rests on the continuation of natural environmental conditions. The postulation of this kind of mechanism, based on action contingent transformation, represents one attempt to demonstrate that *rational* behaviour can be generated from action-based internal representation in conjunction with niche constraints in the absence of a logically-coordinated symbolic model of the world.

This view of 'sub-symbolic' (Smolensky, 1988), pragmatic representation forms the basis of situated robotics implementations[11].

- Extended mind

  – "epistemic action demands spread of epistemic credit" (Clark & Chalmers, 1998, p. 10).

  Clark & Chalmers (1998) describe the tendency of human reasoners to rely on environmental support (McClelland *et al.*, 1986; Clark, 1989; Hutchins, 1995, for example). Delegation of information-processing to both natural and engineered features of the external world is a means of increasing system capacity. The 'loop' of system and environment plays a causal role in

[11] See below section 2.3.1.

> the generation of behaviour and must, therefore, be well characterised. Systems take advantage of niche constraints, and their own bodily constraint, to simplify the solutions to problems. Furthermore, the structure of the environment plays a crucial role in constraining the way in which systems can develop.

The emphasis of the situated perspective on *contextual* representation is supported by recent neurophysiological evidence which suggests that the activity of even 'low-level' areas of cortex is highly context-dependent. Activity of neurons in the visual area V1 (Gallant *et al.*, 1995) is dependent on fixation point location which, in turn, depends on high-level goals. The parietal cortex has been discovered to contain neurons which are sensitive to exocentric (task-relevant) location in space (Stricanne *et al.*, 1994; Snyder & Andersen, 1994), and activity of cells in motor cortex suggests the use of exocentric frames of reference (Helms-Tillery *et al.*, 1991; Tagaris *et al.*, 1994; Pellizzer *et al.*, 1994), prismatic adaptation has also been found to be context sensitive (Flook & McGonigle, 1977).

Lifetime learning

Situated stances are clearly influenced by the dynamic perspective. An appreciation that the designer's ontology should not be pre-installed within the systems at either representational or behavioural levels (Clancey, 1991*b*, for example) reflects the emphasis of dynamicism on the constraints provided by the lower bounds of a system. The coordination mechanisms assumed to underly system behaviour and adaptation are analogous in many respects to the coupled systems of the dynamic perspective.

Although there are areas of convergence between situated and dynamic perspectives, the relationship of situated theories to the dynamic perspective is rarely made explicit[12]. The reader's attention should be drawn, however, to one critical convergence: an important implication of the constructivist emphasis of situated stances is that no end-point to the developmental trajectory is envisaged. Where behaviour-based stances lacked the capacity for non-trivial adaptation, and traditional cognitivism sought to

[12] Although see Clark & Chalmers (1998).

characterise development towards the ideal logico-mathematical end-point, the situated perspective allows for open-ended growth. Whereas the majority of implementations from both symbol-based and behaviour-based approaches are static, situated and dynamic perspectives advocate the progressive reconstruction of systems.

2.2.3 Summary

The situated perspective on cognition emphasises the coordination of perception and action dually constrained by physical embodiment and environment. Adaptivity occurs through coordination in the absence of explicit representation for some theorists. Others regard the function of representation as the provision of contextual information rather than enablement of action through logical manipulation of symbols. It is hoped that the meaning of representations for such systems will be grounded in direct experience of the world and that consequently the frame problem of classical AI can be avoided.

The major implications for the design of artificial systems are that structural and niche constraints should be utilised to coordinate perception and action. Representation, if present, should be restricted to non-manipulable, contextual, agent-centred information.

2.3 Situated and emergent artificial systems

2.3.1 Situated robotics

Situated robotics seeks to retain the reactivity of purely behaviour-based systems whilst reintroducing elements of representation and internal control. Representation becomes, once more, an element of robotic architectures yet the symbolicism of classical AI is replaced by a pragmatism which emphasises egocentric, contextual representation. As characterised herein, 'situated robotics' encompasses situated agent, and some animat approaches. In order to give a flavour of the nature of these systems several will now be briefly described.

PENGI (Agre & Chapman, 1987; Chapman & Agre, 1987) although not a robotic implementation is important for popularising a situated approach within AI. Agre (1988) had come to the conclusion that everyday activity is more accurately characterised as routine, rather than newly deliberated. He maintained that many goal-directed behaviours are improvised on the spot based on current environmental features. PENGI is an implementation inspired by this view of activity. PENGO is a simulated computer game played on a two-dimensional maze of ice-blocks. The player navigates a penguin through this environment which dies if it is hit by either an ice block or a bee. PENGI is a control system whose action derives from interactive routines based on 'indexical-functional', rather than global symbolic, representation. The behaviour of the system, therefore, is contextually-informed — situation specific.

PENGI demonstrated, albeit within a simplified and simulated domain, that representation needs neither to be global nor objective to support adaptive behaviour but can be restricted to primitives which relate perception of the world to activity.

AARON (Cohen, 1988; McCormick, 1991) again, is not a robotic implementation but a drawing program capable of producing an infinite number of different pictures of garden scenes. The system draws on-line; its behaviour is not planned in advance but is rather generated through matching current world state with a desired end state. Representations are indexical-functional, reflecting the immediate dynamics of the drawing process (perception and action) rather than objective states of the world. The system's ontology characterises the relationship of system states and perceived world states.

Although the system is categorically bound (its performance is constrained by a pre-determined ontology of trees, and a small number of people) it is again important for demonstrating effective on-line control based on non-symbolic representation and improvised action.

Situated automata The 'situated automata' of Rosenschein & Kaelbling (1986, 1990) represent an engineering approach to robot construction. They possess no symbolic knowledge of the world, rather contextual 'knowledge' is pre-programmed into the agent by the designer. Agents are specified in terms of perception and action components by two programs. RULER specifies the perception component based on the semantics of the agent's inputs, a set of static facts, and the state transitions possible within the world. The designer specifies the semantics for the output of the agent, and then the compiler synthesises a circuit whose output will have the desired semantics. Declarative knowledge is therefore reduced to a pre-programmed circuit (Kaelbling, 1991). GAPPS accepts a set of goal reduction rules (which encode information about how goals can be achieved) and a top level goal, and then generates a program that can be translated into a digital circuit to achieve the goal.

For these systems knowledge is not incorporated as data structures but is replaced by a description of how the state of the machine should relate to the state of the world. The generated circuit neither represents nor manipulates symbols as all knowledge is installed at compilation time. Knowledge becomes a theoretical construct of the designer used to construct a circuit with the interactive features required. Although the systems which derive from this approach are capable of contextually appropriate action, the ontology of the system is that of the designer.

Toto (Matarić, 1990; Matarić & Brooks, 1990) is a robotic 'dog' that learns about the relative location of landmarks and constructs a map of its environment. The map is not globally available but is accessible only in the context of moving through the environment. This implementation, therefore, dynamically constructs a map of the environment through action.

However the system possesses predefined categories for modelling the world (walls, obstacles, *etc.*) which are stored as landmark descriptions. Learning of novel landmarks occurs through comparison of the current landmark to the stored descriptions of known landmarks. The matching process uses a predefined calculus for manipulating the representation, as in classical systems. This calculus, for example, represents the equivalence of a right wall when heading north, and a left wall when heading south (Matarić & Brooks, 1990).

Situated robotics is based on an alternative to logical AI, procedural or functional semantics (McDermott, 1987; Birnbaum, 1991). These implementations eschew the objective, global, symbolic representations of classical AI, replacing them with subjective descriptions of the inter-relationship of perception and action. Whilst striving to retain the reactivity of BB systems, situated systems hope to add an element of contextual relevancy to behaviour. Memory becomes less important to the system, being replaced by a 'knowledge' of how to do the right thing at the right time. The designers of these systems seek to move away from the pre-installed competences and representations of classical AI.

Situated robotic implementations lag far behind situated theories. This is unsurprising as situated stances provide guidelines rather than principles for design. Although one strong feature of the situated approach is the grounding of the meaning of representations through action, implementations are overly preconstrained by their designer's ontology: meaning is still installed rather than derived. Furthermore a number of important design questions remain unanswered: should systems be provided with ontological primitives or 'grammars' for representation construction; what should be the role of goals and internal control; how should 'coordination' be achieved[13]; how can learning be extended from 'filling-in' pre-installed categories to the construction of novel categories; should system structure change over time through reorganisation; how should scaling[14] of systems be achieved?

Summary

The field of 'situated robotics' is disparate and still young. Compared with the plethora of theorising from situated stances implementations are sparse. Part of the reason for this disparity is that the major emphases of situated theories, the contextual relevancy of behaviour grounded in subjective, agent-centred representations and coordination of perception and action, are not easily translated into clear prescriptions for design. The lack of clarity of the situated perspective on the correct role, if any, for representations only serves to aggravate the problem. Situated robotics, however, does demonstrate that augmentation of purely reactive behaviour-based systems by local, agent-centred representation can support more contextually relevant adaptive behaviour without risk-

[13] Clancey (1995) has recently attacked situated robotic implementations for using stored representations rather than reconstruction of coordination processes. In his view these should involve reactivation of neural networks that coordinate sensing and acting.

[14] As with BB robotics, the aim of the approach is to work incrementally from simple to complex artificial systems: to recapitulate evolution.

ing the combinatorial explosions characteristic of classical AI.

However, many problems remain. How can such systems move from being generally reactive towards internal control? Doing the right thing at the right time involves *internal* as well as environmental criteria and the assessment of whether what has been done is, in fact, right, requires *interpretation* of action in context. So far, situated robotics has not begun to address this issue. Furthermore, systems should be able to construct their own ontological categories through experience, and undergo internal reorganisation throughout the life span. If the goal of artificial intelligence is the construction of *complex* artificial systems an adequate, and *prescriptive*, characterisation of the biological targets is required.

We turn now to emergent approaches which strive to embody the *a*representational coordination of perception and action suggested by some situated theorists, and the internal reorganisational features of dynamic stances.

2.3.2 'Emergent' computation

The 'emergent' approaches — connectionism, artificial life and their conjunction, evolutionary robotics — share an interest in developmental questions, stress system-environment interaction, and claim some biological plausibility. Although elements lacking in either symbol-based or behaviour-based approaches, for example the capacity for change in system structure over time, are present, it is suggested that none of these stances can, as currently instantiated, support the development of artificial intelligent systems.

Connectionism

The connectionist approach originated with the analogy between the all-or-none rule of neuronal firing and the binary logical units of computers which led to the description (McCulloch & Pitts, 1943) of a hypothetical artificial 'neural network' (henceforth ANN). More recently coherent criticisms of the capacity of purely symbolic systems to support complex (human-like) forms of cognition (Dreyfus, 1972; Winograd & Flores, 1986; Clancey, 1987; Rommetveit, 1987, for example) have led to explicit opposition to the dominant symbolic approach exemplified by Newell & Simon (1976) and a motivation to demonstrate intelligent behaviour arising within a *situated* system with a *distributed* system of representation.

Connectionism is often depicted by its advocates as a significant advance over the traditional symbolic stance within AI. ANNs can be viewed analytically for their powers of abstraction and regarded solely as devices for pattern extraction and classification (McCulloch & Pitts, 1943; Minsky & Papert, 1969; Marr, 1982) and also, more significantly, from a synthetic standpoint — as capable of mimicking some essential qualities of human cognition such as generalisation, ambiguity tolerance, graceful degradation, and content addressable memory (Smolensky, 1988; Clark, 1989) as well as providing a more plausible biologically-rooted model of intelligence.

Connectionism moves the focus of interest within AI away from symbol systems and heuristic search methods towards issues of adaptation. Unlike traditional forms of architecture within AI, ANNs do not construct a more or less detailed symbolic model of the world but rather adapt directly to the world through experience (Luger & Stubblefield, 1998*b*).

Implementation

A traditional neural network consists of a number of interlinked units ('neurones'). The connectivity of the network is generally instantiated randomly; the 'weights' of excitatory or inhibitory links between units are modified in the light of experience in accordance with a number of different learning functions. A unit becomes active ('fires') if the sum of its weights exceeds its activation threshold.

Traditional neural network architectures possess a number of layers: an input layer, often a series of binary inputs to a number of units; one or more 'hidden' layers; and an output layer, often directly connected to actuators, with binary output. The performance of a standard neural network clearly depends upon two factors: (1) the properties of the individual units within the network; and (2) the overall architecture — connectivity, whether a given connection is inhibitory or excitatory, and choice of learning algorithm. More recently, ANNs have been developed which constitute an explicit attempt to incorporate ideas from dynamic systems theories (Beer & Gallagher, 1992; Cliff *et al.*, 1993*a*; Harvey *et al.*, 1993; Yamauchi & Beer, 1993), with the assumption that the dynamic profiles of *biological* neural networks might be essential in the explanation of cognitive phenomena. These *dynamical neural networks* possess features such as time-delays on connections between units, recurrent and directionally unrestricted connectivity, non-uniform activation functions, and deliberately introduced noise. These kinds of networks are consequently capable of much richer dynamics than those produced by traditional connectionist systems.

Networks are exposed to a training set of data (or experience in the world in the case of robotic implementations) and produce an output. An error function determines how weights should be altered and the process is repeated. Over the course of training the network converges on its 'target' output. Typical learning algorithms are back-propagation (Rumelhart *et al.*, 1986), reinforcement learning (Barto *et al.*, 1995), classifier systems (Booker *et al.*, 1989), and self-organised maps (Kohonen, 1982). The algorithm chosen imposes constraints on the nature of the architecture and the type of supervision provided by the designer.

Analysis

The distributed form of representation peculiar to connectionist systems has been argued (Harnad, 1990*b*) to be one solution to the symbol grounding problem — through grounding meaning in direct experience of the world. But is there a need for grounding within connectionist architectures? ANN systems are capable of only *implicit* meaning — the network *is* the system. Without modularising the system how can meaning be made explicit? Incorporation of ANNs within a larger, hybridised architecture is one solution. This would allow the implicit meaning lurking within the network to be made explicit to the system at large.

A further consequence of this holistic structure is that, as for many classical symbolic architectures, experiential domains cannot be segmented — systems are *a*contextual. In the absence of explicit, contextual knowledge such systems lack the capacity for both hierarchical and serial control. Although the structure of the network derives from past experience, the behaviour of a network at a given time depends purely on current environmental features. Such systems should be characterised, therefore, as reactive — behaviour is driven by changing external stimulation.

A further difficulty inherent within the connectionist approach is that the lack of system primitives to constrain growth trajectories means that novel behaviours cannot be predicted in advance, restricting the explanatory role of connectionist models to *a posteriori* modelling of behaviour (McGonigle & Chalmers, 1998*a*). Furthermore although connectionist theories aspire towards extendibility, the problem of the *origin* of high-level (epistemic) competences remains (Fodor & Pylyshyn, 1988; Clark, 1989). McGonigle & Chalmers (1996) argue that this is because ANN models reflect a phylogenetically older and more simple form of neural inductive machinery which does not scale up to more complex epistemic forms of adaptation that appear to rely on *explicit* knowledge.

Summary

Connectionism constitutes an advance over traditional stances in its emphasis on change in system structure over time through developmental and interactional processes. This frees the designer from exact specification of the system. However as currently instantiated ANNs cannot support the kinds of explicit, contextual representation which underly adaptive internal control and their behaviour, therefore, is determined by local environmental contingencies — they are *reactive*. The future of connectionist architectures is within hybrid systems where, as encapsulated modules, their implicit representations can come to *mean* for the remainder of the system.

Artificial life

The philosophical background of ALIFE lies in evolutionary epistemology (Popper, 1959; Campbell, 1974). Popper (1959) suggested that systems incorporate both explicit and implicit knowledge and should be regarded as conjectures to be falsified against their adaptive domains. The mechanism by which systems become progressively more adaptive was, for Popper, neo-Darwinian evolution. So the evolution of the adaptive knowledge embedded within a system occurs through a process of trial (mutation) and error elimination (selection) — *via* a blind process of mutation and selection systems become progressively more adaptive over time. ALIFE seeks to simulate this process with the eventual goal of the emulation of complex competences.

von Neumann (1958) was the first to capture the essence of ALIFE in his attempt to abstract the logical form of self-reproduction and subsequently demonstrate artificial self-reproduction in cellular automata. Early ALIFE research which attempted to empirically investigate the functioning of random mutation and selection within simulated artificial 'organisms' (Barricelli, 1962, 1979) was ignored until very recently (Dyson, 1998). More recently still genetic programming techniques have led to a vast area of study with the explicit motivation of synthesising life ('strong ALIFE') within digital computers (Langton, 1987, 1995) with the further aims of the construction of a general theoretical biology, and the application of evolutionary methods to the development of artificial intelligent systems.

The dominant methodology is based on the genetic programming algorithm (Holland, 1975; Koza, 1990). Initially a fitness function is determined. A random population of artificial programs ('genes') is generated which code for 'phenotype' (for example, sensor-actuator wiring). Next, iteration through program execution, assignment of fitness value, generation of a new population and mutation leads to the target solution, or termination after a pre-determined number of iterations. The implicit focus of ALIFE is behavioural adaptation (Langton, 1995) as can be seen in the resulting systems which are currently at a primitive stage; complex competences remain on the distant horizon.

ALIFE is often characterised as avoiding the problems of preinterpretation inherent within symbol-based and behaviour-based stances. Prior knowledge is assumed not to be present in genetic models although a critical implicit assumption, as with all of western science, is the presence of environmental regularity (Christianini, 1995). A more serious criticism concerns the design of fitness functions. The particular function chosen determines which individuals will survive and evolve and consequently controls the direction of the evolutionary process. At present there are few guidelines for the choice of such functions which reflect the chosen target of the evolutionary process and, therefore, the designer's desires.

Of course, synthesising artificial intelligent systems through emulation of neo-Darwinian processes might be successful only if such processes are the driving force behind increased system complexity over time. This is a great area of controversy within contemporary evolutionary biology. The difficulties[15] of explaining the general trend towards increasingly complexity in biological systems through neo-Darwinistic processes alone (Barricelli, 1979; Goodwin, 1994, for examples) have led more recently to suggestions for a number of supplementary mechanisms such as co-evolution[16], and co-adaptation and symbiogenesis[17]. Furthermore, one consequence of adopting a dynamic systems perspective is that evolutionary mechanisms might merely serve to increase the richness of the ontological lower bounds of the system with the critical force behind system complexity being self-organisational processes over ontogenesis. Suggested mechanisms include conflicting constraints[18], meta-system transitions[19], and combinatorial expansion and generative condensation[20].

It is clearly critical that the trend towards complexity over phylogenesis is adequately characterised for biological systems before attempts at its replication as a means of developing complex artificial systems. Currently, with the exception of a small number of studies (Cliff & Miller, 1996, for example), ALIFE has focussed on neo-Darwinistic processes in isolation. The emphasis of the dynamic perspective on inherent self-organisational processes as a critical force underlying progressive system development over both phylogeny and ontogeny has yet to be influential within ALIFE. This must be rectified if it is to make progress in its task.

Summary

[15] Although see Bonner (1988) for a denial of the existence of any such difficulties.

[16] The unpredictable results of interactions between systems, both within and across species, together with the effects of the actions of systems on both their own, and others' niches has been suggested to imply that any increase of complexity in one species (through neo-Darwinistic processes) will tend to increase the complexity of the niches of others who will be forced to evolve more complex control structures themselves in order to cope (Casti, 1979, for example). By repeated iterations of such a process, systems might 'bootstrap' each other towards complexity.

[17] Wimsatt (1972) has suggested that co-adaptation of internal mechanisms might lead to increased complexity. Symbiogenesis (originally proposed by Merezhovsky in 1909, and developed by Kozo-Polyansky in 1924) ascribes the complexity of cellular structure to a series of symbiotic associations between simpler systems. This idea was extended by Barricelli (1962) who formulated a theory of 'symbioorganisms': self-reproducing structures constructed through symbiotic association of several self-reproducing systems of any kind (Dyson, 1998).

[18] Kauffman (1993) suggested that self-organisation depends critically upon the complexity of conflicting constraints over lower-level interactions.

[19] Heylighen (1991) suggested that the complexity of organisms increases through the emergence of meta-levels.

[20] See section 2.1

ALIFE represents, along with some connectionist models, an important advance. Development and environmental interaction are both fundamental. In principle at least the designer is absolved from exact specification of design, and of end-point specification: different primitives can be combined in order to see what emerges. This allows for the development of unexpected (and arguably impossible to pre-design) adaptations. However, although it strives for dynamism, subscribing to an emergent view of development at both ontogenetic (interaction of units leads to more complex behaviour) and, of course, phylogenetic levels an exclusive focus on the behavioural adaptation of whole systems may prove to be a serious weakness. Furthermore, exclusive concentration on neo-Darwinian processes might be found to lack biological plausibility.

It might be that a switch in focus from population-based phylogenetic adaptation toward individual ontogenetic adaptation would be the most promising future direction for this stance. It is possible that, as argued by Mithen (1996), the evolution of independent domain specific modules is much more likely than a whole mind. Over millions of years such intercommunicating units could be constructed which could cooperate to perform increasingly more complex tasks[21]. Such an approach might also enable the incorporation of representational elements within systems thus allowing ALIFE to avoid many of the pitfalls of other behaviour-based approaches.

Evolutionary robotics

> "Insects first, people later." (Cliff, 1991*a*)

As we saw in section 2.3.1, robotics inspired by situated theories constitutes a progression from either behaviour-based or classical AI through its emphases on pragmatic representation and situated action. However, work in robotics remains focussed on systems whose processes exhibit stable organisation over time. Natural complex systems, in contrast, self-organise over time, becoming increasingly more complex and more adaptive. The major research area which seeks to embrace dynamicism within the system is evolutionary robotics. To what extent might this methodology enable the construction of artificial intelligent systems?

[21] This view of mind is reminiscent of that of Minsky (1986).

Evolutionary robotics seeks to apply artificial evolution to the control systems of autonomous mobile robots (Cliff *et al.*, 1993*b*; Harvey *et al.*, 1994; Nolfi *et al.*, 1994). The 'genotype' which is subject to evolution specifies the dynamical neural network controller (Kodjabachian & Meyer, 1998) and also sometimes physical structure such as sensor and actuator characteristics (Cliff & Miller, 1996) and robot morphology (Lund *et al.*, 1997). Initially a random population of genotypes is generated, a pre-determined fitness function selects the more successful controllers which then have a proportionally greater opportunity to contribute genetic material to subsequent generations. It is hoped that adaptive behaviour will develop through a dual process of selection of successful controllers over phylogeny, and by emergence over ontogeny through interactions both internal to the system, and between system and environment. The evolved controller, therefore, coordinates perception and action, without representation as demanded by situated cognition theorists such as Clancey (1995), to generate adaptive behaviour (Cliff, 1991*b*; Chiel & Beer, 1997).

The approach attempts to free the designer from specification of all system features except task and fitness function. It is hoped that the problems associated with explicit design of adaptive and appropriate coordination of behaviour will be overcome by adoption of evolutionary and self-organising processes (Urzelai *et al.*, 1998). The behaviour of the system, as characterised by an observer, reflects the dynamic interaction of the coupled agent-environment system. Complexity of behaviour does not imply corresponding complexity of internal control structures (see Simon 1962, Braitenburg 1984). Evolutionary robotics, it is argued (Nolfi, 1998, for example) frees the designer from both behavioural decomposition by relying on evaluation of the whole system and its global behaviour, and the problems inherent in designing appropriate system-environment interactions. The system, it is hoped, will consequently be untainted by the preinstalled ontology of its designer.

What kind of competences have been developed using this methodology? Some example implementations will now be briefly described.

- Harvey *et al.* (1994) evolved the control system for a robot which was trained to visually discriminate between triangular and rectangular targets. The successful controller solved the task using information from only two pixels of the camera. Presumably such a strategy was effective only through sensory-motor coordination.

- Nolfi (1997) evolved the control system for a Khepera mobile robot with the task being to discriminate walls and cylinders in the environment. It was found that all evolved individuals solved the task by moving back and forth in front of a perceived object. This was interpreted as problem solution through sensory-motor coupling.

- Nolfi & Miglino (In press) evolved robots capable of reaching the upper right hand and bottom left hand corner of a rectangular box starting from eight different positions[22]. The evolved robots solved the task by differing the speed of the two wheels which ensured that long and short walls were encountered at different angles allowing the robot to reach the corners by either following or avoiding a wall dependent on angle of incidence. Again sensory-motor coupling appeared to enable solution of the task.

- Yamauchi & Beer (1993) showed that evolved dynamical neural networks are capable of successfully performing relatively abstract tasks involving sequential behaviour. Their system generated a fixed sequence of outputs in response to a series of external triggers.

Design

Evolutionary roboticists claim to have removed themselves almost entirely from the design process. As few restrictions as possible are placed on the structure of a given network; the designer simply computes a fitness function and retires. In practice, however, in the vast majority of cases the designer determines the genotype to phenotype mapping, the number of individuals in a population, mutation and crossover rates and, in some cases, even the architecture of the controller. It has been argued (Nolfi, 1998, for example) that since all these parameters might, in theory, be subject to the evolutionary process it is sufficient merely to modulate the selection criteria. Indeed, by only modifying these criteria whilst holding all other parameters constant varying forms of

[22] Modelled on an experimental task of Gallistel (1990) with rats.

behaviour have been evolved from the same base system: obstacle avoidance (Floreano & Mondada, 1994), exploration (Lund, 1996), navigation (Nolfi & Miglino, In press), discrimination of different objects (Nolfi, 1997) and predator avoidance (Nolfi & Floreano, In press). All these competences are, however, relatively simple. We shall see in the next section that in the attempt to construct systems capable of more complex competences the designer often becomes *more* rather than less involved.

Scaling

Researchers within evolutionary robotics stress that through system-environment interaction agents evolve which exhibit adaptive behaviour. All these competences, however, are both simple and under direct control of the environment. Even the most advanced (Yamauchi & Beer, 1993) which demonstrates sequential behaviour, does so only on the provision of appropriate environmental stimulation. How are such systems to be scaled up to exhibit more complex behaviours?

Evolutionary roboticists (see Nolfi, 1998, for example) envisage three paths towards increased complexity. The first is *incremental* development over phylogeny. Harvey *et al.* (1997) suggest that selection criteria should be progressively modified by the designer to impose increasing task requirements on the system and consequently, it is hoped, produce more complex competences[23]. However, this solution demands the supervision of the designer and risks reintroducing inappropriate constraints on the solutions to tasks — something that this approach prides itself in avoiding.

A second solution is based on *co-evolution* (Casti, 1979; Dawkins & Krebs, 1979) which leads, through self-organisation, to an evolutionary 'arms race': populations bootstrapping each other towards complexity. Preliminary investigations (Nolfi & Floreano, In press) have shown that co-evolution of populations can generate solutions impossible to evolve with one population. The task studied, however, was again simple: a predator and prey scenario.

Evolutionary robotics, though, unlike ALIFE as a whole, also emphasises self-organisation over the life-span as a critical driving force behind increased complexity. Where phy-

[23] This concept is based on shaping (Skinner, 1938), and was applied to classifier systems by Dorigo & Colombetti (1994, 1997). It is also analogous to the paradigm employed by McGonigle & Chalmers (1998*a*) to motivate increasingly complex behaviours over ontogeny in primate.

logenetic scaling relies on selection of the whole system, once, according to relatively gross criteria, ontogenetic development can, in theory, take advantage of the rich information available in the environment to provide feedback on the effectiveness of behaviour leading ultimately to increased adaptiveness.

Preliminary studies (Ackley & Littman, 1991; Nolfi & Parisi, 1997, for example) used a fixed architecture split into two sub-networks. The first determined how to respond to a given sensory state; the second generated a teaching signal for the first. The weights of the two sub-networks were subjected to an evolutionary process. Results seemed to indicate that systems were able to learn over ontogenesis by translating sensory information into useful teaching signals. Similar results were found for networks with evolving topologies but which learn over ontogenesis through unsupervised learning (Miller & Todd, 1990) and architectures in which the evolutionary mechanism selects different learning rules for different connections (Floreano & Mondada, 1996).

Again however, the competences learned through such processes were limited to sensory-motor coordination. Complex competences were neither evolved nor learned. A serious problem with this approach for evolutionary robotics is that systems are incapable of determining for themselves using internal criteria the utility of their actions yet the information present in the environment only indirectly specifies how successful given actions are. Indeed at the present time no evidence has been presented to indicate that the addition of learning mechanisms to evolved robot control systems results in more complex competences than evolution alone (Nolfi, 1998).

A third approach is to attempt to mimic the developmental (morphogenetic) properties of biological systems. The implicit assumption appears to be a view that this might hold the key to developing more complex competences through increasing the interactional complexity within the system: richer system dynamics might provide the scope for increased self-organisation leading to more adaptive systems. Natural biological systems are *evolvable* — random variations can sometimes produce improvement. Evolvability depends on the mapping of the variation within the genotype to that in the phenotype (Wagner & Altenberg, 1996). The most simple mapping is one-to-one where one gene codes for one feature of the system. Evolutionary robotics includes: one-many mappings (Floreano & Mondada, 1994, for example); recursive instructions for growth which are applied to their own results, analogous to cell duplication and differentiation (Kodjabachian & Meyer, 1998); and genotypes which can vary in length (Harvey, 1992). These attempts to more accurately model the workings of biological evolution are still in their infancy but little indication that more accurate modelling of biological evolution in this way leads to the development of more complex systems has been reported.

Currently the competences evolved within this paradigm remain low-level. Maybe this is partly to do with the time involved in conducting such experiments on real robots (Matarić & Cliff, 1996) and, indeed, Nolfi (1998) suggests that ultimately this might prevent the paradigm ever scaling up to more complex behaviours. Notwithstanding this pragmatic difficulty more serious problems exist. As with ALIFE it remains unclear how to select effective evaluative criteria for the *whole* system, and whether progressive system development can be achieved through evolutionary mechanisms alone. A tacit recognition of this problem has inspired hybrid systems designed to also allow ontogenetic development but currently the competences learned by such systems remain low-level. Furthermore, the reliance of evolutionary robotics on dynamic neural network architectures renders it susceptible to the same criticism levelled against connectionism: lack of explicit, contextual knowledge makes the system reliant on external triggers rather than internal control.

Summary

Evolutionary robotics strives to develop adaptive artificial systems through evolutionary and ontogenetic, neural network based, self-organisational mechanisms. As such it shares the benefits of connectionism and ALIFE: a focus on developmental and interactive processes. However, it is also subject to the same criticisms: excessive reactivity and an inability to scale.

Summary of emergent approaches

Connectionism, ALIFE, and their conjunction evolutionary robotics, all share an important emphasis on system reorganisation over time through processes of environmental interaction which was lacking within both classical and behaviour-based AI. Connectionist systems adapt directly to the world through experience and may indeed be one method of resolving the symbol-grounding problem as Harnad (1990*b*, 1995) suggests. However such systems lack explicit, contextual knowledge and therefore the means for adaptive internal control. Within ALIFE the focus on the evolution of reactive behaviour neglects the importance of control exerted from within the system. Furthermore, the ability of neo-Darwinistic processes, in isolation, to support the progressive adaptation of whole systems, as opposed to domain specific sub-systems, is unclear. Evolutionary robotics combines the best and worse of both prior approaches yet despite initial optimism systems remain at a primitive level of adaptation.

Some of these difficulties, however, might be surmounted by incorporation within hybridised systems. The scenario might include evolutionary methods being used for behavioural adaption, search, and perceptuo-motor coupling; ANNs could ground knowledge within limited domains that could be made explicit to the remainder of the system.

2.4 Summary: two contenders

We saw in chapter one that the symbol-based stances attempted to characterise and replicate the competences of complex biological systems. Many problems arose however, due to the mischaracterisation of a logical, symbolic competence. The artificial systems inspired by this mischaracterisation were slow and unreactive. Behaviour-based approaches attempted to rectify this problem by focusing on tightly-coupled, 'emergent' reactive competences with the aim of eventually scaling-up to complexity. Serious problems arose: these emergent competences had to be engineered after appropriate and accurate decomposition; and attempts to scale systems were unsuccessful. Attempts to hybridise these two approaches provided little extra headway.

The emergent stances have become popular as possible solutions to these problems. Maybe focusing on system-environment, and internal system interactions might provide a way out of the impasse. Systems became characterised as fundamentally coupled to their environments; representation was reduced to non-manipulable system-centred information or discarded altogether in favour of direct transactions between system and environment. Situated robots seemed to demonstrate that low-level adaptive behaviours could be achieved in the absence of global, symbolic, manipulable representations; their behaviour, however, remained highly reactive. Connectionist, alife, and evolutionary robotic approaches stressed the dynamics of inter-system and system-environment interaction leading to self-organisation of competences over time. These approaches, though, have currently supported only relatively simple competences and are finding it difficult to scale to more complex abilities.

It seems that in the desperation to escape from the lumbering systems of classical AI the headlong rush into a-representational, reactive systems, albeit well coordinated and coupled to adaptive niches, has served primarily to lead us away from both characterising, through synthesising, the complex intelligent systems that are, or should be, the research targets of psychology and artificial intelligence. As we shall see in the next chapter such systems *are* well-coupled to their environments and *do* self-organise over the life-span. they are also, however, capable of high-levels of internal control and not merely reactive.

The next chapter describes a research programme which strives to characterise intelligent biological systems as they self-organise towards ever-more complex competences over the life span. This analysis reveals a number of important properties of such systems which, it is argued, *must* be incorporated within any artificial system with aspirations towards complexity.

Chapter 3

The synthetic stance: characterisation and engineering

We have seen that both symbol-based and behaviour-based characterisations have failed to deliver intelligent artificial systems, and that neither dynamic nor situated perspectives provide clear design principles for artificial systems, and have consequently failed to resolve the impasse. It appears that in the attempt to develop artificial systems which are robust, reactive and behaviourally adaptive we have lost sight of the complex biological systems which should be our experimental targets.

If we are to construct artificial intelligent systems characterisations of the diversity of natural systems are required. Where should we look for such biological characterisations? This chapter briefly examines comparative research. Here, again, an impasse has been reached between symbol-based and behaviour-based characterisations which suggests that new approaches towards fractionating biological systems are required.

Next, a characterisation of complex biological systems informed by situated and dynamic perspectives is described. Central to the characterisation is the notion of the self-organisation of competences by agents on open-ended growth trajectories toward maximally economic information-handling strategies. Although the behaviour of these systems can be well characterised in the absence of logical or symbolic competences, their behaviour is not reactive but is rather generated by *rational* principles.

The characterisation suggests a number of critical design features for artificial systems. Following this description a number of implementations inspired by the biological characterisation, precursors of the architecture presented in chapter 4, are reviewed.

3.1 Biological characterisation

3.1.1 Characterising adaptive behaviour and fractionating biological systems

Reverse engineering of intelligence manifestly requires good characterisation of our biological targets. Comparative research should lead to descriptions of the diversity of biological systems allowing us to fractionate their differing capabilities and potential: which design features distinguish systems which remain at purely reactive behavioural levels of adaptation over the lifespan and those capable of epistemic growth over ontogenesis? The past century of comparative research, however, has failed to adequately fractionate systems. Why?

An impasse

Comparative study of biological systems grew from the evolutionary approach of Spencer (1855) and Darwin (1859, 1871) who both maintained that the search for the origins of human intelligence must involve attempting to find its precursors and analogues within the animal kingdom. Early researchers (Romanes, 1882; Yerkes, 1916, for example) assumed a continuity of adaptation from the lowest organisms to humans. A range of species were observed to exhibit intelligent behaviour but such characterisations were anecdotal and attributive; few paradigms existed within which to experimentally assay the capacities of species.

Criticisms of the method of early researchers (Morgan, 1894; Thorndike, 1898, for example) combined with a desire for repeated observation and objective methods of experimentation set the scene for the rise of behaviourism. We have already seen (section 1.3.1) that, for the behaviourists, tightly-coupled reflexes, trial and error learning, and operant conditioning became the indexical system competences. Associative learning was thought to underly adaptation in all species, in all contexts (Hull, 1945). Behaviourism culminated in the assertion of Skinner (1957) that even the ontogenesis of human language, the epitome of intelligent behaviour, could be explained through behaviour shaping, and that 'verbal behaviour' constituted the only significant difference between humans and other animals.

We have also seen (section 1.1.1) that cognitivism arose motivated by the inability of behaviourism to characterise high-level human competence — especially language. Few psychologists accepted Skinner's (1957) analysis. Surely the richness of the human mind could not have arisen entirely through incremental additions of primitive reflexes? Cognitivism, instead, inspired by the failure of behaviourism to satisfactorily account for linguistic competence, adopted language as the indexical system competence and a key explanatory concept. Humans were qualitatively different from other animals: the logical manipulation of language-like stored representations underlay the complexities of human thought and behaviour.

Threatened by this cognitive revolution, behaviourism retreated and became 'animal learning' still essentially behaviourist in methodology and outlook (see Macphail, 1982, and Mackintosh, 1983, for example) but making no claim to adequately characterise human competence. Although the possession of language distinguished humans from other animals, no real differences pertained between these others (Warren, 1965; Macphail, 1982, for example). The indexical system competence remained a highly optimised (Gallistel, 1995) associationistic learning mechanism, its operation dependent on the tight spatio-temporal contiguity of events (McGonigle & Chalmers, 1996) — despite early criticisms (Lashley, 1929; Maier & Schneirla, 1935) that such a mechanism, in isolation, could not underly the variety of adaptive behaviours. Focus on this mechanism alone, rooted in phylogenetically more ancient areas of the nervous system (Mishkin & Petri, 1984), meant that the range of inductive mechanisms utilised by systems were left ignored. Furthermore, the abstraction of system behaviour to unnaturalistic laboratory settings meant that the contextual operation and predictive qualities of reinforcing stimuli (Rescorla, 1971) have also been relatively under-examined. So, as mentioned previously (section 1.3.1), the narrowing of the inductive space characteristic of behaviourist paradigms, and their focus on tightly-coupled, easily modifiable behaviours abstracted away from their ecological setting, meant that they were incapable, in principle, of capturing the richness of situated adaptive behaviour and consequently of fractionating the varying capacities of different species.

So currently we have an impasse: on one hand symbol grinding humans imbued with the capacity for thought and language, and on the other the levelling of the remainder of the animal kingdom, differing only quantitatively in the speed of, and potential for, association formation. And between these two 'an unbridgeable gap' (McGonigle & Chalmers, 1996). Within the vast majority of contemporary psychology the impasse remains: the focus of cognitivism is abstract, disembodied, ruleful manipulation of symbolic representations, and that of animal learning the associative principles of behavioural modification. Cognitivism[1] ignores the origins of cognition; animal learning paradigms cannot fractionate species, nor explain the ontogenesis of intelligence.

3.1.2 The view from biology

How are we to adequately characterise the adaptiveness of biological systems; and where should we look for the phylogenetic and ontogenetic origins of adaptive intelligent behaviour? Cognitivism rarely motivates such inquiry yet behaviourism is incapable of characterising the richness of adaptive behaviour, and consequently cannot fractionate non-human species, nor explain the ontogenesis of cognition.

The alternative comparative paradigm, ethology, founded on the observation of species within their natural habitats, had identified both commonalities and differences among species. Biological systems were observed to differ both in their endowment of instincts (Tinbergen, 1951, for example) and their capacity for learning (Tinbergen, 1951, 1960; Bateson, 1990). Associationistic mechanisms were observed to form part of the complement of the adaptive machinery of systems yet even these, when operating in their ecological setting, were non-arbitrary[2] — dependent on the existing behavioural repertoire of a system (Lorenz, 1935; Bateson & Klopfer, 1982; Hinde, 1982), and influenced by both internal and external contextual factors (Hinde & Stevenson-Hinde, 1973). Systems seemed to be in possession of a range of learning mechanisms, optimised for a variety of contexts, and in operation at differing parts of their life-cycle (Tinbergen, 1952). Adaptive behaviour was seen to be inherently sequential, from locomotor behaviour to courtship, nest building and ultimately language. Instinctive behaviours

[1] With, of course, the notable exception of Piaget (1971).

[2] Behaviourist research was predominantly concerned with the learning of arbitrary relations between stimuli and reinforcers (McGonigle & Chalmers, 1996).

seemed to be 'triggered' by both internal and external 'releasers', and varied in their susceptibility to modification through learning. Yet although ethologists had concluded that systems avail of a multitude of learning mechanisms, recruited as and when needed, and characterised the richness of sequential adaptive behaviour, their lack of both clearly prescriptive theories and rigorous experimental method meant that each of their findings could be met with post-hoc, counter arguments from the behaviourists.

More recently, cognitive neuroscience and cognitive neuroethology, situated within an evolutionary perspective and computationally informed, have once again begun to challenge the prevalent associationism of psychology and animal learning. Although neuroscience has been dominated by an incremental view of species differences based, essentially, on a 'bigger is better' characterisation of the evolution of the nervous system (Geschwind & Levitsky, 1968; Passingham, 1982, for example), this is beginning to give way before an onslaught of evidence which suggests that phylogeny results in neuronal reconstruction and specialisation (Williamson *et al.*, 1993; Gazzaniga *et al.*, 1998, for example). Furthermore, Lashley's (1951) focus on the serial order of behaviour, and its underlying neural substrate(s) is once again receiving more interest within the neurosciences (Berridge & Wishaw, 1992; Aldridge *et al.*, 1993; Colombo *et al.*, 1993, for example).

The origins of adaptive behaviour

Within the neurosciences, the incremental view of the phylogeny of intelligence is being replaced by an appreciation of the range of *adaptive specialisations* (Gazzaniga *et al.*, 1998) which underly species differences. Learning across the animal kingdom cannot be reduced to the operation of a small number of highly optimised associationistic principles, as maintained by the behaviourists. Rather, species possess a range of *problem specific learning mechanisms* (Marler, 1991), each computationally specialised for solving particular problems, and applicable in some contexts but not others (Gallistel, 1995).

We can now begin to fractionate biological systems. All animals are not equal; rather they possess different adaptive specialisations subserved by functionally distinct neuronal regions (Gazzaniga *et al.*, 1998). Qualitative differences between human and the remainder of the animal kingdom certainly exist; as they do between different members of that kingdom, reflecting their varying evolutionary histories. As Waddington (1969) maintained:

> "The main issue of evolution is how populations deal with unknown futures." (Waddington, 1969, p. 278)

The maintenance of life over successive generations is a problem of *induction*. The solution of the problem requires systems to be adaptive in the face of an uncertain world. Species differ because their different evolutionary histories present differing evidential bases from which to 'induce' the future. Unfolding of systems over morphogenesis is sensitive to environmental conditions (Dobzhansky, 1947) providing a little flexibility in the construction of the ontological lower bounds of the system. This lower bound consists of a set of reactive, system-preserving instincts, together with a complement of learning mechanisms[3] which, through epigenetic and intra-system interactions, allow the system to adapt to an unpredictable world.

The adaptive value of learning

Biological systems must deal with uncertain futures. To this end, evolution has provided them with both 'instinctive' system-preserving reflexes and a variety of inductive mechanisms. These inductive mechanisms range from those optimised for associating stimuli linked closely in space and time, to more abstract and, critically, productive mechanisms (Fodor & Pylyshyn, 1988).

A fundamental difference between simple and complex biological systems is that they possess *qualitatively different* inductive mechanisms (McGonigle & Chalmers, 1996). Whereas the ability of simple systems to manage the future is highly limited, complex systems possess inductive mechanisms whose operation does not merely support the recapitulation of acquired knowledge and competence, but is generative — allowing

[3] Plotkin (1988) speculates that approximately five percent of species possess adaptive learning mechanisms in addition to instincts.

generalisation of acquired skills and experiences to related phenomenon and into new domains of competence.

So we see that systems differ in respect of their lower bounds. These lower bounds reflect the differing evolutionary histories of species: they are 'designed' both to reflect future environmental conditions 'induced' by the evolutionary mechanism whilst making a provision for flexibility in the face of uncertainty. The outcome of these learning mechanisms, associative and non-associative, over ontogenesis critically depends on the experience of the individual system.

The seriality of adaptive behaviour

> "I have devoted so much time to [...] the problem of syntax [...] because the problems raised by the organization of language seem to me to be characteristic of almost all other cerebral activity [...] Not only speech, but all skilled acts seem to involve the same problems of serial ordering, even down to the temporal coordination [of] such a movement as reaching and grasping. Analysis of the nervous mechanisms underlying order in the more primitive acts may contribute ultimately to the solution even of the physiology of logic." (Lashley, 1951, p. 122)

The rich observational data provided by ethology clearly demonstrates that many, if not the majority of, serial behaviours are not dependent on moment to moment response to a serially ordered environment, nor blindly consequent upon previous actions, but rather depend upon organising principles by which systems are capable of *internal control* of behaviour. All systems with any degree of parallel perceptual input, or of commonly 'triggered' behaviours face the problem of response scheduling: how to determine which, of a range of possible actions, to immediately perform and which to hold in abeyance. Serial order is therefore not unique to language but rather many systems must be in possession of mechanisms which select and schedule responses activated in parallel.

What are the mechanisms underlying such serial control; how do systems syntactically recombine their behaviours to achieve their goals; what mechanisms underly selection and arbitration of different behaviours; and how are behavioural sequences learned? Despite the ubiquity of serial behaviours throughout the animal kingdom, the majority of research on serial competence has focused on human language and other cognitive processes such as memory. What kind of explanations have been provided for such serial competences?

Skinner's *'Verbal behaviour'* (1957) attempted to account for linguistic competence using standard behaviourist chaining mechanisms, where each element in a series of actions provides the excitation for the next. Although initial formulations were criticised by Lashley (1951) for being context-insensitive, current behaviourist-inspired research, including animal learning, remains entrenched in an associationistic view of serial behaviour, albeit with the addition of some contextual features to basic associative mechanisms (Wickelgren, 1969, for example)[4]. Despite these elaborations it has become increasingly clear that association chaining does not accurately characterise serial system performance, especially in cognitive domains (Rosenbaum, 1991; Henson *et al.*, 1996, for example). Furthermore, associationistic mechanisms do not address a number of fundamental features of systems' serial competence such as their ability for syntactic recombination of behaviours and their mechanisms of internal arbitration which require contextual interpretation of behavioural success and default.

Cognitivist accounts, on the other hand, tend to rely on symbolic primitives to explain serial, grammatical linguistic competence. Theoretical linguistics, primarily concerned with the internal representation of serial order ('competence') rather than its execution ('performance'), relies on primitives such as strings, and ordered sets, which are used to construct descriptions of grammars (Houghton & Hartley, 1995)[5]. Unfortunately, when symbolic primitives are abandoned, the problems of serial order once more become intractable (Houghton, 1990; Shallice *et al.*, 1995). Consequently cognitive science remains without a grounded theory of serial competence at a neuropsychological level. Cognitively inspired neuropsychological accounts, when advanced, focus on specialised linguistic circuitry (Chomsky, 1957, for example) which have no purchase on serial

[4] Which have more recently inspired neural network models of serial control (Jordan, 1986; Rumelhart *et al.*, 1986; Seidenberg & McClelland, 1989; Elman, 1990, for example).

[5] Analogously, classical AI utilises similar symbolic primitives in conjunction with techniques of recursive serial processing.

behaviour *per se*. The cognitive focus on serial competence as a distinctly linguistic phenomena together with its linguistically-biased experimental methodology makes it extremely difficult for cognitive accounts to disentangle ordering from language.

More recently parallel activation models, as initially proposed by Lashley (1951), have become popular once again. Lashley (1951) had considered a hierarchy of internal mechanisms capable of restructuring behaviour at multiple levels of abstraction. He suggested that the production of serial behaviour involves the parallel activation of a set of actions which together comprise some 'chunk' of possible actions. Responses are therefore internally activated before being externally triggered; a 'schema for action' serves to select which responses from a 'chunk' are produced at a given time — to serially order behaviour. A range of models, have been inspired by Lashley's views, including simple response competition with associated 'strengths' (Shallice, 1972) and activation gradients (Estes, 1972; Grossberg, 1978; Mackay, 1987); response competition under internal modulation such as competitive queueing (Glasspool, 1985; Houghton, 1990; Burgess & Hitch, 1992), and other parallel models (Houghton & Hartley, 1995, for a review). These models, however, remain primarily linguistic in focus.

We see that serial behaviours are ubiquitous — how can they be disentangled from language; is the ability to syntactically recombine action pre-linguistic and, if so, what are the precursors of 'grammar'; does seriality of behaviour rest on a unique ordering mechanism or do task-specific, specialised mechanisms underly different ordering competences as suggested by recent cognitive neuropsychological accounts (Gazzaniga *et al.*, 1998, for example); and how do systems select between competing responses?

Summary

In order to replicate the intelligent systems we see around us we need both:

- Accurate characterisations of adaptive behaviour in its ecological setting:
  - which competences are open to modification?
  - what are the effects of internal and external contexts?
  - how are behaviours sequenced?

- Accurate characterisations of differences in the ontological lower bounds of biological systems:
  - what kinds of reflexive, or 'instinctive' behaviours do systems possess; what are their functions?
  - what kind of inductive mechanisms do species possess?
- Accurate assays of the ontogenesis of biological systems:
  - in what internal and external contexts do inductive mechanisms operate?
  - how do species' inductive mechanisms interact with environments of differing complexity in order to scaffold the system?

We have seen that the behaviourist focus, now exemplified in contemporary animal learning paradigms, on an optimised inductive mechanism effective only for events closely linked in space and time has resulted in a relative neglect of complex adaptive behaviour. Associationistic mechanisms *alone* do not underlie progressive system adaptation over ontogenesis, but rather form only part of a complement of *adaptive specialisations*. The focus of cognitivism on linguistic competence tends to neglect the origins of intelligence and, relies on symbolic primitives, preinterpretation, or peculiarly human neurolinguistic machinery to account for such abilities. In-between we appear to have an 'unbridgeable gap' (McGonigle & Chalmers, 1996).

Turning to biology we find rich characterisations of the varieties of adaptive behaviour originating from ethological research. System behaviour is inherently ordered in space and time, 'grammatical' sequencing of productions exists throughout the animal kingdom and is not a peculiarly linguistic phenomenon. Systems differ in the range of inductive mechanisms at their disposal. Learning mechanisms are many and varied; computationally optimised for different contexts, internal and external. Cognitive neuroscience informs us that these *adaptive specialisations* are subserved by specialised neuronal circuitry: phylogeny primarily results in *qualitative* changes in system lower bounds rather than progressive incrementation of existing circuitry. These lower bounds constrain the developmental trajectories open to systems through both inter-system, and system-environment interactions.

Only when the questions above have been answered will we able to begin to address the fundamental issue: which differences in ontological lower bounds allow some systems to become increasingly more adaptive over ontogenesis, culminating for some in abstract knowledge, whilst others remain limited to the refinement of tight sensory-motor couplings.

The next section outlines an experimental paradigm, unique in both theoretical motivation and success, which strives to address the comparative and developmental questions outlined above. As a result this paradigm provides the potential for much richer characterisations of intelligent biological systems with which to inspire reverse engineering than both the mischaracterised logico-deductive end-state of traditional cognitivism and the one-mechanism associationism of behaviour-based approaches.

3.1.3 A new perspective on the origins of complex systems

The research programme of McGonigle and Chalmers at Edinburgh constitutes the most comprehensive current attempt to address the phylogenetic and ontogenetic origins of complex systems. The research, now in its third decade, conjoins developmental and comparative assays of species, from a dynamic perspective which stresses longitudinal observation of individual systems over ontogenesis. The experimental paradigms developed within this programme seek to assess the capacity of systems to self-regulate their behaviour, in essence scaffolding themselves, 'learning to learn', in the face of incrementation of task difficulty. Their findings allow us to fractionate different species, individual systems over ontogenesis, and to assess the generative benefits of non-associative, non-arbitrary inductive mechanisms.

Methodology and perspectives

Tasks and subjects

McGonigle & Chalmers' research programme centres on the relational abilities of systems. They speculate (McGonigle & Chalmers, 1996) that there may be at least two fundamental types of learning mechanism: (1) a phylogenetically more ancient associative inductive mechanism — such a mechanism is inductively weak as it requires events to be closely linked in both space and time and is unable to distinguish arbitrary from non-arbitrary connectives; and (2) a phylogenetically more recent relational mechanism that allows systems to perceive the non-arbitrary links between events, and enables generative strategies.

Initially research centred around reinterpretation of system performance on the standard five term transitive inference task (McGonigle & Chalmers, 1977; Chalmers & McGonigle, 1984; McGonigle & Chalmers, 1984, for example) but has since diversified to investigate the seriation and categorisation skills of systems (McGonigle & Chalmers, 1993; McGonigle *et al.*, 1994, for example). The results of these studies have prompted reviews of the fundamental operating assumptions, conclusions, and implications of traditional cognitive, developmental, and comparative research (McGonigle & Chalmers, 1996, 1998*a*, In press*b*, for example).

The systems assayed over the course of this research include human, brown capuchin monkey (*Cebus apella*), squirrel monkey (*Gen.* SAIMIRI), and pigeon (*Columba livia*).

A dynamic, life-historical approach

It was suggested above (section 2.1.3) that one of the most critical contributions of the dynamic perspective is the emphasis on the constraints imposed by the ontological lower bounds of a system on its potential ontogenetic trajectories. We have seen that part of the endowment of biological systems is a complement of inductive mechanisms, 'design primitives', which provide flexibility for an unpredictable future. This ontological lower bound ultimately constrains the kind of growth trajectories available to a system through inter-system and system-environment interactions.

The richness of this lower bound is a critical feature for fractionating different biological systems. William James (1981) speculated that humans possess *more* instincts than other animals and indeed this turns out to be close to the truth — as McGonigle & Chalmers (1996) suggest, the richness of that part of the ontological lower bounds of a system that is its complement of adaptive, inductive mechanisms determines its ontogenetic growth potential. The greater the number of inductive mechanisms a system possesses the greater its scope for development. This analysis enables us to relate the ontogenesis of systems to their lower bound condition and thus allows fractionation of systems at multiple levels of analysis.

This approach, then, necessitates a *life-historical*, longitudinal characterisation of systems. Furthermore, unlike traditional learning paradigms which demand only that systems learn arbitrary connections between events separated only by short spatio-temporal intervals, and take little account of the experiential history of a system, this paradigm strives to determine how the inductive mechanisms of a system interact over ontogenesis with tasks of increasing difficulty to result in progressive adaptation. The ancestry of this approach can be traced back to Harlow (1949) who had realised that performance even on classic behaviourist tasks was to some extent dependent on the life history of a subject. He speculated that trial and error learning might lead over ontogenesis to 'insightful' behaviour; his work on learning sets seemed to indicate progressive improvements in performance which led him to speculate that his subjects were, in fact, "learning how to learn" (Harlow, 1949, p. 53).

Ontogenesis involves interaction: both within a system and between that system and its environment. Growth is dually constrained by the richness of the lower bounds of a system, and the richness of the environment to which it is exposed. In order, therefore, to fractionate systems over ontogenesis a rich, and increasingly complex, environment must be provided. Only then will it be possible to accurately assay what inductive mechanisms 'buy' a system. McGonigle & Chalmers' paradigm achieves what Harlow (1949) could not manage: a hierarchy of tasks of progressive difficulty to which systems are exposed over ontogenesis and, therefore, provides a rigorous methodology for observing the generative (Fodor & Pylyshyn, 1988) benefits of adaptive mechanisms.

Self-regulation

We have seen that over ontogenesis interactions both internal to a system, and between system and environment underly growth. The research reported here depicts the development of systems *in the absence of selective reward* (McGonigle & Chalmers, 1998*a*). This is a critical point for it implies that the driving force behind growth for these systems is *internal*.

Presented with tasks that feature items to be ordered where some orderings are inherently more 'rational' than others, such as monotonic *versus* arbitrary colour sequences, both human and non-human primate systems appear to very rapidly avail themselves of these constraining features. The sole advantage of doing so is *economical*: by utilising ruleful constraints present within a data set a system achieves the "most behaviour for the least effort". The knowledge trajectory for epistemic agents (human and monkey) seems to be based on the degree of constraint present within a given decision space, rather than with the construction, and manipulation of ever more abstract levels of input (McGonigle & Chalmers, 1998*a*, p. 514). Evaluation of behaviour based on an *economy criterion* appears to constitute part of the ontological lower bounds of systems providing, as McGonigle & Chalmers (1998*a*) suggest, a means of surviving, and maximally exploiting, their formative learning experiences.

Evidence strongly suggests, therefore, that systems are internally regulating, using an inbuilt rather than externally imposed performance metric, towards maximally economic, generative, information handling strategies for dealing with ever-larger search spaces.

Rational constraints on development

> "The number of non-trivial solutions to the problem of complexity may be finite and may, indeed, reduce to one." (McGonigle, 1987)

It was mentioned above that systems appear to utilise whatever information-reducing constraints are present within a search space whenever possible. Such a fact reinforces one of the underlying operating assumptions of McGonigle & Chalmers' research programme: that the dynamics of (epistemic) growth possess an *inherent rationale.*

In common with a growing movement in favour of *rational analysis* of task constraints on the nature of their solution (Oaksford & Chater, 1998*b*, is a recent collection), McGonigle & Chalmers argue that analysis of system performance must be informed by appreciation of the informatic demands of tasks.

Findings and interpretations

A relational primitive?

Evidence that a perceptual relational ability might form part of the lower bound armament of non-human primate was discovered by McGonigle & Jones (1978). They found that performance on relational discrimination tasks was superior to that on those requiring absolute discrimination. Furthermore, this finding was robust in the the face of introduced contextual variation leading them to conclude that a perceptual relational ability was irreducible to any lower level of competence.

More recently McGonigle & Chalmers have argued for the primacy of relational skills on *a priori* grounds. A relational competence is logically, and must be ontogenetically, prior to absolute discrimination since:

> "Only by comparing and contrasting stimuli from a known set can the defining features logically and inductively be determined." (McGonigle & Chalmers, 1998*a*)

Relational primitives, therefore, can be seen as contextual learning devices which enable object definition in a given location by supporting active comparative and contrastive interrogation in order to extract the logically defining features of a set of objects (Garner, 1962; McGonigle & Jones, 1978; McGonigle & Chalmers, 1998*a*).

A relational competence, then, might be fundamental to intelligent systems. Against the grain of much psychological research over the past century (Spence, 1949; Inhelder & Piaget, 1964; Bryant & Trabasso, 1971, for example) which conflated relational with linguistic ability[6], evidence suggests that relational skills are pre-linguistic reflecting a rational *design primitive* of intelligent systems which supports later ordinal, serial, and classificatory competences (Tversky, 1969; McGonigle & Chalmers, 1998*a*). A relational competence appears to be a rational, adaptive mechanism utilised by systems in the context of certain problems. How much of seemingly 'logical' competence can be explained through operation of this primitive, and to what extent can multiple relational codes be derived?

Reinterpreting transitivity: rationality vs. *logic*

[6] Which, of course, begs the question of the origin of meaning for relational linguistic terms (Lashley, 1951; Bryant, 1974; Haugeland, 1985; McGonigle & Chalmers, 1998*a*).

The transitive inference task has been widely accepted as a measure of the understanding of necessity (Piaget, 1953, 1955) through the logical coordination of internal symbolic representations (Bryant & Trabasso, 1971, for example). How then are we to interpret evidence which suggests that four year old children (Lawrenson & Bryant, 1972), squirrel monkey (McGonigle & Chalmers, 1977) and even pigeon (Terrace & McGonigle, 1994) are capable of solving transitive inference problems?

Rather than ascribing a logical competence to such subject groups, McGonigle & Chalmers (1996) concluded that the ability to order items transitively must be prelogical and, in fact, reflects the operation of the relational primitive mentioned above. Solution of the standard 'transitive inference' task is possible through a strategy which generates rational, seemingly truly transitive choice, yet which does not depend on the logical coordination of premises. This conclusion is supported by evidence that the production system model developed by Harris & McGonigle (1994) developed, after a short learning curve, perfect 'transitivity'.

In addition to 'transitive' abilities, hierarchical competences traditionally indexed by class inclusion (Inhelder & Piaget, 1964) and use of quantifiers (Johnson-Laird, 1983) are generally interpreted as linguistically based reflecting a subject's understanding, through coordinative mechanisms, of the task demands. However evidence suggests that this competence, too, is based on the system's generation of economic, information management procedures (McGonigle & Chalmers, 1998*a*) which presuppose no logical abilities on the part of the agent whatsoever. Hierarchical structure underlies rational, efficient information storage and manipulation by helping to minimise the combinatorial problem (Newell *et al.*, 1958; Sokal, 1974).

The dominant characterisation of representation within psychology is, as we have seen (section 1.1.1) that of language-like internal symbols. Phenomena such as the symbolic distance effect (Potts, 1972) have been widely interpreted as evidence for a logically-coordinated seriated internal structure (Paivio, 1975; Trabasso *et al.*, 1975, for example). Contrary to such interpretations, McGonigle & Chalmers' (1984, 1994) findings of binary transitive choice in non-human primate, and of differing effects of presentation mode in human (1984) suggest that such competences might not be based in internal logical coordination mechanisms. Again it seems that a relational competence underlies representation construction, through integration of pairwise comparisons between items into an internal representation of series.

In fact, however, an even more elegant possibility suggests itself. Strong directional effects in performance on five term series problems indicate that the two end-items possess privileged status. It seems that systems do not construct a bidirectional manipulable representation of a known series but rather possess a ranking mechanism with unidirectional properties (McGonigle & Chalmers, 1994). Evidence to support this conclusion is again provided by the production system model of transitive choice developed by Harris & McGonigle (1994) which was capable of accounting for the co-existence of binary transitive choice and the symbolic distance effect in the absence of logical coordination mechanisms and, with only a small subset of rules (16 out of a possible 1920), accurately modelled both binary and triadic phases of the transitive inference data. Furthermore, all rule stacks that performed correctly with the initial training pairs extended to remote binary comparisons without the addition of further assumptions or expansion of the rule stack.

So transitive data from both human and non-human primate indicates that operation of a relational primitive is fundamental. Systems appear to utilise a unidirectional ranking mechanism in order to construct a representation of the serial order of a known set of items. What would happen when task complexity was increased towards explicit seriation, hierarchical search, and categorisation?

Linear seriation

The ten-item monotonic seriation task developed by Piaget and Szeminska (Inhelder & Piaget, 1964) is commonly used as an index of human cognitive growth. Instead of requiring a binary choice as in the transitive inference task, monotonic ordering of all ten items is demanded. This expansion of set size demands effective on-line serial control because of the massive combinatorial problem — insurmountable unless a system is capable of taking advantage of the inherent monotonic constraint within the set. How would performance differ between sets featuring arbitrary (e.g. colour) and non-arbitrary (classically, size) connectives?

Biological systems, both human and non-human primate, seem to automatically take advantage of constraints, when present in the data set. A large corpus of evidence has now been gathered which demonstrates that performance on tasks requiring non-arbitrary monotonic size ordering is far superior to that on those demanding non-monotonic size, or arbitrary connectives, as indexed by both acquisition measures and error profiles (McGonigle & Chalmers, 1992, 1993, 1996).

Faced with this combinatorial problem the system has two options: reliance on brute force memory strategies; or reduction of cognitive load. The relational primitive, once more, becomes critical. What appears to be happening is that the system utilises the monotonic constraint present in the data set — success is possible simply by iteration of a relational rule: take biggest, take biggest *etc.* Such a mechanism is maximally economic for data sets featuring this kind of monotonic constraint as it provides the potential for generation and prediction of successor items following seriation of a few antecedent elements. Such strategies are generalisable to sets of any size, eliminating the problem of combinatorial expansion and therefore serve a data-reducing purpose for resource limited biological agents (McGonigle & Chalmers, 1998*a*).

Hierarchical classification

Given that systems' performance in the face of non-arbitrary connectives far exceeds that when arbitrary connections are required, reflecting use of a relational, generative mechanism based on iteration of a simple rule to provide a path through a large search space, what would happen when task difficulty was once again incremented by introducing a further demand: hierarchical classification?

The experimental task required systems to exhaustively search increasingly large sets, with subsequent decomposition of the set into categories, within which items must be further ordered (see figure 3.1). The task provides a rich scenario within which to assess system competences as its allows conflation of both arbitrary and non-arbitrary connectives at different levels of the search hierarchy. Furthermore search sets can be extended in both breadth (addition of novel categories) and depth. Once more efficient and generative serial control is required. Would systems again exploit data-reducing constraints within the target data set in order to reduce cognitive load?

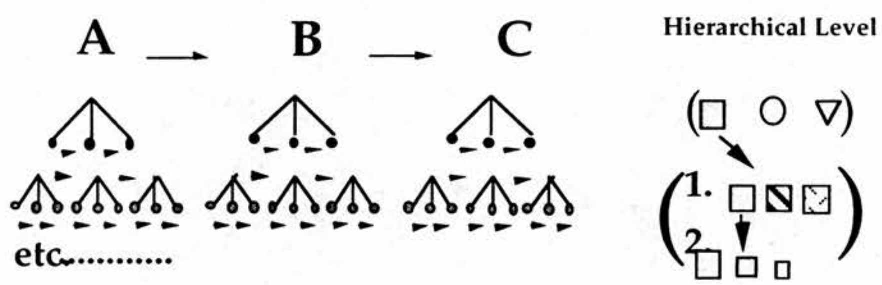

Figure 3.1: From seriation to categorisation (after McGonigle *et. al.*, in preparation)

Indeed results suggest that non-human primate, after a training period, is capable of ordered productions of up to twelve item sequences, in a two-level hierarchy using data-reducing classification and chunking strategies in conjunction with generative, relationally-based, mechanisms (McGonigle & Jaswal, 1993; Dickinson, 1997; McGonigle & Chalmers, 1998*a*). Figure 3.2 depicts data obtained from three adult *Cebus apella*. "Within classes" refers to a condition where the only constraint within the data set is monotonic variation in size; "between classes" refers to a condition where such monotonic size variation is augmented with different stimuli types[7]. The figure shows clearly that subjects are availing themselves of the added constraint provided by the different categories — demonstrating the adaptive value of their generative inductive mechanisms. As McGonigle & Chalmers (1992) suggested, in the face of increasingly complex and difficult tasks presented over a protracted period, systems would avail of constraints present within the data set to produce progressively economic organisational structures.

Furthermore, both supervised and unsupervised phases of research, together with preliminary evidence from a study involving simultaneous learning of arbitrary and non-arbitrary connectives (McGonigle & Ravenscroft, pers. comm.), indicate that these data-reducing strategies emerge through internal processes of self-regulation, motivated by economy. The benefit of a longitudinal analysis can be clearly seen: progressive increases in task difficulty stimulate the emergence of increasingly effective information-handling strategies on the part of the system.

Self-organisation: economical arbitration

Tasks demanding exhaustive free search of a data set with the only requirement being a non-reiterative path provide rich opportunities for assessing system self-organisation over time. Using this experimental paradigm the first evidence of spontaneous self regulation in both human and non-human primate has been reported (McGonigle *et al.*, 1992; de Lillo, 1994). Despite differences in the discriminable features of items, and no requirement for a particular path through the search space, systems appeared to utilise a spatial strategy to generate a unique path through the space.

[7] The subjects have previously been trained to order the different stimuli.

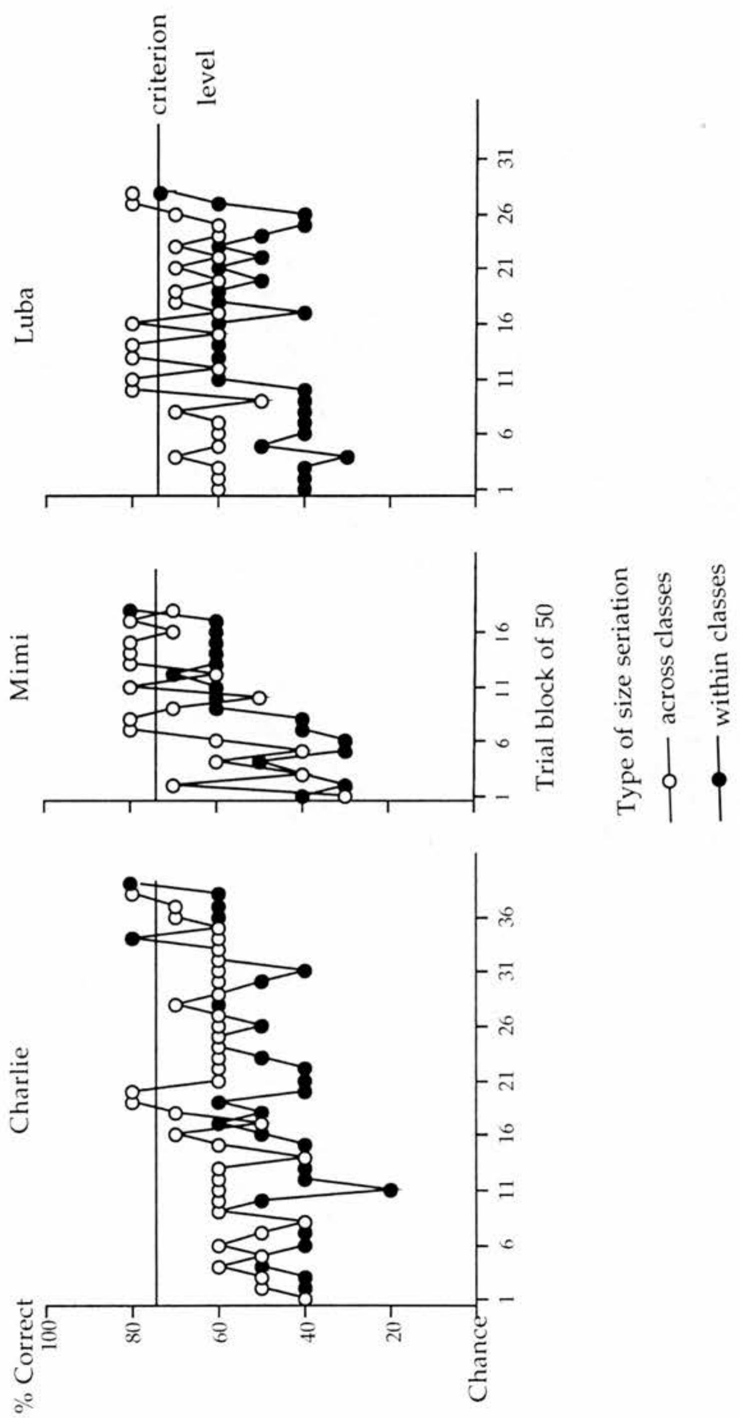

Figure 3.2: Utility function: categorical *vs.* linear constraints

Ontogenetic and comparative fractionation of systems has been achieved within this paradigm. By incrementing set size towards a maximum of nine both the largest set manageable and the solution strategy can be measured. In human, a path based on item adjacency was found to become both more common and more consistent with age and was clearly correlated with task success (McGonigle *et al.*, 1992). A similar developmental profile was observed in adult *Cebus apella* following exposure to increasing set sizes (McGonigle *et al.*, 1992; de Lillo, 1994). Although the paths generated by these latter systems were neither as consistent nor as economical as those of four year old children, their behaviour was clearly self-regulated and based upon internal factors rather than any environmental contingency. The behaviour of pigeon, however, when confronted with this task is clearly divergent: these systems did not self-regulate towards economic generative strategies and could not, in fact, manage sets of greater than three items.

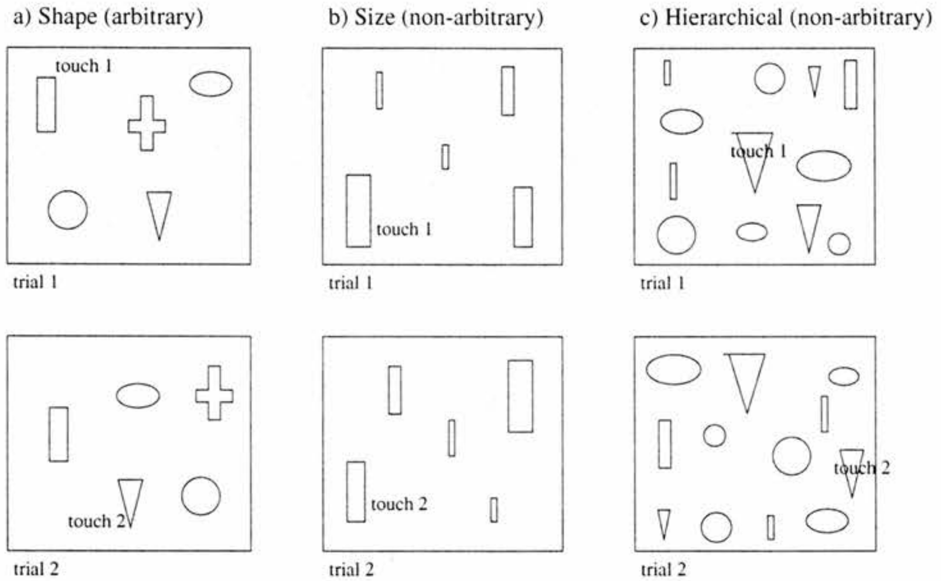

Figure 3.3: The free search paradigm. Feedback (climbing man graphic for child/peanut for monkey) is given to reward any and all exhaustive and non-reiterative searches through the set (After McGonigle, in press)

What would happen when task difficulty was incremented by repositioning stimuli between trials (see figure 3.3)? In the absence of both spatial cues and non-arbitrary connectives systems have been found to spontaneously generate and maintain a privileged path through the search space. In both cases success is achieved through *imposition* of ordered structure on an otherwise arbitrary set.

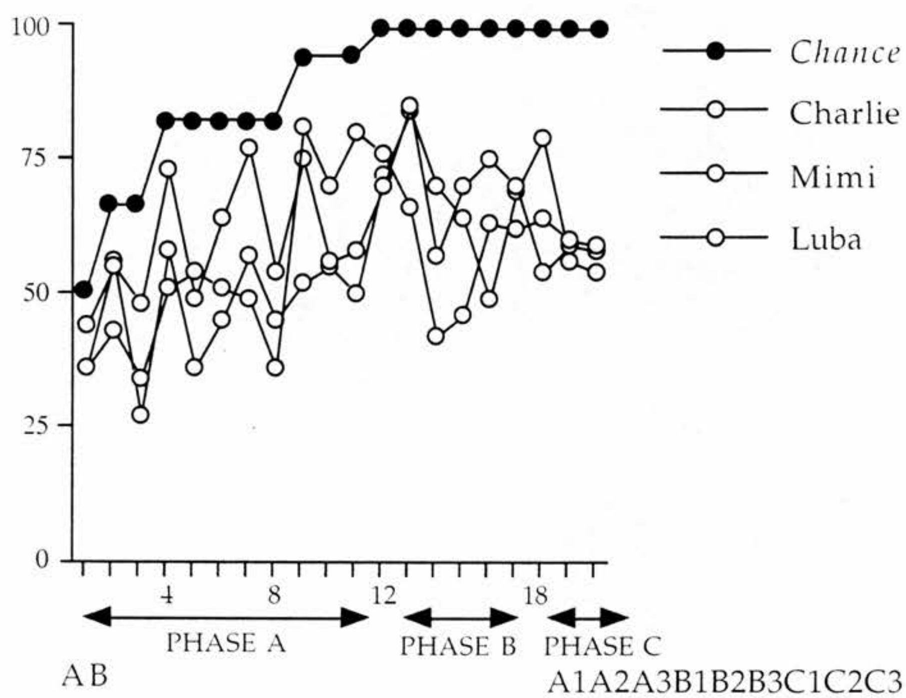

Figure 3.4: Self-organisation in action — increase in task difficulty is compensated by utilisation of constraints (after McGonigle *et. al.*, in preparation)

Figure 3.4 depicts a generic utility function. As task difficulty is incremented, the performance of three adult *Cebus apella* clearly diverges from that expected by chance. Moreover, we can see that as phase C is entered performance actually *improves* — and in the absence of selective reward. McGonigle & Chalmers have developed a paradigm which allows us to observe the adaptive value of learning. Here systems are *self-organising* towards efficient information-handling strategies. Internal arbitration based on economy results in refinement of performance and *generation* of ever more successful cognitive strategies. By taking advantage of the constraints present in the data set, these subjects can learn strategies which are predictive and generative — and thus equally as successful as set size is incremented.

Section 5.2 describes an experiment in robotic navigation motivated by these findings. Over ontogenesis the operation of installed design primitives (sets of very simple locomotor behaviours in this case) is arbitrated by inbuilt economy metrics which enable the system, with experience, to develop system-centred representations of the environment which could support later, more complex, derivations.

Although currently at a preliminary stage, this new extension of the paradigm (McGonigle & Ravenscroft, pers. comm.) provides the potential for detailed fractionation of species. Systems can be assessed on arbitrary, non-arbitrary and hierarchical conditions, with clear predictions of their performance, based on a relational competence constrained by economic data reducing principles.

Summary

Studies of epistemic growth in human and non-human primate suggest a novel characterisation of intelligent biological systems, and a number of important prescriptions for the design of artificial systems. Critically, conjoining comparative and developmental perspectives has overcome the problems of the interdependence of language and thought, and of the origins of meaning for linguistic agents, which have contaminated previous characterisations of cognitive growth.

We have seen that these experimental paradigms reveal agents self-regulating toward maximally economic information handling strategies, utilising whatever constraints are present in task to construct generative and predictive search mechanisms and, in the absence of external constraint, imposing organisation on search sets. The resulting characterisation shows the epistemic agent on an open-ended growth trajectory, from a number of *rational design primitives* which form the ontological lower bounds of the system, developing ever-more powerful information handling strategies in the face of cognitive challenge through internal *self-regulation* based on economy. Serial control of behaviour reflects an internal hierarchical organisation at multiple levels of abstraction as suggested by Lashley (1951) and does not depend, as implicitly assumed throughout much of psychology, on the possession of language, nor on immediate environmental contingencies.

This characterisation is inherently dynamic and situated, sharing with these perspectives an emphasis on development through intra-system and system environment interactions, but diverging from both in the characterisation of the key feature of adaptation. Where, as we saw in the previous chapter (sections 2.1 & 2.2), both dynamic and situated perspectives characterise progressive adaptation as a refinement of the coordination of primarily sensory-motor behaviour, here we have a view of systems self-regulating and becoming increasingly more *epistemically* adaptive through processes of *internal arbitration* of their behaviours according to inbuilt and learned utility metrics. Grounded in behavioural assays of system competence, from this perspective we can characterise progressive system adaptation at *both* behavioural and epistemic levels, where these levels are viewed as fundamentally intertwined: systems are both embodied and situated. Their ontological lower bounds, and their resulting developmental trajectories are observable and, critically, measurable.

Finally, this developmental and comparative paradigm, allows us to bridge the gap between cognitive accounts focusing on a mischaracterised logico-deductive competence, and traditional behaviour-based accounts which emphasise associationistic modification of tightly-coupled behaviours. The origin of high-level competences is no longer a miracle: we can see their precursors in the variety of specialised inductive mechanisms which form the ontological lower bounds of systems and, furthermore, gather rich data (see figures 3.2 & 3.4) which allows us to characterise the adaptive value of these mechanisms.

3.2 Engineering

3.2.1 Design prescriptions

The biological characterisation outlined in the previous section (3.1.3) suggests that our attempts to construct intelligent artificial systems must incorporate a number of critical design features. The following features obviously do not constitute an exhaustive list but rather embody a number of principles which have emerged from our analysis and which must be incorporated within any artificial architecture which aspires to ontogenetic extendibility — the critical feature of intelligent biological systems (McGonigle & Chalmers, 1998*a*; McFarland, 1999), and one of the central problems facing contemporary AI (Kirsh, 1991, for example).

The system should be both embodied and situated. Its *ontological lower bounds* should be sufficiently rich, both in hardware and software (e.g. inductive primitives, arbitration criteria) to enable development of the system, both through self-organisation and designer incrementation, towards its design goals. The system should possess memory enabling progressive adaptation over ontogenesis: it should have a *life-history*. The system should be state-based, modular and both hierarchically and serially organised.

Lower bounds The ontogenetic lower bounds of biological systems ultimately determine their ontogenetic potential. The richness of both hardware and wetware distinguishes systems stuck at primitive levels of adaptation from those capable of progressive adaptation.

Embodiment An architecture should be implemented on a robotic platform. Although simulation certainly has its place in the quest to synthesise intelligence, the ultimate testing ground is the real-world. Whether simulated or embodied, relevant structural properties of systems should be utilised to constrain inductive problems where possible (Ballard, 1993; Ballard *et al.*, 1997, for example).

Although engineering accurate sensors and actuators is not the sticking point[8] systems should be endowed with sensors rich enough to deliver per-

[8] Disregarding the *apologia* of some roboticists!

ceptual data of sufficient quality to enable task solution. They must also ideally be able to act upon the world: action scaffolds development (Miller, 1966, for example).

Situatedness Systems should be implemented in real, rather than simplified 'toy', worlds. Niches should be sufficiently dynamic to provide a rich error space, and to challenge the system to develop over time. Ideally the meaning of representations to the system should emerge through self-organisation and environmental interaction not from the ontological stance of the designer.

Limited niche engineering, as ubiquitous in human societies, has an important role to play in constraining computation. Furthermore, specific features of the niche can be utilised to constrain the solutions to problems which might otherwise be intractable (Donnett & McGonigle, 1991; Ballard, 1993; Horswill, 1995).

Behavioural and cognitive primitives Biological systems possess a range of reflexes and 'instincts', often system-preserving behaviours, critical at early developmental stages. There should be no objection, therefore, to installing reactive system-preserving behaviours in our artificial systems.

Many biological systems appear to be endowed with relational primitives which underlie later sequential competences, and inbuilt judgement criteria, such as economy principles, to support arbitration between behaviours (McGonigle & Chalmers, 1996, 1998*a*). Artificial systems' dependence on limited computational and power resources make economy particularly vital. Our artificial systems should be supplied with such mechanisms in order to constrain induction, and thereby enable development.

As we have seen, biological systems possess a range of adaptive specialisations: neural inductive machinery optimised for particular internal and external contexts. These range from optimised associationistic mechanisms to much slower abstracted inductive devices. Similarly our artificial systems should be endowed with a variety of inductive mechanisms optimised for operation in different contexts, at different time scales, and at different levels of abstraction.

Modular Modularity of design provides robustness, but also allows a system to possess multiple encapsulated, and sometimes logically incompatible, competences simultaneously (Fodor, 1983; Sherry & Schacter, 1987; McGonigle, 1991). Furthermore, ensuring that low-level behaviours are encapsulated, with clearly defined initialisation, success and failure conditions, provides the potential for syntactic recombination by controllers at higher levels.

Hierarchical Hierarchical organisation underlies systems' ability to recombine behaviours (the 'keyboard metaphor' of serial control — McGonigle & St Johnston, 1995) at multiple levels to achieve goals. Furthermore, an hierarchical organisation allows for multiple mechanisms (control, inductive, monitoring *etc.*), operating simultaneously at multiple time scales and levels of abstraction, providing the potential for abstraction of statistical regularities invisible to modules at lower levels.

Serial The problem of serial control inheres to all actions of systems (Lashley, 1951). A serial controller capable of accessing modular behaviours provides the opportunity for the development of artificial systems with a true *behavioural syntax*. Furthermore, a serial architecture based on state transitions immediately confers contextuality upon a system. Behaviours can be instigated and interpreted within specific contexts indexed by both internal and external features.

Autonomous Systems should not merely be 'flexless' in the sense of autonomy used by Brooks (1991*c*) but should rather be able to *self-select*, and ultimately *self-regulate*, at initially behavioural, and later epistemic, levels. Systems should not be entirely reactive but should be capable of self-determining appropriate behaviour(s) given environmental and task demands. Furthermore, behaviours should be arbitrated by internal criteria as well as the local environment.

Such autonomy requires systems to be capable of recombining behaviours to achieve goals (and hence an hierarchical organisation), and to interpret the effects of behaviours within given contexts (provided by seriality and state). Self-selection, implementation, termination, and arbitration of behaviour requires representation. A pragmatic view of *state* should be adopted which incorporates external (niche engineered) and internal (endocrine, distributed, or symbolic) forms of representation according to task demands. The critical adaptive role of state is regarded as the provision of *contextual information* vital for internal control, and interpretation, and reinterpretation of internal and external phenomena as demonstrated throughout the biological realm (Hinde, 1982).

Systems must therefore be innately endowed with, or acquire, arbitration criteria, and be cognisant of both behavioural success and default. Learning from error is the prime motor of abduction and fundamental to internal arbitration. Without cognisance of error there can be no scope for progressive adaptation and, particularly, epistemic growth over the life span.

Self-organisation through inter-system, and system-environment interactions over the course of ontogenesis should lead to progressive system adaptations based on installed cognitive primitives, recombination of, and arbitration between different behaviours.

Life history As with all complex biological systems, artificial systems should possess memory allowing progressive adaptation such that later behaviours rest on earlier achievements (McGonigle, 1995), not merely recapitulating previous successes but also allowing the generative benefits (Fodor & Pylyshyn, 1988) to emerge,and so that epistemic derivations of initially purely behavioural competences can emerge.

3.2.2 Robotic implementations

The biological characterisation outlined in section 3.1.3 and the design features it suggests, together with the mission statement of McGonigle (1991), are the the inspiration behind a number of robotic implementations within the Laboratory for Cognitive Neuroscience and Intelligent Systems at the University of Edinburgh. All of these experiments are motivated by adoption of a *'logical hierarchy of design'* (first detailed by McGonigle 1991 — see figure 3.5) which has grown out of *a priori* consideration of a design strategy, in conjunction with appreciation of the *rational constraints* on the solutions to robotic tasks.

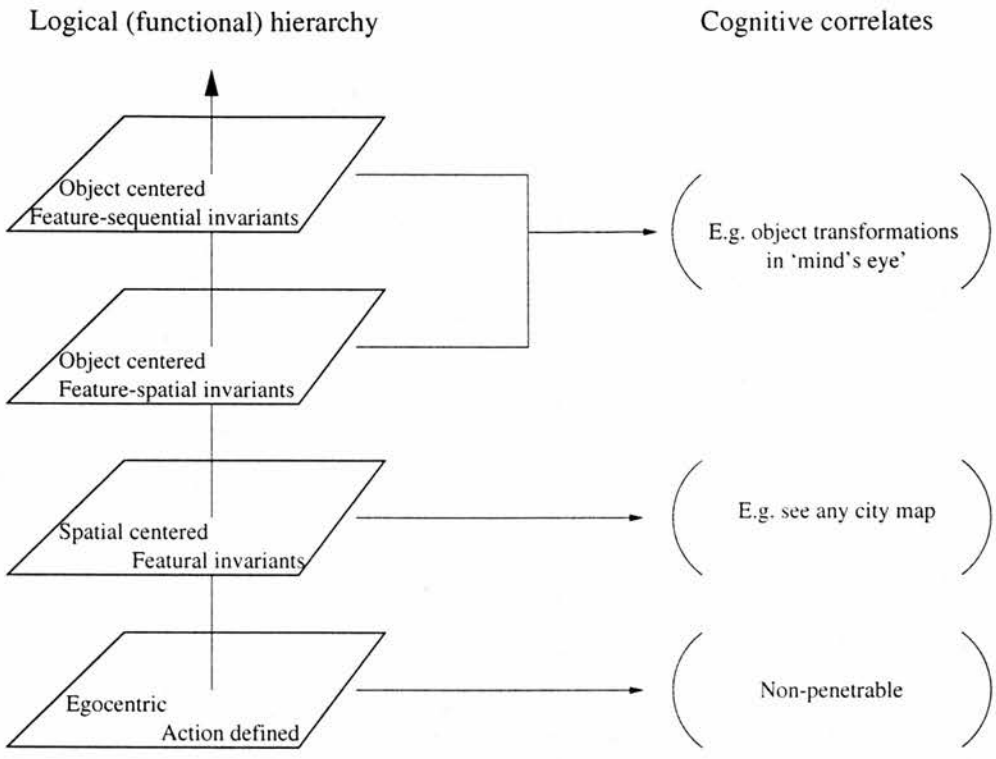

Figure 3.5: The logical hierarchy (after McGonigle, 1990).

Several of these implementations will be briefly described in order to demonstrate implementation of some of the design features outlined in section 3.2.1 and to depict the history of architectures motivated by our synthetic stance. 'Incrementing by design' (McGonigle, 1991) is a core feature of this approach.

The hearing robot

Donnet & McGonigle's (1991) hearing robot was the first from this perspective. It was designed to navigate through an office environment using a distant sound source for orientational and contextual cues.

Agent movement creates a demand, not merely for avoidance, but also for a *locative* competence in order to allow the system to return to a previously occupied location. Donnett & McGonigle (1991) developed a navigational algorithm for the hearing robot using an acoustic beacon and the resulting phase and intensity differences characteristic of various parts of the niche combined with robotic tropotaxis. Exploitation of niche constraints was therefore a key element of the navigational strategy adopted.

Loss of signal might arise either due to a failure of the system, or due to baffling from objects. This implementation was invested with an *interpreter*, essentially a set of failure diagnosis procedures (McGonigle, 1998) which enabled it to interpret absence of signal and take appropriate remedial action: check acoustic sensors; move at random 'hunting' for sound; or check acoustic bins. The value of an on-board state-based interpreter in helping to solve critical control problems is clearly demonstrated by this implementation. Furthermore, the state based interpreter allowed reinterpretation of infra-red readings: on approach to the sound source the tightly-coupled avoidance behaviour was overridden by an approach behaviour.

Control of the system swung between target approach and obstacle avoidance behaviours, especially in areas of high object density. *Prospective control* was achieved through constructing a map of 'clutter' in the environment by logging the frequency of interrupt between goal seeking and local obstacle avoidance. A cumulative record of the distributions of objects along the route allowed the robot to gear its speed using a look-up table. Assumptions concerning stability of initial position (a 'home' or 'base' location) and the locative invariance of objects in the environment were made to simplify the task.

Behaviour shaping

The Edinburgh R2, a compact autonomous robot, was used for the next phase of research (Nehmzow, 1991; Nemhzow *et al.*, 1993; Nehmzow & McGonigle, 1995). A neural network controller mapped sensors to actuators and received feedback from a human supervisor. Light sensors provided orientational information; infra-red detectors sensed object proximity. The system was designed to examine whether multiple competences could be supported by a single NN architecture, and to determine if multiple competences could be synthesised to form newer and more complex behaviours — currently a major problem facing AI (Brooks, 1986*a*; Malcolm *et al.*, 1989).

Through external supervision the R2 was rapidly taught competences of obstacle avoidance, contour following, phototaxis and (simple) route learning. Learning typically took around 10 minutes even for the most complex combined competence. *Sensor fusion* formed a necessary precondition for success on a maze learning task which combined object avoidance, wall following, and phototaxis.

Although competence acquisition was very rapid, training was still necessary for the controller to switch between tasks — a consequence of the adoption of a neural network controller, although rapid learning of fundamental robotic competences through external shaping was demonstrated.

Light compass navigation and pole sorting

Using a distant light source algorithms were developed (Nehmzow & McGonigle, 1992, 1993; McGonigle & St Johnston, 1995) to allow the robot to return to near a base location and exploit the structure of the environment in order to get to base by means of local landmarks. The R2's task was to search for, retrieve, and order a set of targets upon a workbench of known location.

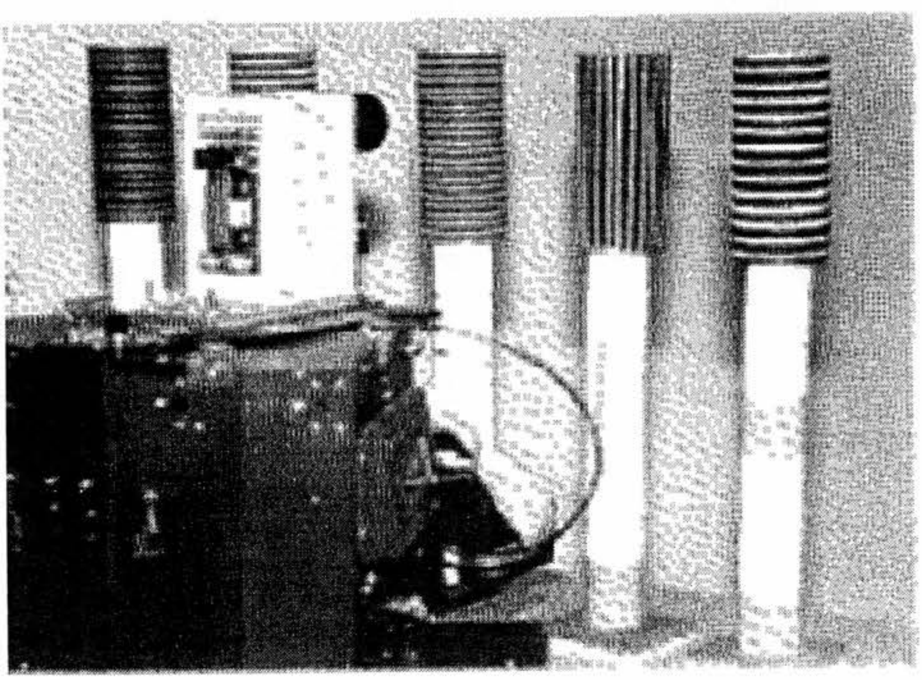

Figure 3.6: The R2 ordering poles

Navigation

The navigational algorithm relied on two sources of information: light gradients throughout the niche which provided a general orientational frame of reference; and on-board odometric information based on wheel actuators which supported dead reckoning. Each run started from a fixed point against the workbench where the poles were to be sorted. The distance from workbench to light source was preinstalled in the robot to provide a means of eliminating the cumulative odometric error inevitably consequent on movement. Between sorties, the R2 pushed against the workbench, and all odometric information was recalibrated. Target search was decomposed into two stages differing in their level of *granularity* — greatly helping to constrain the solution to the navigation problem. Initially the R2 attempted to attain proximity to the target site, upon which a fine grained search for targets began.

The task illustrates the importance of mapping task stages to internal robot *states* to enable it to do the right thing at the right time — once the target region had been attained, object signals could be reinterpreted: the targets were to be lifted not avoided.

Active vision

Each target was returned to the workbench thereby establishing a set of targets in close proximity to each other. This strategy greatly reduced the time spent identifying and sorting the targets. Active vision could now commence based on a compare-and-contrast mechanism identified within both rodent and primate (McGonigle & Jones, 1975, 1978). Now the set of targets had been locatively encoded, the defining features of the targets were extracted on the basis of comparing and contrasting them with others from the set (see figure 3.6).

Rational *economic* principles were used to reduce data load. Selective attention was achieved by peripheral blocking (McGonigle, 1998): high level vision was only instigated when a target was present in the gripper. This strategy helped to overcome the fixed focus limitation of the Imputer camera, varying illumination directions, and ensured that only the centre of the obtained image had to be processed — overcoming the figure-ground problem. Once each target had been successfully discriminated and identified a bubble sort algorithm was used to order them.

Summary

The implementation highlights the importance of a state-based system. Mapping of different task stages to internal state allows for signal reinterpretation in different contexts. The task grammar allows recombination of behaviours to meet adaptive needs.

3.2.3 Summary

Following the hierarchy of design (McGonigle, 1991), based on the logical dependencies between competences, the systems described above have become increasingly more adaptive — first reactive behaviours were developed, followed by a locative competence supported by navigation, and only then vision. Throughout a criterion of adequacy has been used — the implementations are designed not to 'solve' particular engineering problems, but rather to support the progressive development of the system.

Some of the critical design features outlined in section 3.2.1 have been implemented and demonstrated to be a key engineering feature of these successful implementations. Critical elements include:

Exploitation of niche constraints Engineering of a system to its niche, and the niche to the system (*niche engineering*) helps to simplify problems which might otherwise be intractable (Donnett & McGonigle, 1991; Ballard, 1993; Horswill, 1995).

- Donnett & McGonigle (1991) developed a navigational algorithm for a 'hearing' robot using an acoustic beacon and the resulting phase and intensity differences characteristic of various parts of the niche combined with robotic tropotaxis.
- Nehmzow & McGonigle (1992, 1993) and McGonigle & St Johnston (1995) provided an interlinked series of implementations which utilised a distant light source in conjunction with niche constraints to allow a robot to navigate to base by means of local landmarks.

Prospective control Was achieved as a derivation of Donnett & McGonigle's (1991) implementation through building a map of obstacle density based on signal interrupt.

State A feature of these synthetic approaches is the importance of context both for behaviour initialisation, and signal (re-)interpretation.

- Donnett & McGonigle's (1991) hearing robot reinterpreted signals based on internal context allowing both object avoidance, and docking, within the same implementation.
- McGonigle & St Johnstons' (1995) light compass navigation system featured a *task grammar*. Decomposition of the task into segments with varying *granularity*, and consequent mapping of task stages to internal robot states, constrained, and thus simplified, the navigational problem allowing the robot to 'do the right thing at the right time'.

Learning An example of an early learning approach was provided by Nemhzow *et al.* (1993) and Nehmzow & McGonigle (1995) who applied a method analogous to 'shaping' through external supervision to teach multiple competences to a neural-network controlled mobile robot. Unfortunately training was still necessary for the controller to switch between tasks, although external shaping of fundamental robotic competences was demonstrated.

Rationality and economy McGonigle & St Johnstons' (1995) active vision implementation, heavily informed by characterisation of primate visual competence (McGonigle & Jones, 1975, 1978), followed logically in the design hierarchy from earlier navigational competences. Visual identification and ordering of targets was achieved using an active compare-and-contrast principle. Furthermore *economic* principles were utilised in order to reduce computational load.

With a history of implementations incorporating a number of important design features, and a prototype state-based, functional architecture capable of supporting multiple competences within the same system together with the concept of task grammars which allowed recruitment and recombination of modularised behaviours, the task became twofold:

- To more fully incorporate these design features within a single architecture, constructed to be easily extendible by the designer. This architecture should also be state based and hierarchically and serially motivated whilst incorporating greater potential for selection and arbitration of behaviour.

- To introduce *self-organisation* into the architecture, both to more accurately model our biological targets, but also in order to circumvent some of the problems of the hand-crafting of behaviour.

The next chapter outlines an architecture which strives to meet these objectives, before a description is provided in chapter 5 of some experiments in self-organised navigation.

3.3 Summary: the synthetic stance

This chapter initially considered some requirements for the reverse engineering of biological systems. Comparative research suffers from an impasse: cognitive accounts are focussed on primarily linguistic achievements, and consequently ignore phylogeny, ontogeny, and the range of adaptive behaviours found across many systems; behaviourist accounts tend to focus exclusively on an optimised associationistic inductive mechanism which, although certainly part of the complement of learning mechanisms at systems' disposal, cannot in isolation account for the complexity of adaptive behaviours, nor explain their ontogenetic origins.

Ethological research depicts the richness of adaptive behaviours. Systems possess both evolutionary engineered instinctive behaviours, plus a variety of specialised inductive mechanisms, optimised for use in particular contexts. Adaptive behaviours are inherently serial in nature — a factor unaddressed by the traditional behaviourist, animal learning, and cognitivist perspectives on systems. Evidence from cognitive neuroscience indicates that phylogeny results in qualitative changes in neural structure. Systems possess a range of *adaptive specialisations* as their ontological lower bound which constrain their ontogenetic growth through processes of inter-system and system-environment interaction.

Determining the interaction of lower bounds and environment over epigenesis, particularly with respect to serial competences, requires comparative and developmental paradigms untainted by language. The research programme of McGonigle & Chalmers, is an exemplary example of such a paradigm. Their evidence suggests that relational primitives might form part of the ontological lower bounds of many complex systems, both human and non-human primate, which in interaction with environments of increasing complexity result, over ontogeny, in increasingly complex serial, and categorical skills. Furthermore, their comprehensive longitudinal studies indicate that complex systems *self-organise* by internally arbitrating their behaviours according to metrics such as economy principles, which might also be built in to systems. They suggest that the associationist inductive mechanism might be supplemented by a more powerful relational mechanism in complex systems.

Their detailed comparative and developmental evidence suggests that artificial systems should be designed with rich ontological lower bounds, consisting of reactive, system-preserving behaviours, specialised inductive mechanisms, and arbitration criteria. Their analysis suggests that artificial systems should be state based and both hierarchically and serially organised. These systems should be able to select, and arbitrate between behaviours, becoming more adaptive over time through self-organisational processes.

Finally, a number of implementations inspired by this characterisation, incremented over the past decade through a logical hierarchy of design, were described. Between them these implementations incorporate a number of key design features including, particularly, *state-based interpreters* which allow signal reinterpretation in differing contexts and, in combination with *task grammars* allow for selection and recombination of behaviours from a core repertoire, and *rational design primitives* which reduce computational load and enable simplification of the solutions to complex problems.

The next chapter describes an architecture designed to supplement these precursors by introducing some more key design features, whilst striving for easy extendibility.

Chapter 4

A synthetic architecture

We saw in the last chapter that the extensive research programme of McGonigle & Chalmers at the Laboratory for Cognitive Neuroscience and Intelligent Systems has inspired a number of implementations of artificial systems. This chapter describes an extension of these architectures which strives to incorporate many of the key design features outlined in section 3.2.1 within a modular construction designed for easy extendibility.

4.1 Introduction and aims

Arguably the most critical problem currently facing robotics is that of *extendibility* (McGonigle, 1991; Kirsh, 1991) — how to design a robotic control architecture which can scale, either through design, or by learning? This implementation has grown out of a research programme at the Laboratory for Cognitive Neuroscience and Intelligent Systems at Edinburgh, whose fundamental tenets were set out in a mission statement by McGonigle (1991). A critical feature of the position here adopted is that *design* is critical in the attempt to develop complex artificial agents. Furthermore there exists a *logical hierarchy* of design which helps to provide constraints on how an artificial system should be constructed. Our dynamic perspective suggests that systems should also scale themselves over ontogenesis. Scaling of our systems, therefore, is envisaged to occur through both explicit hand-design and self-organisation over ontogenesis through interaction of its complement of design primitives with its environment.

Just as the forms and processes of complex biological agents do not derive purely from random mutation coupled with local environmental selection and arbitration but also from the capacity for self-organisation over ontogenesis, artificial systems should be designed to become progressively more adaptive over the course of their lifetime. An important, and unique, part of our research programme is that its target is *epistemic* rather than purely behavioural adaptation over the long term. We strive to construct our artificial systems such that their capabilities are *extendible* so that complete re-design is not required in order to achieve progressively greater levels of adaptation.

Artificial cognitive systems should be viewed as *dynamical systems* in continuous interaction with their dynamic environments. Taking such a systems view indicates that the ontological lower bound of a system (specifically in terms of perceptual and cognitive primitives) should be sufficiently rich to allow the development of the desired behaviours. A system's behaviour should always be viewed as the result of environmental interactions, where categorisation falls naturally out of correlating internal and external variables, where cross-modal associations are a logical consequence of interaction, and where categorisation and concepts are immediately grounded in the experienced environment. Such considerations also help to reduce the computational load of a system (cf. Ballard 1989; Ballard *et al.* 1997; McGonigle & St Johnston 1995; McGonigle & Chalmers 1998*b*).

4.1.1 Design primitives

The artificial systems constructed as part of the research programme so far are all *embodied* and *situated* within a non-trivialised environment — in an attempt to avoid the problems consequent upon environmental simplification which afflicted much early robotic research (Nilsson, 1984, for example). Critically, our systems are not designed with some abstract generalised environment in mind — all biological organisms exist within, and are adapted to, ecological niches. Tailoring an agent to its environment, and the niche to the agent (*niche engineering*) often enables an agent to exploit constraints embedded within the niche in order to solve problems which might otherwise be intractable (Donnett & McGonigle, 1991; Ballard, 1993; Horswill, 1995; Pfeifer, 1997).

Artificial systems should also be *autonomous* — able to select, maintain and, critically, *evaluate* their own behaviours. Parallel research with children and brown capuchin monkeys (*Cebus apella*) within the research group (McGonigle & Chalmers, 1996; Chalmers & McGonigle, 1997; McGonigle & Chalmers, 1998*a*, for reviews) indicates that *economy* acts as a fundamental constraint on both the solutions to a variety of problems relating to the serial control of action, and to internal arbitration between behaviours (McGonigle & Chalmers, 1996; Chalmers & McGonigle, 1997; McGonigle & Chalmers, 1998*a*) helping an agent to become progressively inductively more powerful over the course of its existence. Such arbitration criteria must therefore be a core feature of our systems, whether preinstalled (as with the economy criterion adopted herein) or acquired through experience. A further critical point is that agents should be studied over a period of time, in order to allow development to occur, later behaviours resting on earlier achievements. Agents, therefore, should have a *life-history* (McGonigle, 1995).

Systems should be *hierarchical, serial* and *state-dependent.* Ethological research clearly demonstrates that context modulates behaviour throughout the biological world (Hinde, 1982). By progressively adding greater constraint, and further interpretation at successive levels, an hierarchical design allows an agent to break free from tightly-coupled behaviours subject only to local environmental arbitration into a new realm of self-selection and internal arbitration of behaviour which is the most striking characteristic of intelligent systems.

Modularity of information-processing is the clearest result of neuropsychological work to date, both empirically and theoretically (Kandel *et al.*, 1991; Gazzaniga *et al.*, 1998). Furthermore a modular system, as well as making a system more *robust* supports multiple behavioural capacities some of which might potentially be logically inconsistent with one another (Fodor, 1983; Sherry & Schacter, 1987; McGonigle, 1991). This evidence contrasts strongly with the traditional AI view of a central organisation of the representation of knowledge and the production of behaviour as a sense-think-act cycle (Luger & Stubblefield, 1998*a*).

4.1.2 Summary

The aim of this research is to provide a control architecture which supports recombination of action and behaviour primitives without necessitating redesign and which is capable of supporting competences in such a way that new competences can easily be designed from lower level **atoms, acts**, and **behaviours** without requiring redesign of the system. The system should support multiple competences and, critically, the system should possess memory, enabling development of competences across its life-span. Along with incrementation through design, the system should be capable of self-organisation over the life span (see chapter 5), arbitrating between behaviours according to criteria such as economy, resulting in progressive adaptation over time.

4.2 A brief description

Hardware

The robot used throughout these experiments was a Nomad 200 (see Figure 4.1), supplied by Nomadic Technologies Inc., California. It is equipped with sixteen infrared detectors, and sixteen sonars (see Figure 4.2), located around the periphery of the chassis, along with base-mounted odometry and a magnetic compass (KVH C100 Compass Engine). The on-board computer is a 80486 DX2 processor running at 66MHz, with a fully suspended 1 Gb hard disk, floppy disk drive, and 20Mb of RAM. Remote control of the Nomad is achieved *via* a radio ethernet card, allowing programme and data transfer from any one of a number of machines situated throughout the niche.

The niche

All of the experiments detailed herein took place within the Laboratory for Cognitive Neuroscience and Intelligent Systems, Level 8 Appleton Tower, Edinburgh — a network of rooms and corridors all of which contain desks, chairs, machinery, and of course humans, providing a rich niche (see Figure 4.3) in which to develop navigational competences.

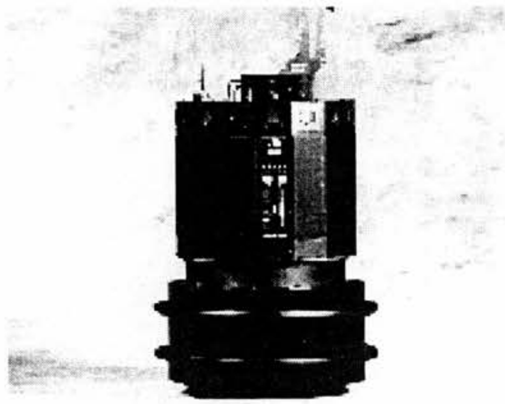

Figure 4.1: The Nomad.

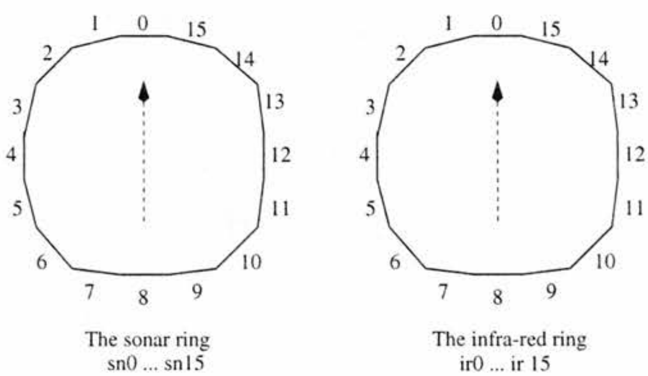

Figure 4.2: Sonar and IR numbering

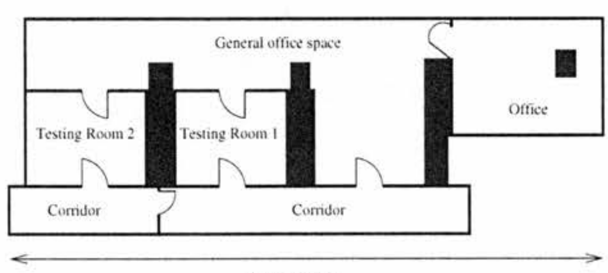

Figure 4.3: The Nomad's niche.

Software

The Nomad was programmed in a mixture of the C/C++ and Perl (release 5.003) programming languages, running under the Linux (Red Hat 4.2) operating system. Each language provides its own benefits. C/C++ was used to communicate with the on-board robot control system, and is particularly useful in the control of time-critical processes. Since Perl is an interpreted programming language it need not be pre-compiled and can therefore be debugged dynamically thereby helping to shorten the development-test-debug cycle. The strength of Perl running under Linux lies in the control of concurrent processes, and the dynamic construction of both programs and data structures.

State arrays

Contextually appropriate behaviour is impossible without some form of state. This architecture utilises two arrays of data, one referring to *external* factors, and the other to *internal* context (see Figure 4.4).

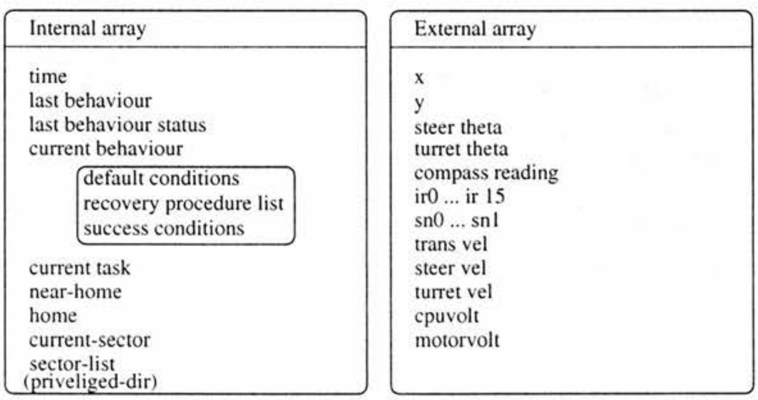

Figure 4.4: Internal and external states.

Each behaviour on start-up and, critically, on completion modifies relevant state variables.

Behaviours

Each behaviour, although not an object in the strict programming sense, was treated as a distinguishable data structure (see Figure 4.5) sharing a range of routines and relevant data which was invisible to other behaviours. Data stored refers to a behaviour's name, class, number of times called, and success rate. Local procedures test initialisation and default conditions, as well as actually executing the behaviour and altering relevant internal and external state variables on completion. A list is maintained of relevant *recovery procedures.*

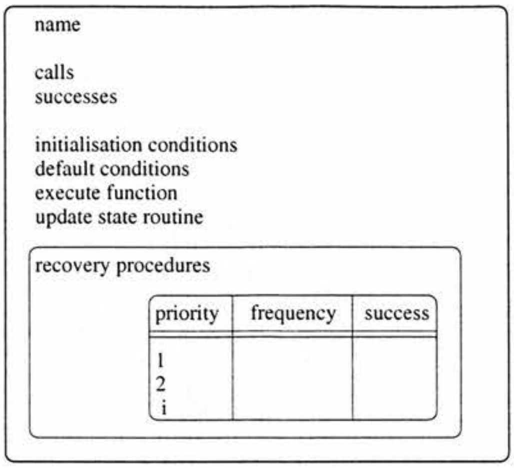

Figure 4.5: Components of a behaviour.

Recovery procedures can be regarded as less elaborate behaviours storing information about their name, success and default conditions, status on completion, and the recovery procedure itself.

Controller

The central routine of the architecture (see figure 4.6) selects a behaviour according to current whereabouts in task space, and then monitors its progress using the behaviour's own default detection routine. As the Nomad uses the Linux operating system it is possible to mimic concurrence by using the system call *'fork'*, and then run the behaviour as a 'child' process with monitoring of default occurring within the 'parent' process.

Following the fork, a copy of the original program is made and the two essentially identical processes begin to diverge. As a child process cannot access data structures within the parent a signal handling system allows the monitoring parent process to 'kill' its child if a default or timeout occurs, and for the child to send a success signal to the parent on completion.

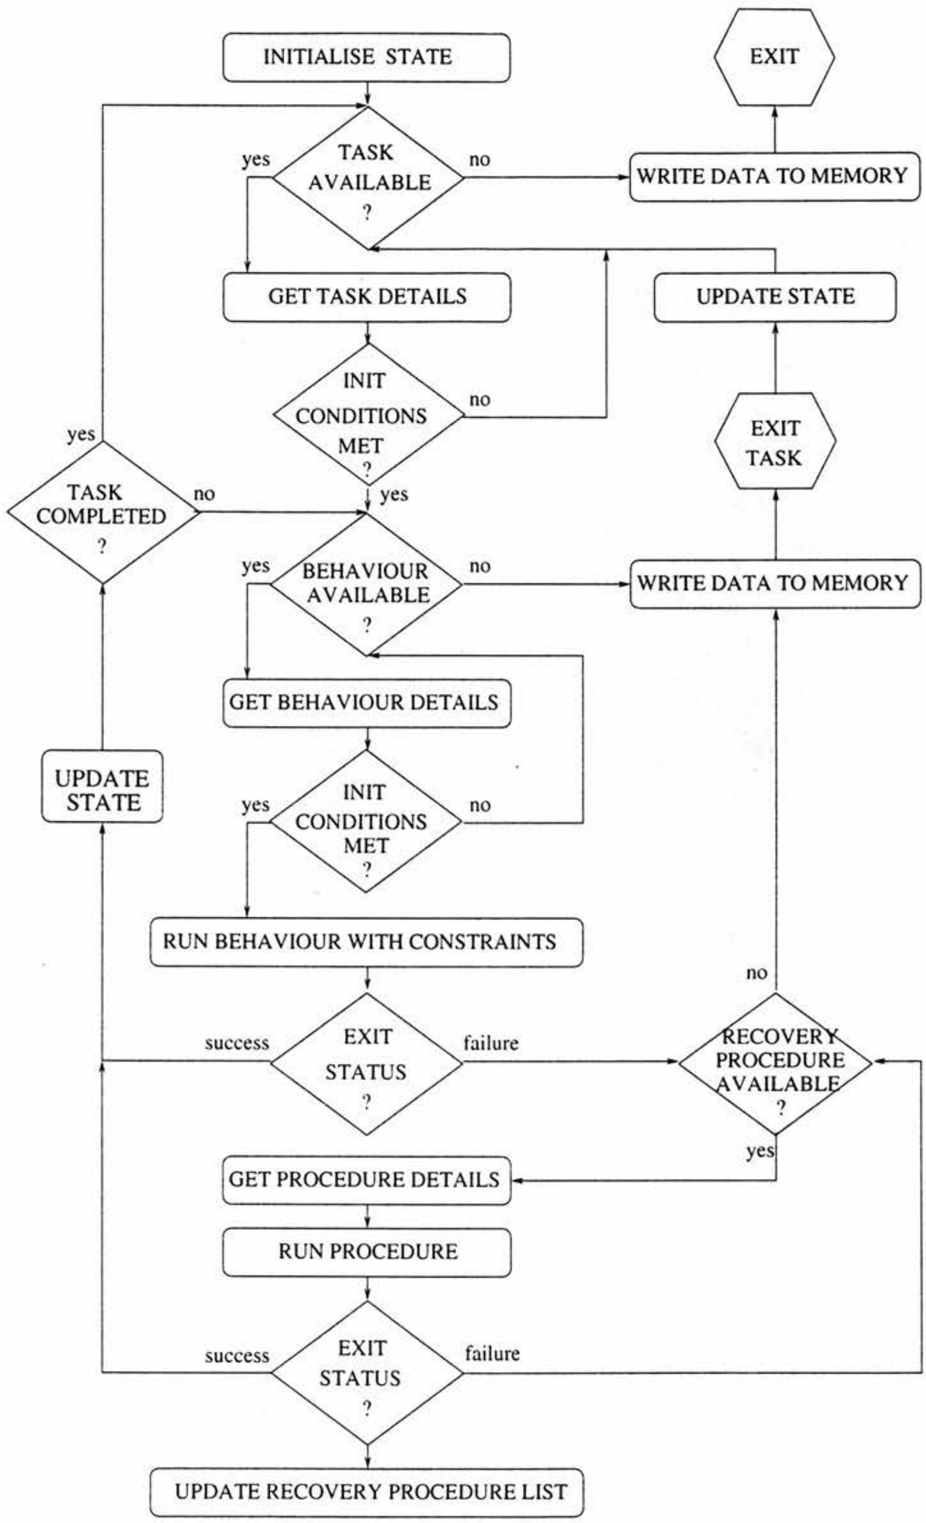

Figure 4.6: Controller.

Memory

For biological agents development is crucial — early stages constrain later ones. The history of an organism, its experiences and how these influence its developmental trajectory is critical. Any attempt to develop a complex artificial system must take this fundamental property of biological agents into account. Giving a life history (see McGonigle, 1991) to a robot is imperative if we want to obtain development of the system not just within individual runs but across the entire life-span of the agent.

This architecture uses the module 'Storable', available from the Comprehensive Perl Archive Network (CPAN)[1], to recursively store all information relating to behaviours, tasks, and the state arrays. At any point in time information from earlier in the current run or from any point in the past can be re-obtained. The Nomad is therefore freed from the short-termism characteristic of many artificial agents!

Summary

The Nomad has access to data structures and programs which are arranged hierarchically. The foundation of the hierarchy is a set of atomic action units, **atoms**, such as `vm(translation, base, turret)[0,1]` (moves the Nomad according to the three parameters, returning 1 on success, 0 on failure), `gs()[0,1]` (reads all sensor data, returning 1 on success, 0 on failure), `tk(speech-stream)[0,1]` (sends `speech-stream` to the Nomad's voice synthesiser, returning 1 on success and 0 on failure). These atoms come pre-installed (see section 4.3.1).

Above this level the atoms are combined into sequences of actions: **acts**. A familiar example of such an act is `avoid(distance)[0,1]`, which combines movement and sensing atoms. See section 4.3.4 for a comprehensive list of available acts.

The next level of the hierarchy consists of combinations of actions into more meaningful, or rather non-trivial, sequences: **behaviours** (section 4.3.5). Behaviours are complex data structures (see Figure 4.5). Each behaviour includes information concerning: the number of times it has been called; the number of successful calls; initialisation

[1] `ftp.funet.fi`.

conditions (both external and internal criteria); the actual running of the behaviour; constraints operating on the behaviour; a list of implementable recovery procedures (each containing information about the number of times they have been called and their success rate); and finally procedures for modifying relevant state variables.

The final layer of the action hierarchy consists of **tasks** (section 4.3.6). These are essentially *meaningful sequences* of behaviours, where initialisation conditions of behaviours determine lawful sequences. Each defined task maintains its own set of internal variables. An example of a complex task involves running the three behaviours `learn`, `retrace`, and `dead-rec` (section 4.3.6).

The complete architecture is divided into a controller (see Figure 4.6), a memory, and lists of tasks and behaviours. Upon start-up the Nomad picks its task and updates the appropriate state variable accordingly. Next, the behaviours in the current task will be run sequentially, with simultaneous logging of constraint status (default and success). Upon successful completion of the current behaviour, transition to the next becomes possible. Hence a sequence of behaviours is implemented, default being dealt with by (currently pre-installed[2]) recovery procedures.

4.3 A more detailed description

4.3.1 Pre-installed behavioural atoms

The Nomad comes with a range of functions which allow communication with the on-board robot motor controllers, and sensor boards[3].

Configuration

`ac(trans,steer,turret)[0,1]` Sets the translation, steering and rotational accelerations of the robot (in units of $\frac{1}{10}$ inch and $\frac{1}{10}$ degree per second squared). Returns 1

[2] It is proposed to extend self organising principles to error recovery itself, by having the system establish its own error typology which would then support expert abductive diagnosis — see section 6.3.2.

[3] Throughout this chapter functions are defined using the following convention:

`function-name(parameters)[return values].`

upon success, 0 upon failure.

`connect-robot(id)[0,id]` Requests a connection to the robot with `id`. Returns the `id` of the robot upon success, 0 upon failure.

`conf-cp(status)[0,1]` Configures the compass system (0 =off, 1 = on, 2 = self-calibrate). Returns 1 upon success, 0 upon failure.

`conf-ir(dependency,order)[0,1]` Configures the infra-red sensor system (`dependency` is the percentage dependency of the current returned reading on the previous returned reading, between 0 and 10; `order` is a list specifying which sensors are to be switched on, and in what order). Returns 1 upon success, 0 upon failure.

`conf-sn(firing-rate,order)[0,1]` Configures the sonar sensing system (`firing-rate` specifies the interval in ms between the firing of sonars; `order` is a list specifying which sensors are to be switched on, and in what order). Returns 1 upon success, 0 upon failure.

`conf-tm(timeout)[0,1]` Sets the timeout period of the robot in seconds such that if no command has been received for longer than `timeout` all current motion will be aborted.

da(θ_{steer},θ_{turret})[0,1] Sets external state variables to the value of θ_{steer} and θ_{turret}. Returns 1 upon success, 0 upon failure.

`disconnect-robot(id)[0,1]` Requests the connection with the robot with `id` to be closed. Returns 1 upon success, 0 upon failure.

`dp(x,y)[0,1]` Sets external state variables to the values of x and y. Returns 1 upon success, 0 upon failure.

`sp(trans,steer,turret)[0,1]` Sets the translation, steering and rotational velocities of the robot (in units of $\frac{1}{10}$ inch within the range 0 to 200, and $\frac{1}{10}$ degree per second within the range 0 to 450). Returns 1 upon success, 0 upon failure.

Sensing

`get-cp()` `[0,1]` Updates external state with the current compass reading of the robot in units of $\frac{1}{10}$ of a degree. Returns 1 upon success, 0 upon failure.

`get-ir()` `[0,1]` Updates external state with the values of the active infra-red sensors. Returns 1 upon success, 0 upon failure.

`get-rc()` `[0,1]` Updates external state with the current configuration of the robot (x, y, θ_{steer}, θ_{turret}). Returns 1 upon success, 0 upon failure.

`get-rv()` `[0,1]` Updates external state with the current translation, steering and turret velocities of the robot. Returns 1 upon success, 0 upon failure.

`get-sn()` `[0,1]` Updates external state with the values of the active sonar sensors. Returns 1 upon success, 0 upon failure.

`gs()` `[0,1]` Updates external state with the current x, y, θ_{steer}, θ_{turret}, translation velocity, steering velocity, turret velocity, CPU voltage, motor voltage, infra-red, sonar and compass values. Returns 1 upon success, 0 upon failure.

Acting

`lp()` `[0,1]` Stops all the motors of the robot. Returns 1 upon success, 0 upon failure.

`pa(tpa,spa,rpa)` `[0,1]` Moves the motors of the robot to the absolute positions specified (`tpa` is a translational step, `spa` and `rpa` are steering and rotational steps respectively) at the speed specified by `sp` at the acceleration specified by `ac`. Depending on the value of `timeout` motion may terminate before the motors have moved the specified distances. Returns 1 upon success, 0 upon failure.

`pr(tpr,spr,rpr)` `[0,1]` Moves the motors of the robot by the distances specified (`tpr` is a translational step in units of $\frac{1}{10}$ in. within the range -32000 to 32000, `spr` and `rpr` are steering and rotational steps respectively, in units of $\frac{1}{10}$ degree within the range -32000 to 32000) at the speed specified by `sp` at the acceleration specified by `ac`. Depending on the value of `timeout` motion may terminate before the motors have moved the specified distances. Returns 1 upon success, 0 upon failure.

`st()[0,1]` Brings the robot to a controlled stop with appropriate accelerations. Returns 1 upon success, 0 upon failure.

`tk(speech-stream)[0,1]` Sends the character string `speech-stream` to the robot's voice synthesiser. Returns 1 upon success, 0 upon failure.

`vm(tv,sv,tv)[0,1]` Moves the robot at the velocities specified (`tv` is the desired translation velocity in units of $\frac{1}{10}$ inch per second within the range -240 to 240, `sv` and `rv` are the desired steering and turret velocities respectively, in units of $\frac{1}{10}$ degrees per second within the range -450 to 450). Returns 1 upon success, 0 upon failure.

`ws(trans,steer,turret,timeout)[0,1]` Waits for the specified motor(s) to stop within the given `timeout` period. Returns 1 upon success, 0 upon failure.

`zr()[0,1]` Aligns the steering and turret motors with the front of the robot, and sets external state variables x, y, θ_{steer}, and θ_{turret} to zero. Returns 1 upon success, 0 upon failure.

4.3.2 Controller functions

The following functions execute and monitor behaviours.

`constrain(behaviour,constraint,timeout)[0,1]` This routine (see Figure 4.7) functions as the central control process of the architecture. It is called with references to a `behaviour` (the behaviour's activate routine), a `constraint` (the behaviour's execution function), and a `timeout`.

It instigates a *fork* system call, running the constraint function, monitoring timeout, and handling all signals within the parent, and running the behaviour itself within the child.

At all times the parent process must check timeout, default, and signal conditions. If default or timeout occurs, the parent sends a signal to the child process, and waits for the child to die before continuing (allowing for safe termination of the current behaviour before a new one starting). If the child process terminates successfully before this time, it sends a 'success' signal to the parent, if it recognises default internally it sends a 'failure' signal.

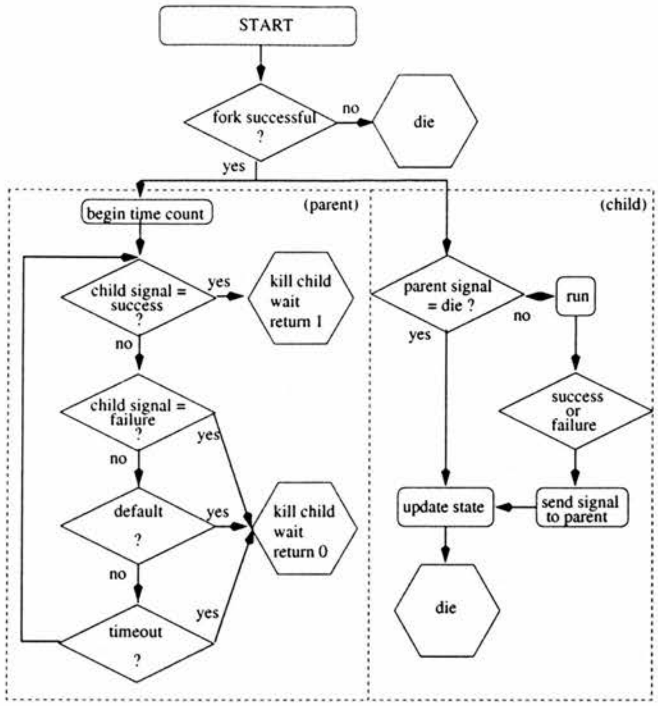

Figure 4.7: constrain(behaviour,constraint,timeout)[0,1]

timeout(behaviour,timeout)[] A general-purpose timeout function (see Figure 4.8) designed to operate within the *parent* of a *fork* process, which takes the name of a **behaviour** and a **timeout**, and terminates the behaviour once the timeout has expired.

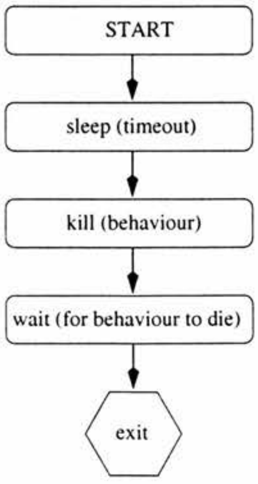

Figure 4.8: timeout(behaviour,timeout)[]

4.3.3 Memory functions

mem::save(filename,data-list)[0,1] Writes the selected **data-list** to **filename** returning 1 on success and 0 otherwise.

mem::retrieve(filename)[0,reference] Attempts to load all the data from **filename** into a **reference** to a data-list returning 0 on failure.

4.3.4 Acts

The low-level commands listed above (section 4.3.1) were combined into **acts** which could be recruited to form more complex behaviours.

Sensing

clean-cp()[theta] This routine simply divides the compass reading (given in $\frac{1}{10}$ degree by **get-cp**) by 10, and returns the value (**theta**) as an integer.

compass-discrepancy(desired)[theta] This function (see Figure A.1) calculates and returns the discrepancy (**theta**) between the current bearing and a **desired** bearing.

det-progress(x,y,frequency)[δ] This function determines whether over a period of time (**frequency**), the Nomad is approaching or receding from the coordinates x,y. Using the equations:

$$distance_i = \sqrt[2]{(x_{initial} - x_{home})\ 2 + (y_{initial} - y_{home})\ 2}$$

$$distance_f = \sqrt[2]{(x_{final} - x_{home})\ 2 + (y_{final} - y_{home})\ 2}$$

$$\triangle = distance_f - distance_i$$

It returns $\triangle$, which is negative if the Nomad is approaching x,y.

get-angle(x,y)[θ] This routine (see Figure A.2) calculates and returns a new steering angle (θ), based on the current position, and the position of the target (x,y). The external state variables **max-theta** and **min-theta** are used in the case where θ falls outside the range **min-theta** $< \theta <$ **max-theta**.

`get-speed(stopping-distance)[speed]` Returns a value for translational velocity using the equation:

$$speed = (\frac{sn_{15} + sn_0 + sn_1}{3} - stoppingdistance) * 5$$

`stopping-distance` is an optional parameter, the default value is 20 inches.

The external state variables `max-speed` and `min-speed` were adopted in the case of speed falling outside the range: `min-speed` < `speed` < `max-speed`.

`get-swerve()[`θ`]` Returns a steering angle based on proximity to lateral objects (see Figure A.3). This routine causes the Nomad to gently swerve away from objects visible by the sonar sensors to either side.

`ir-object-left()[0,1]` Determines whether an object is detected by the left-side IR sensors. Returns 1 if an object is present, otherwise returns 0 (see Figure A.4).

`ir-object-right()[0,1]` Determines whether an object is detected by the right-side IR sensors. Returns 1 if an object is present, otherwise returns 0 (see Figure A.5).

`object-left(threshold)[0,1]` Determines whether an object is detected by the left-side sonar sensors (see Figure A.6). Returns 1 if an object is below `threshold` distance, otherwise returns 0.

`object-right(threshold)[0,1]` Determines whether an object is detected by the right-side sonar sensors (see Figure A.7). Returns 1 if an object is below `threshold` distance, otherwise returns 0.

`obstacle(threshold)[0,1]` Returns 1 if an obstacle is detected within (optional `threshold`) distance or the default value of 20 inches, 0 otherwise (see Figure A.8).

`power()[0,1]` Determines current power levels `cpuvolt`, and `motorvolt`, returning 0 when low, and 1 otherwise.

`recovery-crit(x,y,range)[0,1]` This function assesses whether the Nomad is near (default is 50 inches, but can be modified using the optional parameter $\frac{\texttt{range}}{10}$) the defined position x,y, using the equation:

$$distance = \sqrt[2]{(x_{home} - x_{current})\ 2 + (y_{home} - y_{current})\ 2}$$

It returns 1 when `distance` < `range`, 0 otherwise.

`smooth-cp()[mean]` This function eliminates spikes in the reading of the magnetic compass by calling `get-cp` ten times and eliminating any reading which differs from the previous or subsequent readings by more than 5 degrees, and then returns the mean of the remaining values as an integer. It *does not* update external state.

Orientation and alignment

`align-left(distance)[1]` Aligns the Nomad to a rear surface (see Figure A.9), returning 1 on completion.

`align-rear(distance)[1]` Aligns the Nomad to a rear surface (see Figure A.10), returning 1 on completion.

`align-right(distance)[1]` Aligns the Nomad to a right-hand surface (see Figure A.11), returning 1 on completion.

`conv-to-degrees(rads)[theta]` This function (see Figure A.12) is passed a value in radians (`rads`) and returns the equivalent value in degrees (`theta`).

`conv-to-radians(theta)[rads]` This function (see Figure A.13) is passed a value in degrees (`theta`) and returns the equivalent value in radians (`rads`).

`fine-align(bearing)[1]` This function (see Figure A.14) is called subsequently to `turn-to-bearing(bearing)[disc]` and rotates base and turret very slowly ($1\ \frac{inch}{second}$) towards the desired `bearing`, returning 1 on completion.

`orient-ahead()[1]` Aligns the front sonar (sn_0) to the greatest amount of free space. The function returns 1 on completion.

`orient-left(distance)[1]` Moves the Nomad to a position `distance` inches away from a left-side surface (see Figure A.15), returning 1 on completion.

`orient-rear(distance)[1]` Moves the Nomad to a position `distance` inches away from a rear surface (see Figure A.16), returning 1 on completion.

`orient-right(distance)[1]` Moves the Nomad to a position **`distance`** inches away from a right-side surface (see Figure A.17), returning 1 on completion.

`turn-to-bearing(bearing)[disc]` This function (see Figure A.18) rotates turret and base to ± 1 degree of the supplied **`bearing`**, returning the final discrepancy (**`disc`**) on completion.

Avoidance

`avoid(threshold)[]` **`threshold`** is an optional parameter, the default value is 20 inches. This function is designed to run as a *child* process and implements (see Figure A.19) an eternal avoid loop unless killed.

`rotate()[]` Rotates to left or right whilst an obstacle is detected (see Figure A.20). The function exits when sensor readings indicate that no object is within a range of 15 inches from an arc of the front three sonar sensors (sn_{15}, sn_0, sn_1).

Wall-following

`follow-left(distance)[]` A wall-following loop for use as a *child* process (see Figure A.21). Allows the Nomad to shadow obstacles at a range of **`distance`** (an optional parameter, the default is 22 inches).

`follow-right(distance)[]` A wall-following loop for use as a *child* process (see Figure A.22). Allows the Nomad to shadow obstacles at a range of **`distance`** (an optional parameter, the default is 22 inches).

`get-adjustl(distance)`[θ_{steer}] Returns the steering angle for following a right-side object. **`distance`** is an optional parameter, the default is 22 inches (see Figure A.23). The external state variables **`max-adjustl`** and **`min-adjustl`** are used in the case of θ_{steer} falling outside the range **`min-adjustl`** $< \theta_{steer} <$ **`max-adjustl`**.

`get-adjustr(distance)`[θ_{steer}] Returns the steering angle for following a left-side object (see Figure A.24). **`distance`** is an optional parameter, the default is 22 inches. The external state variables **`max-adjustr`** and **`min-adjustr`** are used in the case of

θ_{steer} falling outside the range `min-adjustr` $< \theta_{steer} <$ `max-adjustr`.

`wall-monitor(`x`,`y`,frequency)[0,1]` This function is used when wall following is invoked as a recovery procedure for navigation. It accepts the x and y coordinates of a goal location, checks current x and y over a period of `frequency` and determines whether either $\triangle x$ or $\triangle y$ are increasing. If both are increasing 1 is returned, otherwise 0.

Navigation

`go(`x`,`y`,range)[0,1]` A navigation routine (see Figure A.25) designed for use independently or as a *child* process. The algorithm moves the Nomad directly from its current position to within the specified `range` of (x,y).

`go-monitor(`x`,`y`,timeout,range,frequency)[0,1]` This function (see Figure A.26) serves to constrain `go` when invoked as a *parent* process. It accepts five compulsory arguments: x and y are the goal coordinates; `timeout` is the function's timeout period in seconds; `range` within which recovery is possible; `frequency` is that of checking progress towards the goal position. It returns 1 if the child process exits successfully before `timeout` has expired, or if upon timeout the Nomad is within recovery distance (`range`), 0 is returned in any other case.

Learning and retracing vector routes

The first learning implementation The following routines (the `corridor` and `direction` objects) were designed as part of the first attempt to allow the Nomad to learn a privileged vector in a corridor environment (see section 5.2.3 for a description of this attempt).

The direction object

Each bearing which the Nomad attempts is instantiated as an independent programming object — a `direction` (see Figure 4.9) which stores information relating to `theta, inverse-theta`, and the `time` spent moving along that bearing.

Initialisation

A direction is instantiated by invocation of `new Direction(theta)[reference]`, which creates a new direction object with bearing `theta` and returns a reference to the object. This function automatically initialises the object *via* the command `dir::init(theta)[]`.

The Direction object	
Variables	Functions
theta inverse-theta time	init(theta)[reference-to-dir] run()[]

Figure 4.9: Initialisation of a direction object

`dir::init(theta)[]` Initialises the new object with `theta, inverse-theta`, and `time` which is initialised as 0.

Functions

`dir::run()[]` Runs the direction object (see Figure 4.10) and updates the private state variable `time` on completion. The general library function `obstacle()[0,1]` is used to check for obstructions.

The corridor object

Upon start-up a `corridor` object was created by invocation of the command `new Corridor` which stored all data relating to the environment, in combination with routines for creating, modifying, and running `directions`.

Initialisation

`new Corridor(theta)[reference]` Creates a corridor object, with an initial bearing of `theta`, and returns a reference to the new object. The routine automatically initialises the new object using `corridor::init(theta)[]`.

`corridor::init(theta)[]` Is called by `new Corridor(theta)[reference]` and creates the private state variables depicted in Figure 4.11. A new `corridor` is initialised with the current compass reading `theta` which is a seed for the creation of a set of `potential-dirs` — a list of references to `direction` objects, initialised with bear-

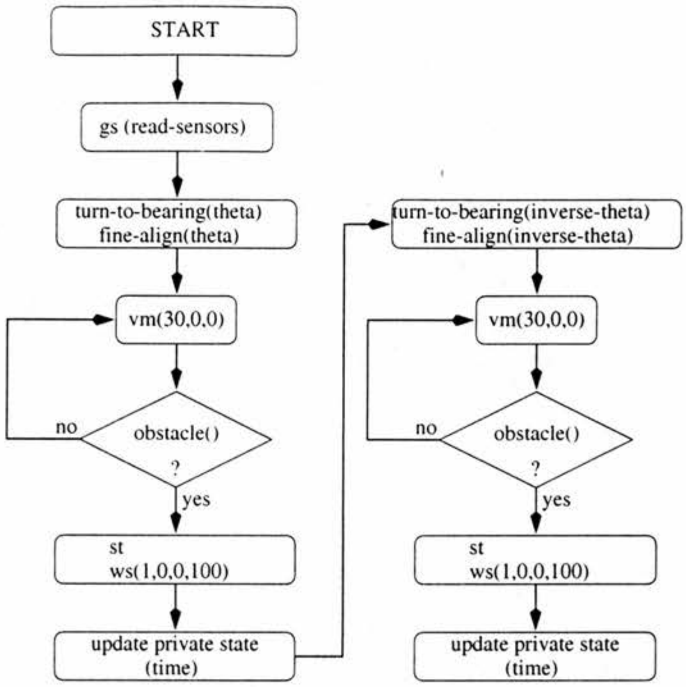

Figure 4.10: dir::run()[]

ings of **theta, u1, u2, d1**, and **d2**, which vary from **theta** by ± 30 and 60 degrees respectively. The list of references to directions **experienced-dirs**, and the variable **best-dir** are initially undefined.

Functions

The following functions can be invoked by a corridor-object, and relate to function activation, and modification of data within the object.

`corridor::create-var(theta, variation)[]` This function accepts a bearing **theta**, and an integer value **variation**, it then updates the private list **potential-dirs** by creating three new directions with bearings of **theta, theta + variation** and **theta - variation**.

`corridor::get-dir()[0,reference-to-dir]` This function selects a direction at random from the list **potential-dirs**, returning a reference to that direction or 0 if no unexperienced directions remain.

`corridor::get-best-dir()[]` Updates the privates state variable **best-dir** with the bearing in which most time has currently been spent. This becomes the seed for a new

The Corridor object	
Variables	Functions
theta theta + 30 (u1) theta + 60 (u2) theta - 30 (d1) theta - 60 (d2) potential-dirs = (new Direction(theta), new Direction(u1), new Direction(u2), new Direction(d1), new Direction(d2)) experienced-dirs best-exp	init(theta)[reference-to-dir] create-var(theta,variation)[] get-dir()[0,reference-to-dir] get-best-dir()[]

Figure 4.11: Initialisation of a corridor object

set of `potential-dirs`.

`corridor::run()[]` Looks up the private state variable **`potential-dirs`**, runs each one in a random order using the routine **`dir::run()[]`**, and updates the private state variable `experienced-dirs`.

The final learning implementation The following programming objects were developed from those described above and were designed as part of the final vector learning implementation (described in section 5.2.4). They are an integral part of the behaviours `learn` (see section 4.3.5) and `retrace` (section 4.3.5) and deal with the creation and usage of representations of regions of physical space (`sectors`), and different compass bearings (`vectors`) within those regions.

The vector object

Each `vector` which the Nomad has experienced, or will attempt, is instantiated as a programming object, possessing encapsulated functions and data relating to its specific bearing, success rate, and x and y coordinates.

Initialisation

A `vector` is instantiated and initialised with a set of private state variables (see Figure 4.12), by a `sector` object invoking the '`new Vector`' command.

The Vector object	
Variables	Functions
theta (xi)	init(theta)[]
initial x (yi)	run()[time]
initial y (xf)	go-back()[0,1]
final x (yf)	go-to-edge()[0,1]
final y	out-of-range()[0,1]
compass value (cp)	am-i-centered()[0,1]
compass discrepancy	
time	near-obj-det()[0,1]
calls	near-behind-det()[0,1]
number of successful calls	
number of failed calls	success()[0,1]
centered	invert(theta)[inverse]
obj-behind	print
reached edge	

Figure 4.12: Initialisation variables for a vector object.

new Vector(theta)[reference-to-vec] This function creates a new **vector** object with angle **theta** within the calling package, and returns a reference to the new object. The routine automatically initialises the new object using **vec::init(theta)[]**.

vec::init(theta)[] Is called by **new Vector(theta)[reference-to-vec]** and initialises the new object with the variables depicted in Figure 4.12.

Functions

The following functions can be invoked by a vector-object, and relate to activation, and modification of data within the object.

vec::invert(theta)[inverse] Accepts an angle (**theta**) as an argument, and returns the inverse of the angle.

vec::am-i-centered()[0,1] Determines whether Nomad position falls within the range:

$$(x_{initial} - error) < x_{current} < (x_{initial} + error)$$

and,

$$(y_{initial} - error) < y_{current} < (y_{initial} + error)$$

where **error** equals 1 inch. If so 1 is returned, otherwise 0. This function updates the

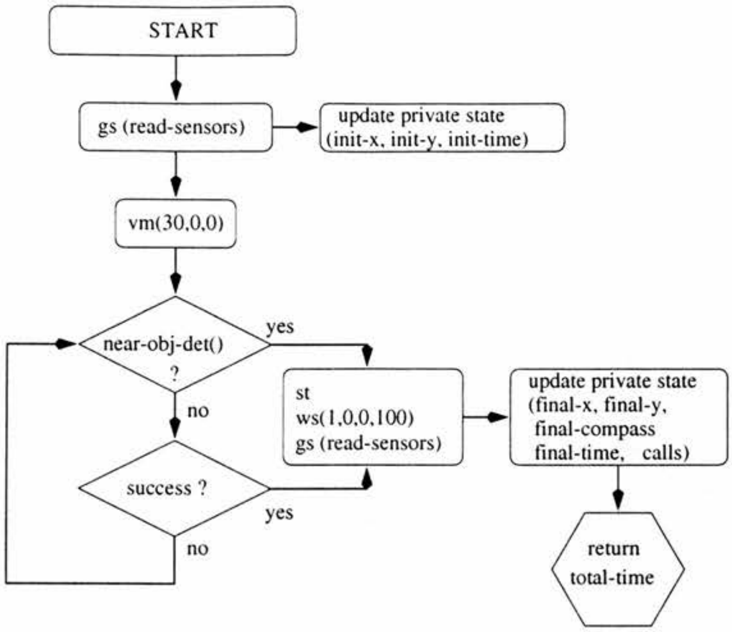

Figure 4.13: vec::run()[time]

private state variable `centered` accordingly.

`vec::near-obj-det()[0,1]` Determines whether any object can be detected at close range by *either* an arc of the front five sonars (sn_{14} (15in.), sn_{15} (20in.), sn_0 (20in.), sn_1 (20in.), sn_2 (15in.)) or an arc of the front three IR detectors (ir_{15}, ir_0, ir_1) and returns 1 if this is the case and 0 otherwise. The private state variable `object` is updated accordingly.

`vec::near-behind-det()[0,1]` Determines whether an object can be detected to the rear of the Nomad at very close range using an arc of the three rear IR detectors (ir_7, ir_8, ir_9), and returns 1 if an object is detected and 0 otherwise. The private state variable `near-obj-behind` is updated accordingly.

`vec::run()[time]` Activates the behaviour: the Nomad moves along `vector(`θ`)` until either an object is detected or the success criterion is met (see Figure 4.13) .The function returns the total time spent moving along the vector and update the private state variables `initial-x, initial-y, final-x, final-y, final-theta, calls,` and `time`.

`vec::go-back()[0,1]` Returns the Nomad to the origin of the `vector` (see Fig-

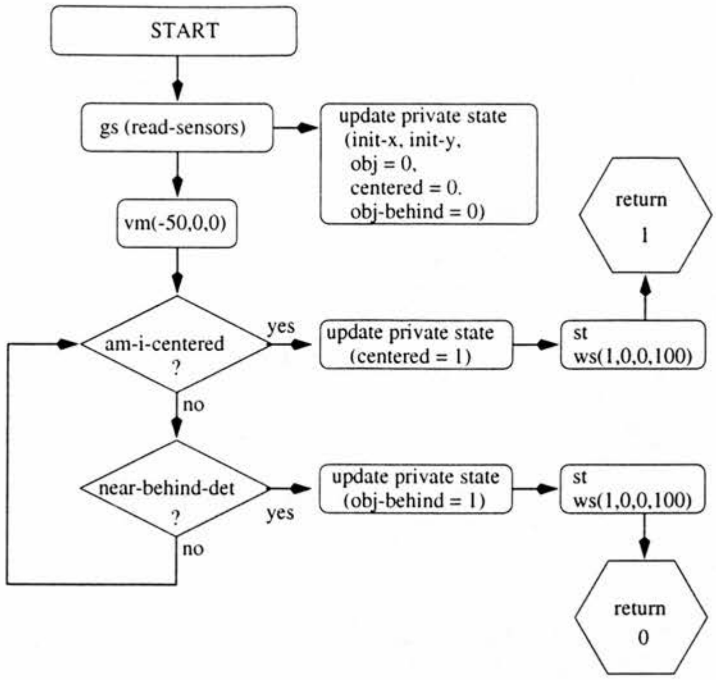

Figure 4.14: vec::go-back()[0,1]

ure 4.14), returning 1 on success, 0 otherwise. This function update the private state variable `centered` accordingly.

`vec::success()[0,1]` Determines whether the current `vector` has met the success criteria. If the current compass reading differs from the angle (`theta`) with which the `vector` was initialised by $\pm$ 15 degrees, the function returns 1, otherwise 0. The function update the private state variable `success` accordingly.

`vec::go-to-edge()[0,1]` This function moves the Nomad to the furthest point achieved by the `vector` (see Figure 4.15). It updates the private state variable `reached-edge` and returns 1 on success, otherwise 0 is returned.

`vec::out-of-range()[0,1]` Assesses whether the current position of the Nomad falls outside the range: initial-x < current-x < final-x and initial-y < current-y < final-y, returning 1 if this is the case and 0 otherwise. The private state variable `reached-edge` is updated accordingly.

`vec::print()[]` Prints all private state variables to a file for debugging purposes.

The sector object

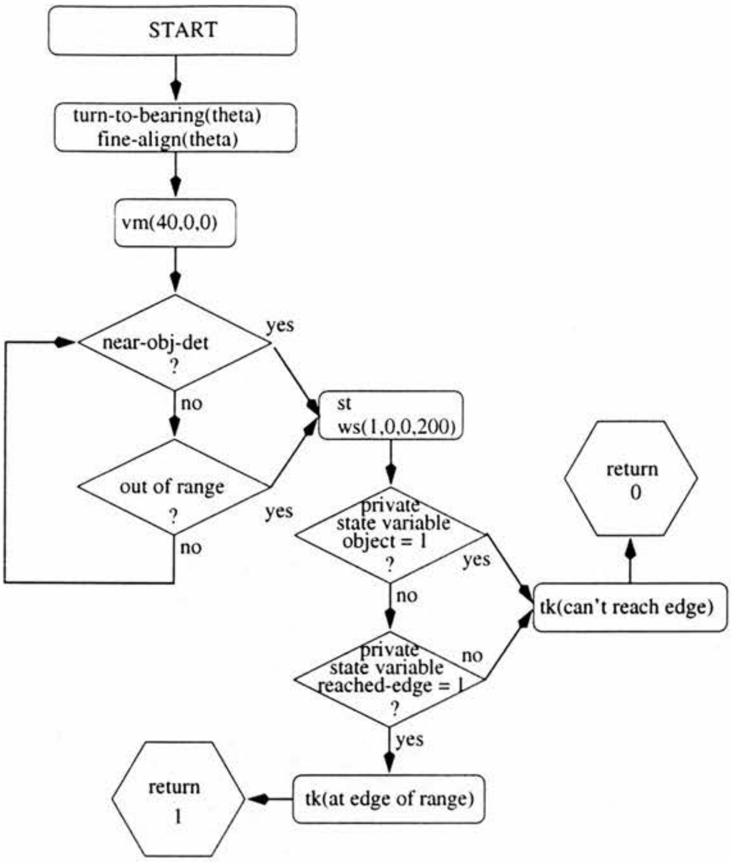

Figure 4.15: vec::go-to-edge()[]

Every novel, separable region of physical space which the Nomad encounters is instantiated as a programming object, possessing encapsulated functions combined with data (see Figure 4.16) relating to all relevant information about that sector (e.g. vectors, coordinates *etc*).

Initialisation

A sector is instantiated and initialised with a set of private state variables (see Figure 4.16), by the controller invoking the '`new Sector`' command:

`new Sector(theta)[reference-to-sec]` This function creates a new `sector` object, with an initial bearing of `theta`, and returns a reference to the new object. The routine automatically initialises the new object using `sec::init(theta)[]`.

`sec::init(theta)[]` Is called by `new Sector(theta)[reference-to-sec]` and creates the private state variables depicted in Figure 4.16. A new `sector` is initialised

The Sector object	
Variables	Functions
min-x = undef	init(theta)[reference-to-sec]
min-y = undef	bin-sonars()[]
max-x = undef	select-vec()[vec-ref,0]
max-y = undef	
	learn ()[0,1]
reached-edge = undef	det-suc-vec-list()[]
tolerance = undef	ord-suc-vec()[]
variation = 30	get-best-fail()[]
	get-best-vec ()[]
current-compass (cur-cp)	dimensions()[]
initial-theta = theta	go-to-edge-state()[0,1]
initial-theta + variation (u1)	nav-origin()[0,1]
initial-theta + (2*variation) (u2)	
initial-theta - variation (d1)	
initial-theta -(2*variation) (d2)	
potential-vecs = (new Vector(init-cp), new Vector(u1), new Vector(u2), new Vector(d1), new Vector(d2))	
experienced-vecs = ()	
successful-vecs = ()	
best-fail = ()	
best-vec = ()	
sonar-bin = (sn0 ... sn15)	

Figure 4.16: Initialisation variables for a Sector object

with the current compass reading (`theta`) which is a seed for the creation of a set of `potential-vecs` — a list of references to vector objects with bearings of `u1, u2, d1,` and `d2`, which vary from `theta` by ± `variation` and ± 2 * `variation`. The list of references to `experienced-vecs` and the variables `best-fail,` and `best-vec` are initially undefined, as are `min-x, min-y, max-x, max-y` and `reached-edge`. The list `sonar-bin` is initialised with the values of all 16 sonars.

Functions

The following functions can be invoked by a sector-object, and relate to function activation, and modification of data within the object.

`sec::bin-sonars()[]` Is used at initialisation of the new `sector`. Five calls are made to `get-sn()[0,1]` the values are averaged and stored in the private state variable `sonar-bin`.

`sec::select-vec()[0,vec-ref]` Selects a `vector` *at random* from the private list

potential-vecs returning either a reference to the selected **vector**, or 0 if no unattempted vectors remain.

sec::learn()[0,1] Is the central learning algorithm (see Figure 4.17), it cycles through the list of **potential-vecs** updating **experienced-vecs** on each iteration. If all vectors are attempted without an unrecoverable error occurring **sec::dimensions()[]** is used to update **min-x, min-y, max-x,** and **max-y**, the routines **sec::det-suc-vec-list, sec::ord-suc-vec, sec::get-best-fail,** and **sec::get-best-vec** are used to determine which is the most promising **vector**, and 1 is returned. If an unrecoverable error occurs 0 is returned.

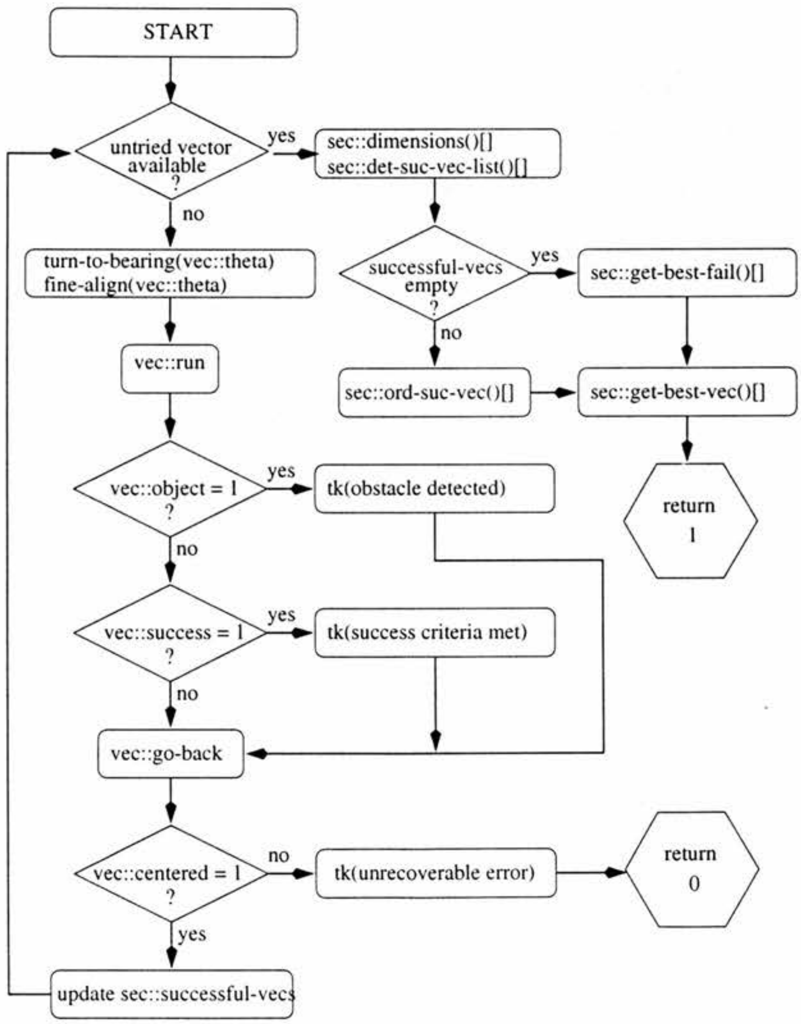

Figure 4.17: sec::learn()[0,1]

`sec::det-suc-vec-list()[]` Constructs a private state variable (`successful-vecs`) which is a list of references to vectors which have met the success criterion.

`sec::ord-suc-vec()[]` Modifies the private state variable `successful-vecs` by ordering the successful directions in terms of the total time spent moving along that `vector`.

`sec::get-best-fail()[]` Is called if no vectors have met the success criteria. In this eventuality the routine constructs a private state variable `best-fail` which is a list of references to all experienced vectors in order of locomotion time on each.

`sec::get-best-vec()[]` Consults the private state variables `successful-vecs` or `best-fail` and updates the variable `best-vec` with a reference to the most successful experienced `vector`.

`sec::dimensions()[]` Calculates the extremities of the `sector` in Cartesian space by querying each of the vectors listed in `experienced-vecs` for their initial and final x and y coordinates. The function updates the private state variables `min-x, min-y, max-x, max-y`.

`sec::go-to-edge-state()[0,1]` Moves the Nomad to the edge of the `sector` along the most successful `vector`, using the vector's own routine `vec::go-to-edge()[0,1]`, returning 1 on success, and 0 otherwise. The private state variable `reached-edge` is updated accordingly.

`sec::nav-origin()[0,1]` Is used to move the Nomad from any point in space within the current state to the initial position within that state. The function returns 1 on success, and 0 otherwise.

4.3.5 Behaviours

Behaviours functions are classified in terms of: initialisation, constraint, activation, state modification, and data storage (printing).

Align

This behaviour is an orthotaxic mechanism designed for use near the home position and ensures that the Nomad is positioned at exactly the same position in physical space at the beginning of all runs, and also sometimes between consecutive behaviours within a task.

Initialisation `align::init()[0,1]` Align may only be run if the internal state variable `near-home` is true in which case 1 is returned, otherwise 0.

Constraint `align::constraint()[0,1]` The only constraint acting on `align` concerns battery power. If `cpuvolt` or `motorvolt` fall too low during execution of the behaviour 0 is returned, otherwise 1.

Activation `align::run(distance)[]` Utilises the library routines `orient-ahead()[1]`, `orient-rear(distance)[1]`, `orient-left(distance)[1]`, `orient-right(distance)[1]`, `align-rear(distance)[1]`, `align-left(distance)[1]`, and `align-right(distance)[1]` to position the Nomad an exact `distance` from surrounding surfaces.

State modification `align::mod-state(status,time,timeout)[]` updates the internal state variables `home`, `last-behaviour`, `last-behaviour-status`, `last-behaviour-timeout` and `last-behaviour-time-taken` along with all external state variables.

Print `align::print()[]` Writes all positional and temporal information to a file for debugging purposes.

Navigate

Initialisation `navigate::init()[1]` The behaviour `navigate` can be run subsequently to any other behaviour whether it exited successfully or not, ensuring that it can be used both to navigate outward from home, and then back again. This function currently returns 1 at all times.

Constraint `navigate::constraint(x,y,timeout,range,frequency)[0,1]` The general constraint function for `navigate` runs as the parent component of a fork system call and consists of the `go-monitor(x,y,timeout,range,frequency)[0,1]` function called with the parameters of (x,y) of the desired position, a `timeout` which is varied depending on the distance to be navigated (default 400 seconds), a recovery `range` of 130 inches, and a checking `frequency` of 1 second, in combination with the `power()[0,1]` library function. If either `power()[0,1]` or `go-monitor(x,y,timeout,range,frequency,)[0,1]` return 0, then this is passed on by `navigate::constraint(x,y,timeout,range,frequency)[0,1]`.

Activation `navigate::run(x,y,range)[0,1]` The function which executes `navigate` operates as a child process, using the library function `go(x,y,range)[0,1]`. The default value of `range` is 10 inches.

State modification `navigate::mod-state(status,time,timeout)[]` On completion `navigate` updates the internal state variables `last-behaviour`, `last-behaviour-status`, `last-behaviour-time-taken`, and `last-behaviour-timeout`, and all external state variables.

Print `navigate::print()[]` This routine prints all coordinate and time information to a file for debugging purposes.

Dead-reckon

This behaviour (referred to as `dead-rec`) is used for dead-reckoning back to the home position after invocation of `retrace`.

Initialisation `dead-rec::init()[1]` The behaviour `dead-rec` can only be run subsequently to `retrace` whether it was successful or not. This function currently returns 1 at all times.

Constraint `dead-rec::constraint(x,y,timeout)[0,1]` The constraint function for `dead-rec` runs as the parent component of a fork system call and consists of the `go-monitor(x,y,timeout,range,frequency)[0,1]` function called with the parameters of (x,y) of the desired position, and `timeout` with the default values of range and frequency, in combination with the power()[0,1] library function. If either go-monitor(x,y,timeout,range,frequency,)[0,1] or power()[0,1] return 0, then this is passed on by dead-rec::constraint(x,y,timeout)[0,1].

Activation dead-rec::run(x,y,range)[0,1] The function which runs dead-rec operates as a child process, using the library function go(x,y,range)[0,1]. The default value of range is 10 inches.

State modification dead-rec::mod-state(status,time,timeout)[] On completion dead-rec updates the internal state variables last-behaviour, last-behaviour-status, last-behaviour-time-taken, last-behaviour-timeout, near-home, and sector, and all external state variables.

Print dead-rec::print()[] This routine prints all coordinate and time information to a file for debugging purposes.

Axis

This behaviour uses the objects corridor and direction (section 4.3.4 in an attempt to get the Nomad to learn the long axis of a corridor environment (see section 5.2.3 for a full description).

Initialisation axis::init()[0,1] The behaviour axis can be run only if the previous behaviour was align and it exited successfully. In this case the function returns 1, otherwise 0.

Constraint axis::constraint()[0,1] Only two constraints are imposed on axis — this function returns 1 if timeout has been exceeded or power()[0,1] indicates that the batteries are running low, otherwise 0 is returned.

Activation axis::run()[] Is a complex function (see Figure 4.18) designed to enable the Nomad to discover the long axis of its corridor environment. It iterates 4 times using the corridor::run()[] routine, narrowing down towards the privileged vector within the environment through use of the corridor::get-best-dir and `corridor::create-var(theta,variation)[]` routines with the value of `variation` becoming progressively smaller (30, 10, 5, and finally 2 degrees).

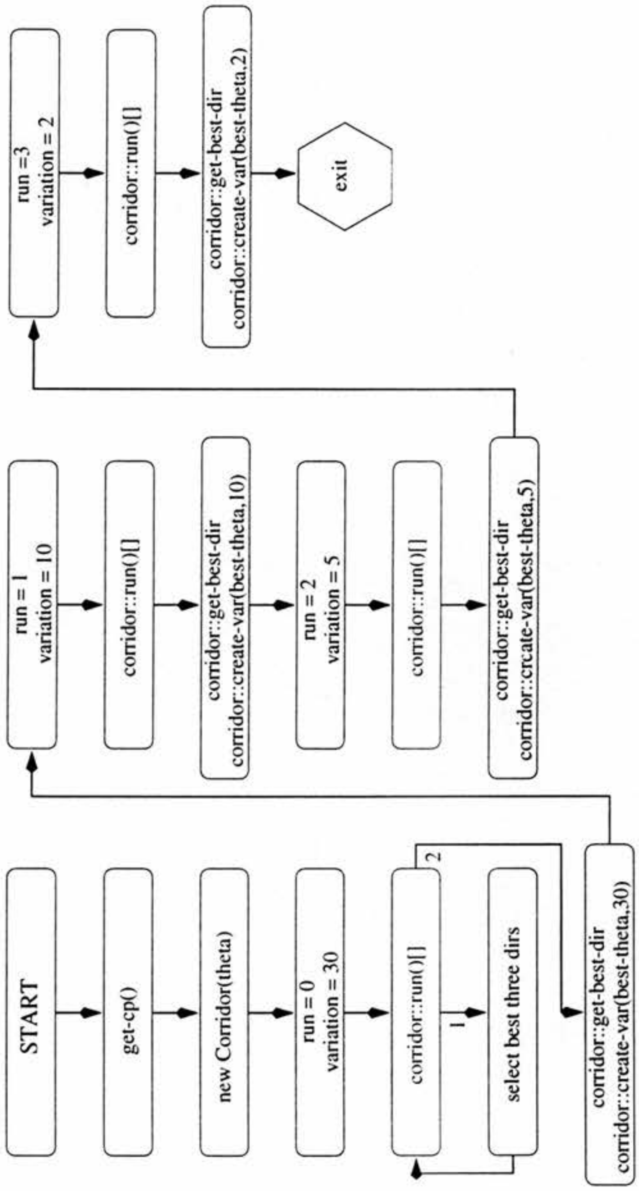

Figure 4.18: axis::run()[]

State modification `axis::mod-state(status,time,timeout)[]` Updates internal state variables `last-behaviour`, `last-behaviour-status`, `last-behaviour-time-taken`, `last-behaviour-timeout`, private variable `privileged-dir`, and all external state variables.

Print `axis::print()[]` Writes all time, corridor, and direction related information to a file for debugging purposes.

Learn-route

This behaviour (referred to as `learn`) uses the objects `sector` and `vector` (section 4.3.4) and causes the Nomad to learn about its physical environment through a combination of low-level sensing and trial and error (see section 5.2.4).

Initialisation `learn::init()[0,1]` The behaviour `learn` can only be run if the previous behaviour was `align` and it exited successfully. In this eventuality the function returns 1, otherwise 0.

Constraint `learn::constraint()[0,1]` As much of the monitoring of the success and default conditions of `learn` is accomplished within the behaviour itself only two constraints are imposed by the controller. This function returns 1 if `timeout` has been exceeded or `power()[0,1]` indicates that the batteries are running low. Otherwise 0 is returned.

Activation `learn::run(number)[0,1]` Is a complex function (see Figure 4.19) enabling the Nomad to learn about `number` of discrete physical spaces within its environment. It iterates `number` of times calling the `sec::learn()[0,1]` and `sec::go-to-edge-state()[0,1]` functions on each iteration. The function returns 1 if all iterations have been completed successfully, 0 otherwise.

State modification `learn::mod-state(status,time,timeout)[]` Updates internal (`last-behaviour`, `last-behaviour-status`, `last-behaviour-time-taken`, `last-behaviour-timeout`, `current-sector`, and `sector-list`) and all external state variables (rather than the private state variables belonging to its component sectors and vectors).

Print `learn::print()[]` Writes all coordinate, time, `sector`, and `vector` related information to a file for debugging purposes.

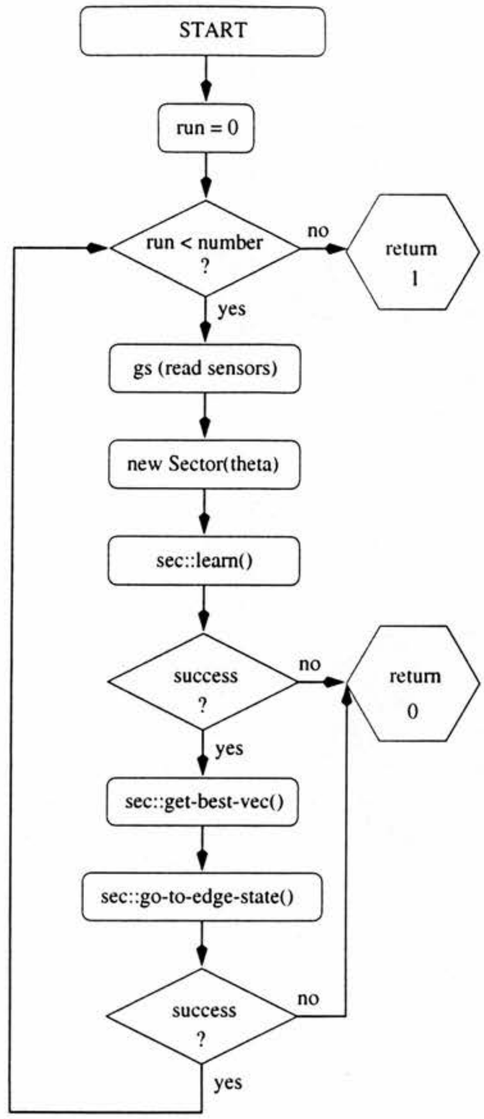

Figure 4.19: learn::run(number)[0,1]

Retrace-route

This behaviour (referred to as **retrace**) allows the Nomad to retrace a previously learned route, following each **sector** on its outward journey, but using *dead-reckoning* to return to its home position. The result looks very much like the navigation of the desert ant (*Cataglyphis*).

Initialisation retrace::init()[0,1] The behaviour **retrace** can run at any time as long as the internal state variable `home` indicates that the Nomad is positioned

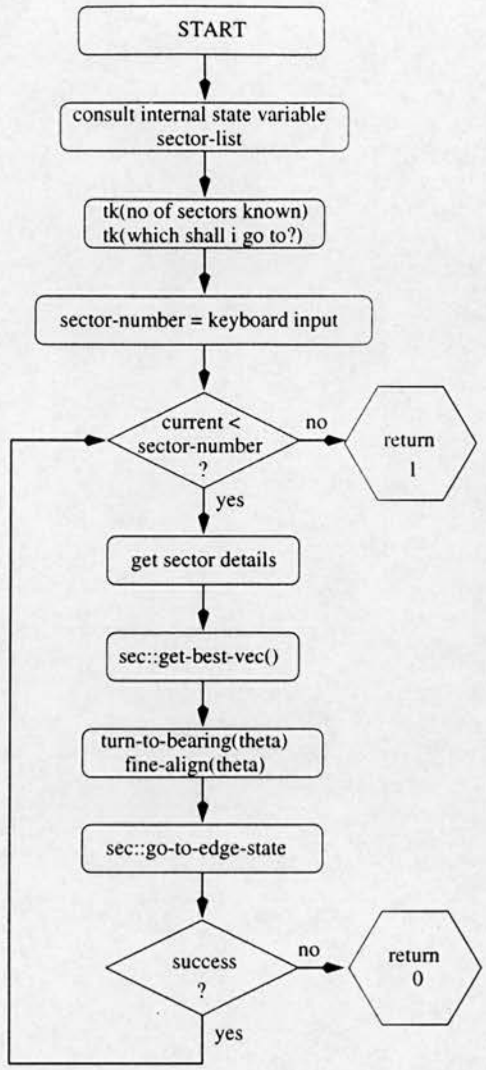

Figure 4.20: retrace::run(sector-number)[0,1]

at the base location.

Constraint `retrace::constraint()[0,1]` The only constraints operating over `retrace` are a `timeout`, and battery power (`power()[0,1]`). If the `timeout` has not been exceeded and power is sufficient the routine returns 0, otherwise 1.

Activation `retrace::run(sector-number)[0,1]` This function (see Figure 4.20) queries the user for `sector-number` (keyboard input) and causes the Nomad to move to the origin of `sector-number` through all previously learned sectors on the route. The function returns 1 on successful completion, 0 otherwise.

State modification `retrace::mod-state(status,time,timeout)[]` Updates the internal state variables `last-behaviour`, `last-behaviour-status`, `last-behaviour-time-taken`, `last-behaviour-timeout`, and `sector`.

Print `retrace::print()[]` Prints all coordinate and time information for each sector to a file for debugging purposes.

4.3.6 An example task

A *task* is a sequence of behaviours, designed to achieve some non-trivial aim. The example here (see Figure 4.21) runs the three behaviours `align`, `learn(4)`, `dead-rec`, `align`, `retrace(3)`, `dead-rec`, and `align` in sequence (see sections 5.3 & 5.4 for a full exposition).

Upon start-up the Nomad aligns to its home position, providing a fixed point in physical space from which all subsequent `sectors` will be defined. If this behaviour exits successfully, the next behaviour `learn(4)` is invoked and the Nomad begins to try different vectors, creating `4` sectors of physical space, each characterised by a different set of possible vectors and one privileged `vector` determined by trial and error. Successful execution of `learn(4)` is followed by `dead-rec`, the Nomad uses its on-board odometry to navigate back to *near* its starting position. If this is successful `align` is invoked to position the Nomad exactly at home.

The next stage of the sequence involves `retrace(3)` which causes the Nomad to retrace its previously learned route up to the origin of the third `sector` it has learned, whereupon once more `dead-rec` is called to navigate back to the region of home, and `align` once more positions the Nomad exactly at the origin of the task.

4.4 Summary: a synthetic architecture

The architecture described herein is designed to be a platform for further development. As stated earlier, 'incrementing by design' (McGonigle, 1991), is a core feature of the stance adopted herein. The potential for the refinement of a number of critical control elements has been provided, facilitated by the modular and hierarchical organisation

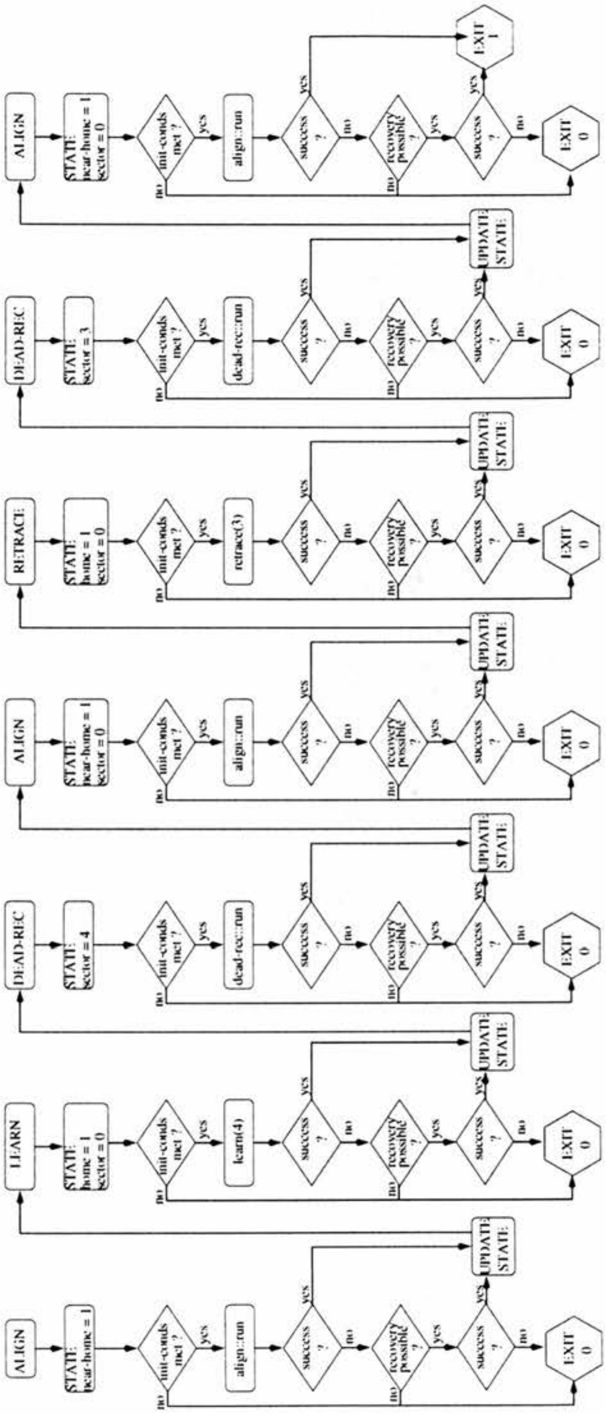

Figure 4.21: An example task: align, learn, dead-rec, align, retrace, dead-rec, align

of the system. Scope exists for much greater development of the architecture than was possible in the current research. The following design features have been incorporated within the architecture; most, if not all, have been constructed to allow further development.

Long-term development of the system is envisaged to involve the extension of learning and self-organisation to action selection and error interpretation and recovery (see section 6.3).

Embodiment and situatedness The architecture controls a Nomad 200 mobile robot in a real world niche[4]. As with all such control architectures the sensory-motor capacities of the robot, and features of the niche, constrain the implementation of behaviours.

Embodiment constraints The clearest example of the use of such constraints was the utilisation of compass error to facilitate the self-selected navigation described in section 5.2.4.

Niche constraints and niche engineering The geometrical structure of the niche suggested, and served to constrain, the method of navigation finally chosen for the system. A 'home' location was provided by construction of a position at one end of a corridor with a unique sensory signature.

Meaning through action The only competence examined in depth — learning to navigate — involves the system segmenting physical space into distinct, internally-represented 'sectors' based on situated action (see section 5.2.4). The meaning of such sectors is therefore derived from the interaction of the system's capabilities and its environment.

Self organisation The navigational implementation described above (section 5.2.4) is inherently self-organised. From a design primitive which is essentially a relatively simple algorithm specifying movement on the basis of current orientation, and an inbuilt economy metric, the system learns about distinct regions of space.

Economy The in-built economy metric is used to choose between experienced 'vectors'. Choosing those which permit the most movement for the least compass error ensures that the system maximises distance travelled whilst simultaneously minimising the scope for navigational error.

[4] Although a suite of simulation software exists for this platform, provided by Nomadic Technologies Inc., it was not used in the present research — largely because of technical difficulties.

State based interpretation Two levels of state-based interpretation are possible within the system as currently implemented.

- Task state space, allows for varying interpretations of signals, events, error, and differential ordering of recovery procedures *etc.* depending on internal context.
- Segmentation of physical space into discrete internally represented sectors (section 5.2.4) based on behavioural capacity, means that the system 'knows' about different areas of the external world. Interpretation of objects and events within these different areas can then occur. Error interpretation and recovery, and later object identification, will depend on an internal context which reflects the external environment.

Hierarchical organisation Although the system is inherently hierarchical, the advantages of this structure are at present only exploited by the designer. Addition of novel, more abstracted and less time-critical control, is envisaged to occur through design. Recombination of behaviours, and interpretation of behavioural default and success should ultimately be subjected to a self-organised learning process.

Serial control The initialisation conditions of each behaviour specify when it can be implemented. These refer to both internal and external criteria. The serial organisation of system behaviour is specified by task demands yet reactivity is retained: specification of default conditions ensures interruptibility and thus general reactivity to changing situations. Again this is an area in need of further investigation. Currently the behavioural syntax (specification of initialisation and default conditions) is pre-installed but learning such a syntax should be a long-range goal.

Error cognisance and recovery Each behaviour's default conditions specify error occurrence. An associated stack of recovery procedures can be implemented in order to recover from error. This aspect of the system has currently received little attention, but preliminary investigations demonstrate that the architecture supports error cognisance and recovery (see section 5.1.4).

Autonomy The system is not merely flexless, but has the capacity to select behaviours based on task demands, implement them and recover from error if necessary. The ability to categorise, and learn to avoid, error could be added to the current system. Extending system autonomy to both the serial organisation of behaviour, and to hierarchical control is a long-range goal of the project.

Life history The system possesses a memory allowing storage of, and later access to, all state information across runs. This provides the potential for the system to become progressively more adaptive over time. Currently memory functions have been exploited only for self-organised navigation, and preliminary examination of error recovery. In the long term, this aspect of the system will be critical, underlying developments in all areas.

As should be clear from the summary above, the architecture strives to incorporate many of the essential features suggested by our synthetic stance and is easily extendible both horizontally (addition and modification of behaviours) and vertically (progressive addition of more abstracted control and learning algorithms).

Chapter 5

Navigation

Navigation was chosen as a suitable competence to engineer for three reasons:

1. It is sufficiently complex to allow testing of the overall architecture in a non-trivial environment.

2. A locative competence is absolutely critical to a locomoting system forming the next layer of the logical hierarchy of design suggested by McGonigle (1991).

3. There exists a history of navigational implementations within the research group on which to build, and with which to compare the current implementation.

This chapter initially describes some preliminary research targeted at navigation without learning — in this case a computationally cheap dead-reckoning competence utilising on-board odometry. The ability to get back home from distant parts of the niche is a vital competence for a locomoting organism. The aim of this part of the implementation was to design a simple and reliable navigational algorithm to achieve this. Odometric information alone was used in order to determine a steering angle back home, this angle was fed into the motors and the Nomad steered, more or less directly, toward home. Detection of obstacles in the path led to navigation being temporarily interrupted by avoidance. No planning was involved in this algorithm. However, as dead-reckoning is notoriously error-prone the task was divided into two stages: navigation toward home, followed by alignment to local landmarks. The combination of these two strategies allowed robust navigation over distances of tens of metres. A corollary to this part of the research was a preliminary investigation of error recovery procedures.

The next phase of research concerned learning to navigate. Using only short range sensing and compass information the task was for the Nomad to learn routes through the niche based on vector information derived from the magnetic compass together with odometry through trial and error learning arbitrated by an inbuilt economy criterion. Two methods were investigated.

A global niche Initially the entire niche space (for this stage the long corridor) was treated as a single entity. Through movement along selected vectors, and their inverse, would the Nomad determine the privileged vector (corresponding to the long axis of the niche) by logging time and distance travelled? Unlawful compass error made this approach impossible and inspired the next phase of the investigation.

Niche segmentation For this phase compass error was utilised to index distance travelled. Now movement along an attempted vector was terminated not only if an obstacle was detected but also once compass error exceeded 15%. Once all vectors had been attempted, the Nomad moved to the final point of the most successful vector and the process was repeated. In this manner segmentation of physical space into discrete internally represented sectors was achieved based dually on compass error and obstacle distribution.

5.1 Navigation without learning

Biological organisms have been shown to utilise a vast range of information sources, both internal and external, for short, middle, and long-range navigation (see McFarland, 1999, for a review). Examples of such strategies are:

Taxes which involve tightly-coupled movement toward, at an angle to, or away from a stimulus source, for example the use of pheromone trails by many insect species (Shorey, 1976).

Compass orientation which involves movement on a particular vector using only idiothetic information (McFarland, 1999).

Dead reckoning (path integration) which requires maintenance of a cumulative record of time spent moving on given vectors. A steering vector towards goal is derived by vector addition. Such a strategy is widespread (Wehner & Srinivasan, 1981; Mittelstaedt & Mittelstaedt, 1982; Etienne *et al.*, 1994, for examples).

Piloting which involves the use of known stimuli located proximally to a goal (Gallistel, 1990), or a sequence of known stimuli for progressive goal approach (Deutsch, 1960).

Homing which involves the use of landmarks with known geometric relationship to a goal to determine a steering course (McFarland, 1999). It is often claimed that this form of navigation requires some form of 'cognitive map'.

The type of navigation initially demanded of the Nomad was a form of dead reckoning, or path integration. In fact, vector addition is not explicitly involved. Rather the system logs starting position in Cartesian coordinates, on-board odometry later provides the current updated position from which a steering vector can be obtained. Cumulative odometric error means that following runs of some 10s of metres, the goal location can be shifted by as much as 2 metres. In order to circumvent this difficulty, *orthokinesis* was used to accurately relocate the Nomad at the goal position.

5.1.1 The task

The first navigational competence developed was an odometrically-based routine for dead-reckoning towards a given (x,y) position, in more or less a straight line. In fact, navigation was always toward the niche-engineered home location at this stage, as the Nomad possessed no routines for accurately distinguishing other parts of the niche.

At the beginning of a run alignment routines (`align` see section 4.3.5) were used to position the Nomad exactly at home so that it was straightforward to determine how successful navigation had been over a series of runs. After a period of time spent `avoiding`, or `wall-following` the navigation routine `navigate` (see section 4.3.5) would be called. Now the task was to get back home using onboard odometric data to determine desired heading.

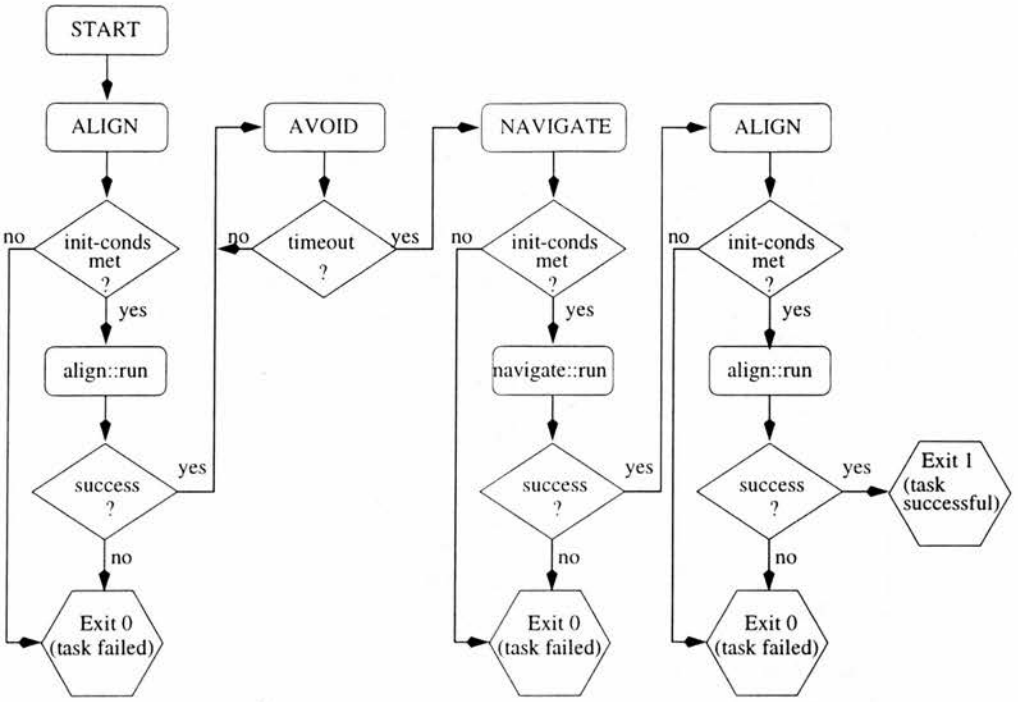

Figure 5.1: The navigation task: align, avoid, navigate, align

Routines and behaviours recruited

Library functions	Behaviours
`avoid(distance)[]` (sec. 4.3.4)	`align` (sec. 4.3.5)
	`navigate` (sec. 4.3.5)

5.1.2 State and the granularity of navigation

As odometric error tends to increase with distance travelled, onboard records of position are generally not sufficiently accurate to allow precise navigation back to a point. Within the biological realm, path integration is often combined with local landmark recognition. The honey bee (*Apis mellifera*), and desert ants (*Formica aquilonia*) both utilise a strategy of dead-reckoning *near* to home, and then matching current retinal image with a stored representation of landmarks (Wehner & Räber, 1979; Cartwright & Collett, 1982, 1983). The desert ant has also been observed to implement a strategy of nest area approach followed by spiralling in circles of ever widening radius until the nest has been located (Gallistel, 1990, p. 61).

Navigation was therefore analogously split into two phases with differing granularity (see Donnett & McGonigle (1991) and McGonigle & St Johnston (1995) for earlier implementations of such a strategy). This served to simplify the development of the main navigation algorithm, and allowed aspects of the niche to aid location finding, rather than depending purely on error-prone odometry. The state-based nature of the architecture greatly facilitated this approach. By defining a task (figure 5.1) as, for example:

$$\texttt{align} \Longrightarrow \texttt{avoid} \Longrightarrow \texttt{navigate} \Longrightarrow \texttt{align}$$

once navigation to the area (defined as a radius of 2 metres) of home had been achieved, the controller switched internal state. The initialisation conditions of `align` ensured that it could be executed only if there was no previous behaviour, or if the previous behaviour was `navigate` and it had exited successfully. This updating of state is critical: whereas during navigation sensor readings are interpreted as objects to be avoided, during the second *orientation* phase (`align` section 4.3.5) sensor readings are interpreted as objects to which the Nomad must approach and orient itself.

5.1.3 The algorithms

Navigation stage

The first implementation of navigation on the Nomad (the behaviour `navigate` described in section 4.3.5) utilised a simple algorithm which calculated desired heading from the disparity between the Nomad's current position in Euclidean space, and the target position (`get-angle(`x`,`y`)[theta]` section 4.3.4). This calculation was repeated twice a second and the new value of θ fed into the motors.

The controller assessed *default* and *timeout* conditions terminating or suspending the navigation behaviour when, or if, appropriate. If the sensor readings indicated that an object was blocking the desired path, `navigate` was temporarily suspended and `avoid` took control. The Nomad could not get stuck in an avoid loop as avoid itself would timeout after 20 seconds, at which point the Nomad would attempt to continue to `navigate`.

Repeated interruptions of `navigate` by `align` inevitably lead to the occurrence of a *timeout*. For this development stage the time regarded as reasonable for successful navigation between current and goal positions was taken to be 80% of the time taken avoiding from home to the current position. If the navigation behaviour was still in operation after this length of time, a *'timeout'* was said to have occurred and recovery procedures (see below) were recruited.

Orientation stage

A combination of odometric error and the sheer bulk of the Nomad made more accurate positioning using only odometric information impossible (see section 5.1.2) therefore the second phase of navigation involved running the behaviour `align` (section 4.3.5) which implemented a form of *orthokinesis* allowing precise positioning.

If `navigate` exited successfully the internal state variables `last-behaviour-status` and `near-home` would equal 1, and `last-behaviour` would equal '`navigate`' in which case the initialisation conditions for `align` would be met and the behaviour could commence. The Nomad now entered a *locative* stage of navigation where niche-engineered features of the environment were exploited to aid precise positioning. Initially rotating base and turret to the region of greatest free space (`orient-ahead()[0,1]` section 4.3.4), meant that orientation and alignment with respect to side and rear surfaces could occur.

A combination of `navigate` and `align` robustly permitted the Nomad to position itself within 1 or 2 inches of its starting position even following runs of 10s of metres from home.

5.1.4 Navigation and error recovery

The structure of the Nomad's niche (see figure 4.3), long corridors, with separate rooms, and dog-leg corners, meant that the strategy of steering more or less directly toward home, in the absence of planning and prospective control, would often fail. Although primitive, one advantage of this algorithm is that it is rapid, requires no form of spatial map, and does not suffer from the problems associated with planning a route through unknown territory. Furthermore, *timeout* occurs often giving a rich error space in which to test recovery procedures both in and of themselves, and as part of the overall architecture.

In situations where the Nomad became stuck behind a corner, `avoid` would repeatedly interrupt `navigate` leading to a *timeout* — navigational failure due to timeout was assumed always to be the result of blockage between the current position of the Nomad and home[1]. Preliminary investigation of recovery procedures centered on simple recovery strategies such as wall following, and creating a temporary goal location midway between current and home positions.

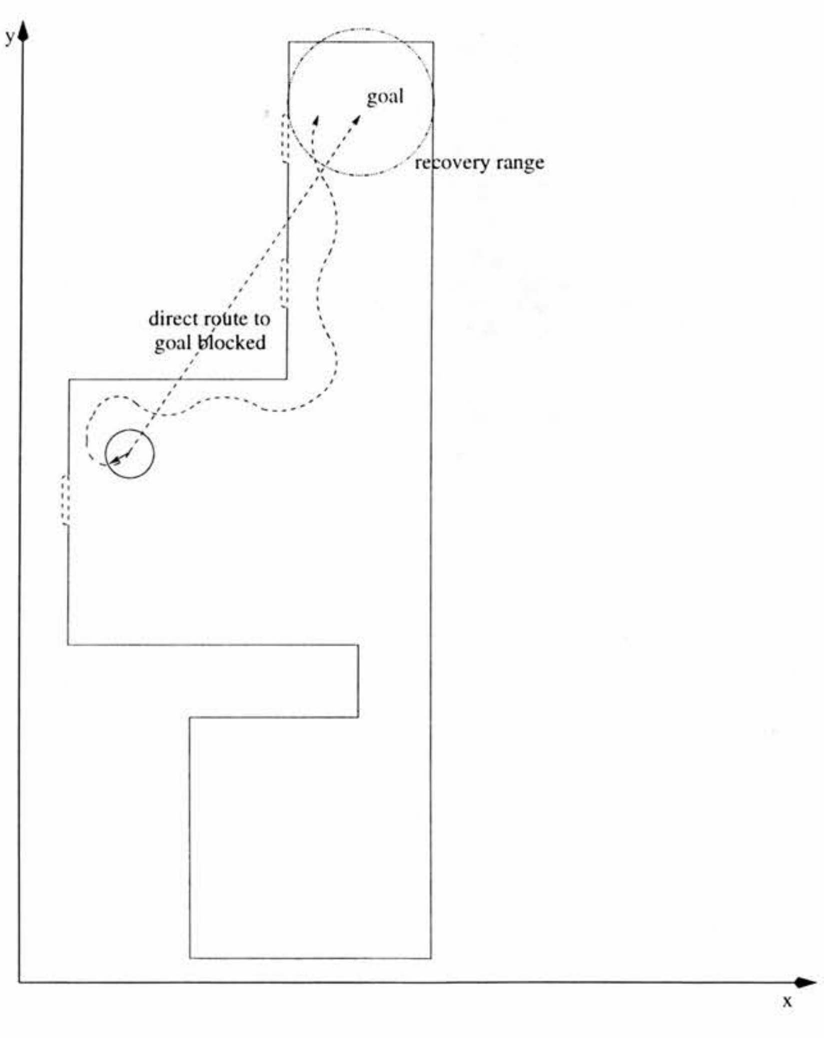

Figure 5.2: A recovery procedure for navigation: wall-following

[1] It must be noted that the Nomad, in its current incarnation is unable to reliably distinguish landmark features of the niche and thus navigational failure resulting from massive odometric error is unrecoverable. In the cases where odometric error became too great no recovery procedures would be successful and the task would return failure.

Wall following The first example of a recovery procedure (Figure 5.2) switches the Nomad into wall-following mode, which continues until either the distance from the goal is less than `recovery-range` in which case `navigate` returns success, or `timeout` has once again occurred (see Figure 5.3) when obviously failure is returned. Given the structure of the niche, long walls with few areas of free space, this has the potential to be a robustly successful strategy.

Obviously a potential difficulty was that wall following might occur in the wrong direction, *i.e.* away from home. This problem was circumvented by addition of a monitoring function (section 4.3.4) that compared current (x, y) with home (x, y) over a small time interval. Because in some cases movement towards home might temporarily result in an increase of:

$$\sqrt[2]{(x_{current} - x_{home})\ 2 + (y_{current} - y_{home})\ 2}$$

wall following is terminated only if both Δx and Δy are increasing. This simple strategy ensured that wall following was always toward home.

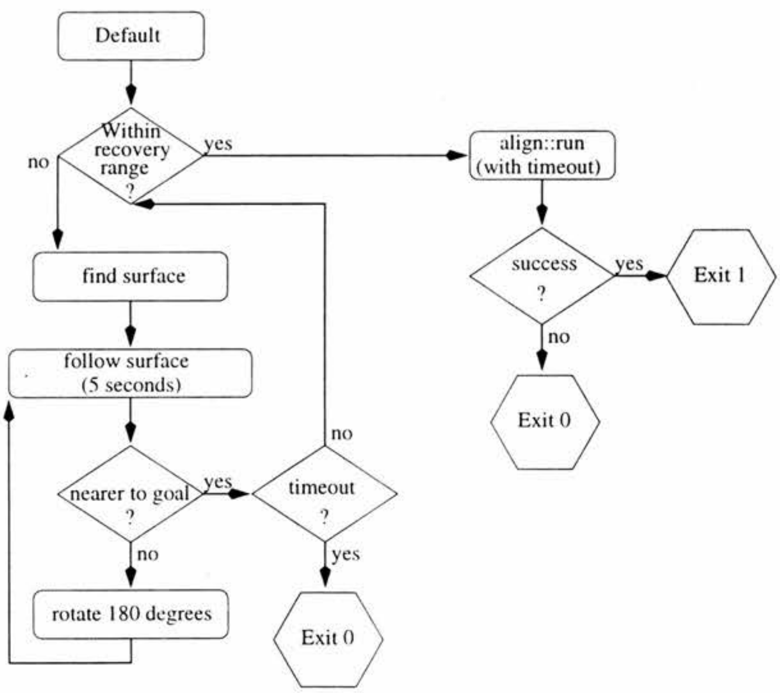

Figure 5.3: A recovery procedure for navigation: Wall-following. The algorithm

Goal creation An alternative strategy involved creation of a new temporary goal (see Figure 5.4). In the case of path obstruction, detouring along either the x or y axes is an effective strategy given the angular construction of the niche. If navigation to this temporary goal was successful an attempt was made to steer directly toward home, if unsuccessful once more goal creation was repeated. This process continued until either the Nomad was within **recovery-range**, or **timeout** occurred (see Figure 5.5).

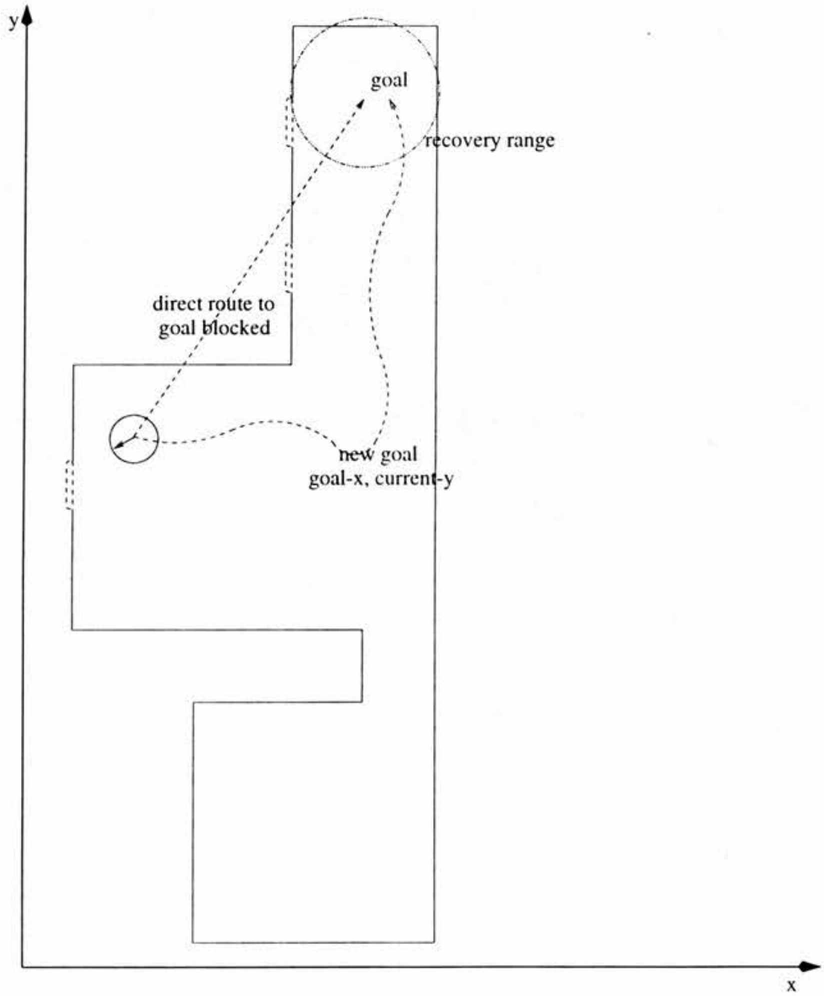

Figure 5.4: A recovery procedure for navigation: Creating a temporary goal

Preliminary investigations with the recovery procedures depicted in Figure 5.3 and Figure 5.5 indicated that the concept of a stack of recovery procedures for a behaviour, each maintaining a record of success and failure, is workable in principle. However, recovery from error was not the central focus of this implementation and no further

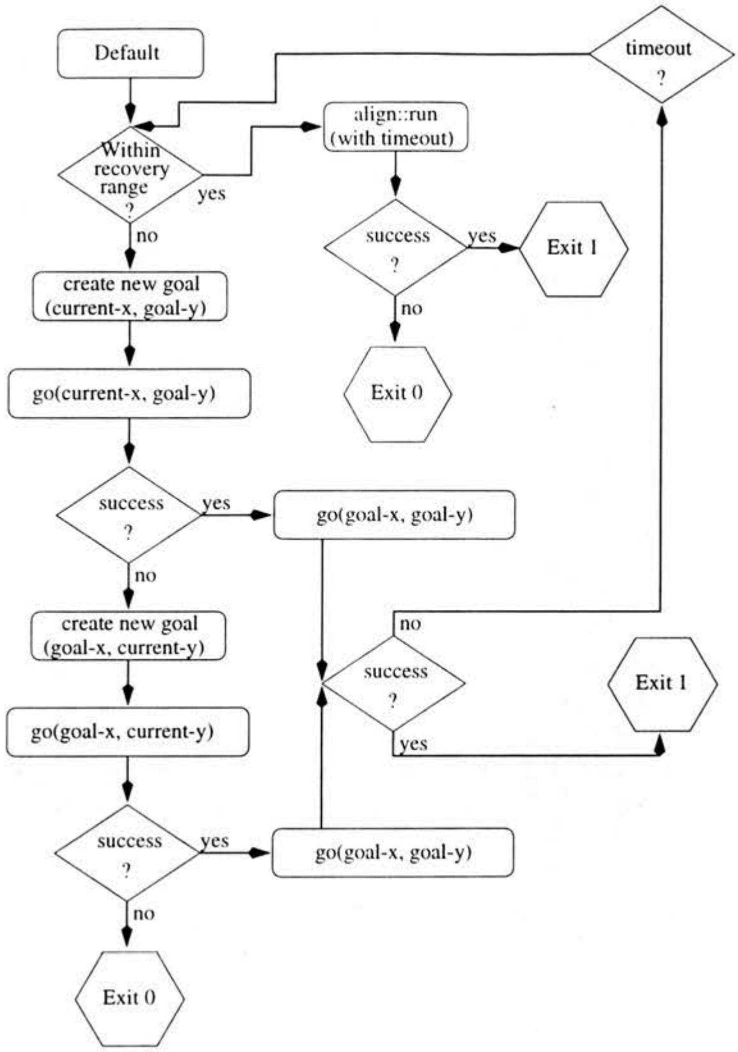

Figure 5.5: A recovery procedure for navigation: Creating a temporary goal. The algorithm

studies were conducted.

5.1.5 Summary

This early phase of navigational development consisted of a task (see Figure 5.1) such as: `align` $\Longrightarrow$ `avoid` $\Longrightarrow$ `navigate` $\Longrightarrow$ `align`. Although not really an adaptive behaviour in the strict sense, for the purposes of this investigation `avoid` could start only if `align` had exited successfully. After a given `timeout` period `navigate` was called and, if `navigate` exited successfully, `align` could be called once more and, if successful, the task would itself return success.

Navigation was split into two distinct phases, *navigation* and *orientation*, analogous to the coarse and fine grained goal seeking behaviour of the desert ant (Wehner & Räber, 1979). The serial, state-based construction of the architecture was critical, allowing for *signal reinterpretation* at different task stages. Niche engineering was utilised to create a recognisable home location from which all runs began. Niche constraints, essentially the corridor environment and the predominance of right-angles, served to simplify recovery procedures. Preliminary investigations showed that the architecture could support multiple stacked recovery procedures, and execute them when required.

5.2 Learning to navigate

5.2.1 Rationale

Although algorithms had been developed which allowed the Nomad to use a combination of the strategies of dead-reckoning, and orientation (see section 5.1.3) to find its way home from distant sectors of its niche, the behaviour obtained lacked the robustness which might be achieved using a learning strategy. Furthermore, the navigation behaviour, as initially implemented, provided little potential for further derivations.

The next phase of research[2] concerned learning to navigate, again constrained by our overriding economy considerations, both in terms of design and internal evaluation criteria. Using only compass information and short range sensing, the Nomad was required to self-select the best trajectories within the niche with the eventual goal being movement along the longest axis of the niche with the least perturbation caused by obstacles, and at the cheapest design and computational cost.

5.2.2 Using a compass for navigation

Compass information provides a global directional frame of reference within which action can occur. Within the biological realm such global compass information can be obtained both internally and externally. Many organisms, for example bacteria, honeybees, and pigeons (Walcott & Walcott, 1982) have been found to possess miniature magnets which might potentially be influenced by the earth's magnetic field. Although

[2] Buoyed by the arrival of a new magnetic compass!

it remains unclear how this information is encoded within the nervous system (McFarland, 1999), such idiothetic compass information does seem to be utilised by a number of organisms. The internal clock sense of honeybees has been shown to be influenced by a variety of magnetic phenomena (Gould, 1980), as is the navigational ability of pigeons (Bookman, 1978). Birds in general seem to be capable of using the declination of the earth's magnetic field, in conjunction with visual information, for navigation (Wallraff, 1978).

Other species, from desert ants (Wehner & Räber, 1979) to starlings (Kramer, 1950), have been found to utilise the polarisation of sunlight by the atmosphere, others, such as the indigo bunting (Emlen, 1972), have been found to utilise rotation of stellar constellations (Schmidt-Koenig, 1979; Wallraff, 1984). Possibly the most common strategy is utilisation of both internal and external compass information, for example constellations and magnetic field are used by many migratory birds (Wiltschko & Wiltschko, 1975, 1976).

This implementation strove to develop 'true' dead reckoning based dually on compass information and odometry, following an initial trial and error learning phase which would demand self-selection of vectors by the system.

The magnetic compass fitted to the Nomad returned the angle from magnetic north in units of one tenth of a degree. A simple function queried the compass for this value (`get-cp()[0,1]` section 4.3.1). Whilst stationary a time-series of the values returned indicated small fluctuations around a central value of $\pm$ 3 degrees with occasional drastic spikes with several hundreds of degrees of error.

These oscillations were overcome simply by taking ten consecutive readings (over approx. 1 sec) disregarding any reading which was more than 5 degrees different from any of the others, and then averaging the remainder of the values. This simple function (`smooth-cp()[compass]` section 4.3.4) served to eliminate spikes and provided a consistent compass value for any single position.

Subroutines (`turn-to-bearing(theta)[0,1]` and `fine-align(theta)[0,1]`, see section 4.3.4) were developed in order to enable the Nomad to orient turret and base to a given (smoothed) compass angle with an accuracy of $<$ 1 degree.

5.2.3 Preliminary investigations — a global niche

The task

Starting from any position along the long corridor of the niche, the Nomad was required to determine the long axis of the corridor by moving along a series of bearings, and their inverse, until an obstruction was detected — progressively narrowing down the set of possible bearings by trial and error, and eventually settling into one privileged direction of movement.

Routines and behaviours recruited

Library functions	Behaviours
`corridor` (sec. 4.3.4) `direction` (sec. 4.3.4)	`axis` (sec. 4.3.5)

The algorithm

The learning task was for the Nomad to use its compass information in order to determine the long axis of the niche (see Figure 4.3). The first attempt to solve this problem utilised the behaviour `axis` (section 4.3.5) and consisted of an algorithm (see Figure 5.6) which moved the Nomad along its current compass bearing until an obstruction was detected, whereupon the Nomad moved along the inverse of the vector until its path was again blocked. A log was maintained of the time spent in free travel along the vector and its inverse.

Initially the starting bearing (`theta`) was used as a seed from which to derive alternative bearings. This initial value was varied by ± 30 and ± 60 degrees. Movement was attempted along each of the resulting five bearings. A log of time and distance information was consulted and an *economy criterion* used to select the best three bearings — those which admitted the longest amounts of uninterrupted movement.

These three bearings were then each tried once more and the best of the three selected. This `direction` then became the seed for a further set of three bearings (± 10 degrees), the best of which became the seed for a further set of three bearings (± 5 degrees), and finally the best of these was chosen as the seed for a set of three

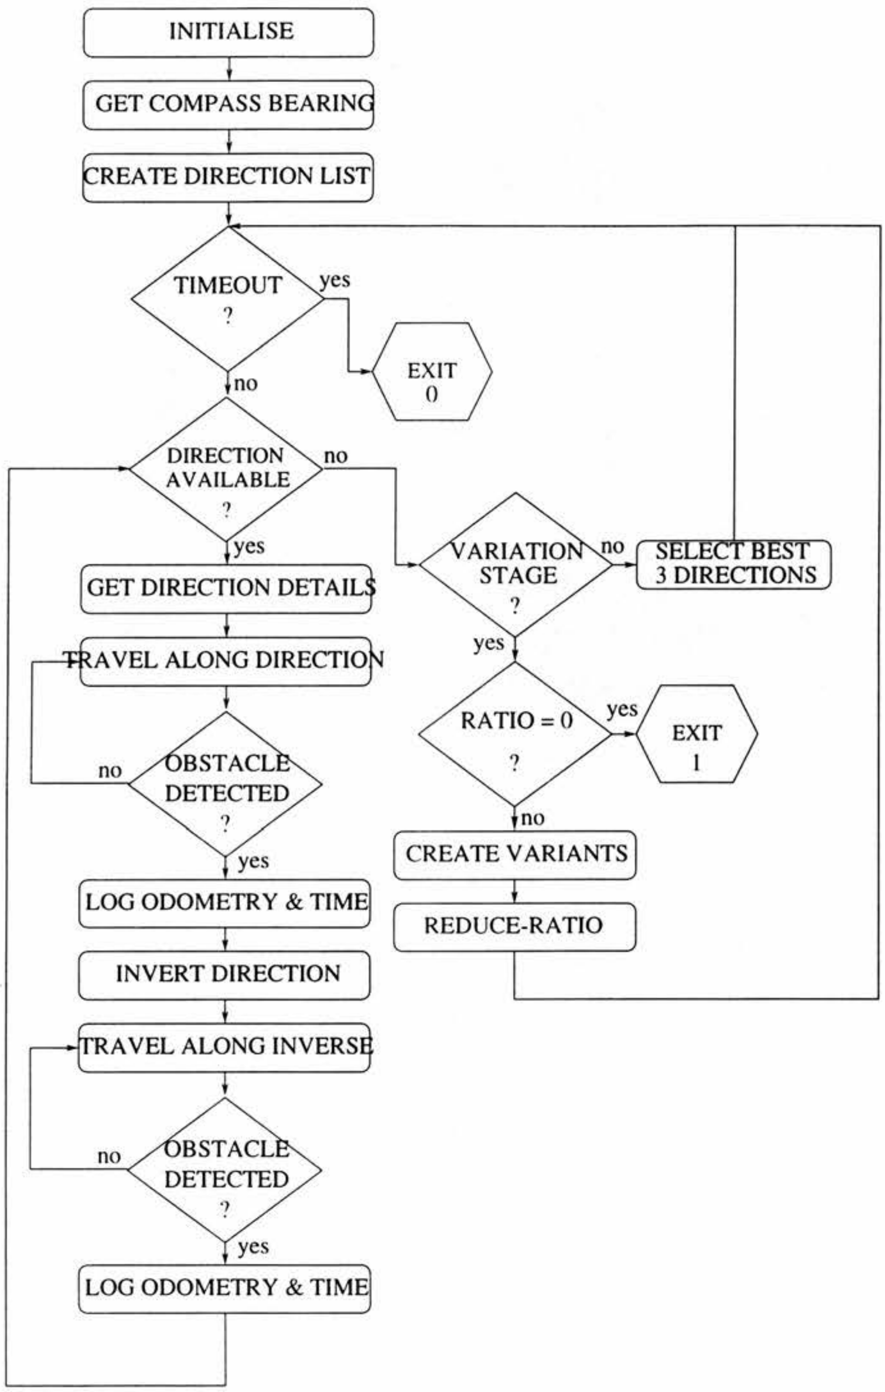

Figure 5.6: The first learning algorithm

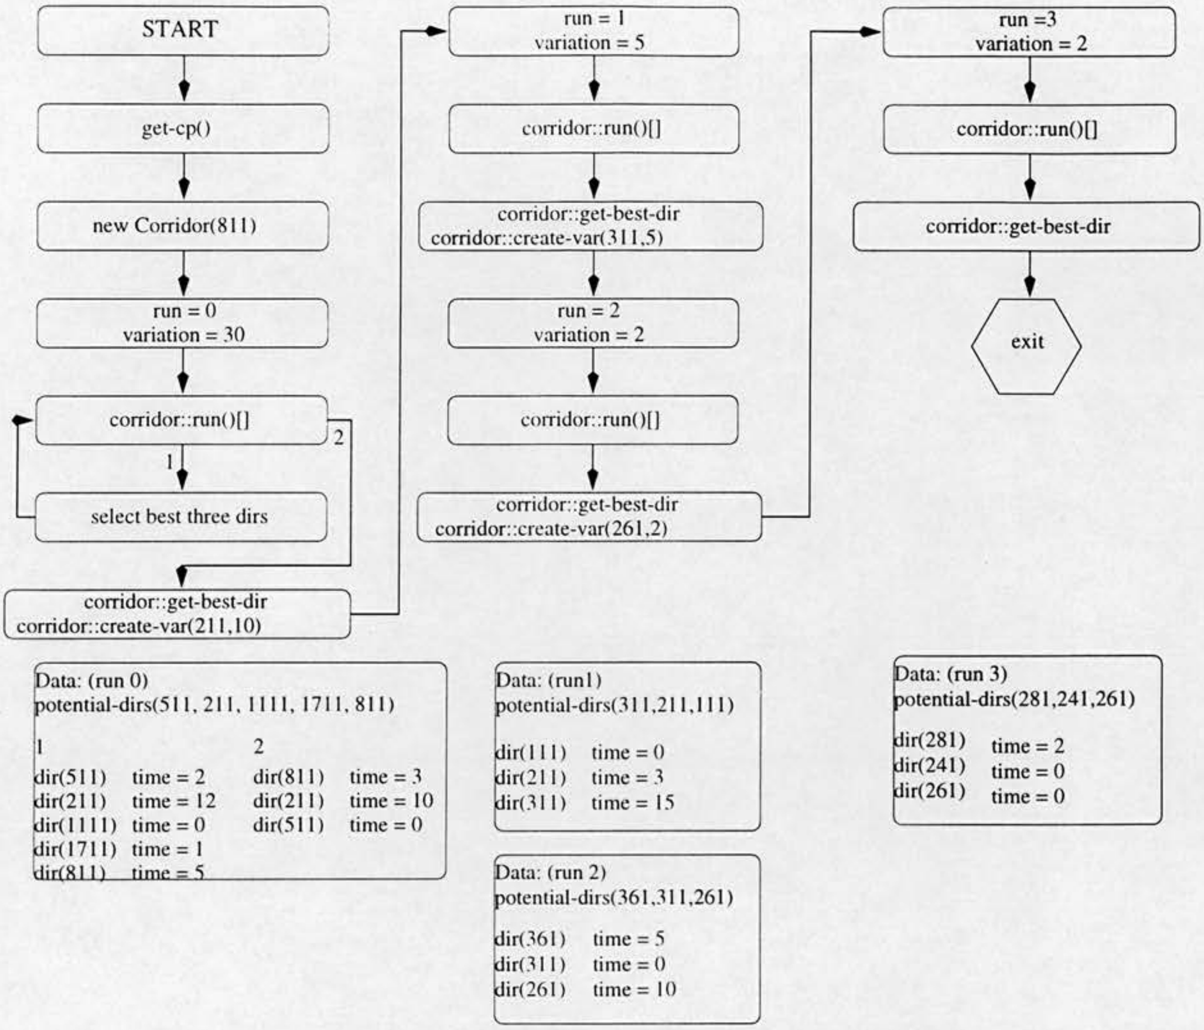

Figure 5.7: A typical run

bearings (± 2 degrees). Although the implementation described here only iterated three times (see section 4.3.5) in theory this process could thence be iterated until no improvements in performance occurred. It was thought that varying the as yet most successful `direction` by progressively smaller amounts (determined by `ratio` and `reduce-ratio`) would result in the Nomad setting upon the longest axis of travel — which should correspond to the corridor of the niche. In this manner it was hoped that, in the absence of both a pre-installed map and long range sensory information, successful navigation in the corridor area of the niche would be obtained.

Initial results

Although the algorithm itself appeared to work as expected (see Figure 5.7 for data from a typical run), a serious problem was that even over the relatively small distance along the corridor of the niche, the compass readings changed quite rapidly. At different points along the axis the readings were too divergent to allow one main free axis to be determined. After travelling along `theta` and discovering an obstacle, when the Nomad attempted to travel back along the bearing (`inverse-theta`) the compass reading was displaced such that movement was instead at an angle to the desired bearing (see Figure 5.8).

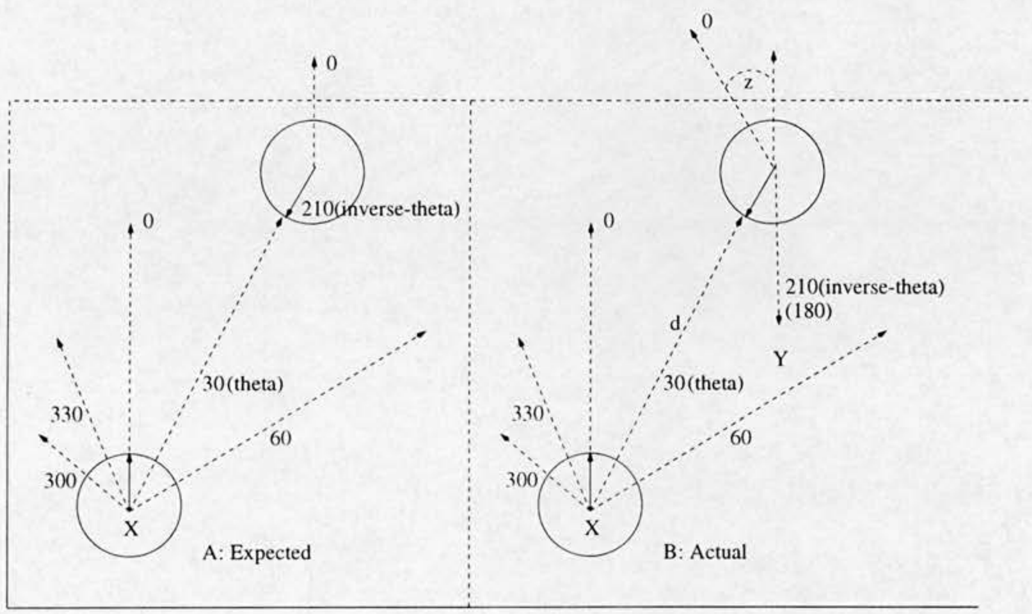

Figure 5.8: Compass error caused movement at an angle to the desired bearing

The left side of the figure depicts what was expected to happen: after travelling along the bearing `theta` (30 degrees) and meeting an obstruction, the library functions `turn-to-bearing(theta)[]` and `fine-align(theta)[]` are called and the Nomad should rotate to the inverse of the original bearing (210 degrees) and start to move back towards the origin x. In reality (right hand figure), due to changes in the compass value after moving the distance (d), 0 degrees has shifted anti-clockwise through the angle `z`. Now when the Nomad attempts to turn to `inverse-theta`, according to the original bearing system it is actually on a bearing of 180 degrees and movement along this bearing results in approach to point Y.

The severity of the compass drift over even apparently very small distances meant that a change in the position of the robot by even a metre or two could occasionally disrupt the returned compass reading by up to 200 degrees. Obviously such large variations were unexpected and quickly became problematic — especially in such a large navigable space.

The initial response to this problem was to see whether or not this drift was lawful — if so, it might have provided unforeseen benefits, for example as a third (along with odometry and time), and independent, distance metric. Such lawful variation might have provided a very useful confirmatory device with respect to actual (as opposed to odometrically derived) distance travelled. Unfortunately this was not to be. The variation of the compass readings seemed to depend primarily on vicinity to the many computers and monitors located around the niche, and (presumed) steel girders holding the building together and did not seem to obey any lawful pattern of variation that could be detected.

The solution to the compass-drift problem led to the development of the algorithm which will be discussed below (section 5.2.4). Breaking up the space into smaller physical regions allowed the Nomad to rely on its compass information within a region, without large-scale movement being compromised by error.

5.2.4 Final implementation — niche segmentation

Rationale and task

Rapid changes in the returned compass value even over small distances meant that the initial approach: learning of the long axis of the niche through trial and error experimentation over the entire space had to be abandoned. The compass, however, provided an invaluable source of information with respect to the orientation of the robot, in that the other sensors (sonar and ir) were simply too inaccurate to provide exact orientation cues in virtually all areas of the niche (with the exception of the niche-engineered home location discussed earlier). Despite quickly mounting compass error upon movement, returned values whilst stationary provided useful orientation information for the Nomad. In a given position the obtained compass reading could be used to reliably position the robot in a desired orientation. Consequently, it was thought that primitives could be installed in the Nomad which would ensure that movement remained within the error tolerance of the compass. By adding a further constraint to trial and error experience of potential vectors, namely termination of movement if compass error reached 15% of its initial value, the movement of the Nomad would never result in unmanageable disorientation. Iteration of the process of vector derivation, experimentation, and self-selection should lead to segmentation of the niche, dually based on the Nomad's experience of obstacle distribution and compass error.

Initially the Nomad was positioned near home, and the orthokinetic behaviour `align` (section 4.3.5) was run to attain an exact, standardised starting position (see Figure 5.9) — it was assumed that all runs would start here providing a constant frame of reference with respect to which all movement in the niche would be interpreted.

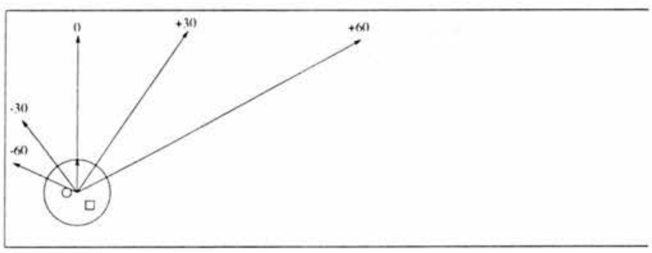

Figure 5.9: Initial position.

From this position the Nomad was required to progress down the long axis of the niche by attempting movement along different `vectors`, segmenting the niche into distinct `sectors` as it went.

Routines and behaviours recruited

Library functions	Behaviours
`mem::save(filename)[0,1]` (sec. 4.3.3)	`align` (sec. 4.3.5)
`sector` (sec. 4.3.4)	`dead-rec` (sec. 4.3.5)
`vector` (sec. 4.3.4)	`learn` (sec. 4.3.5)

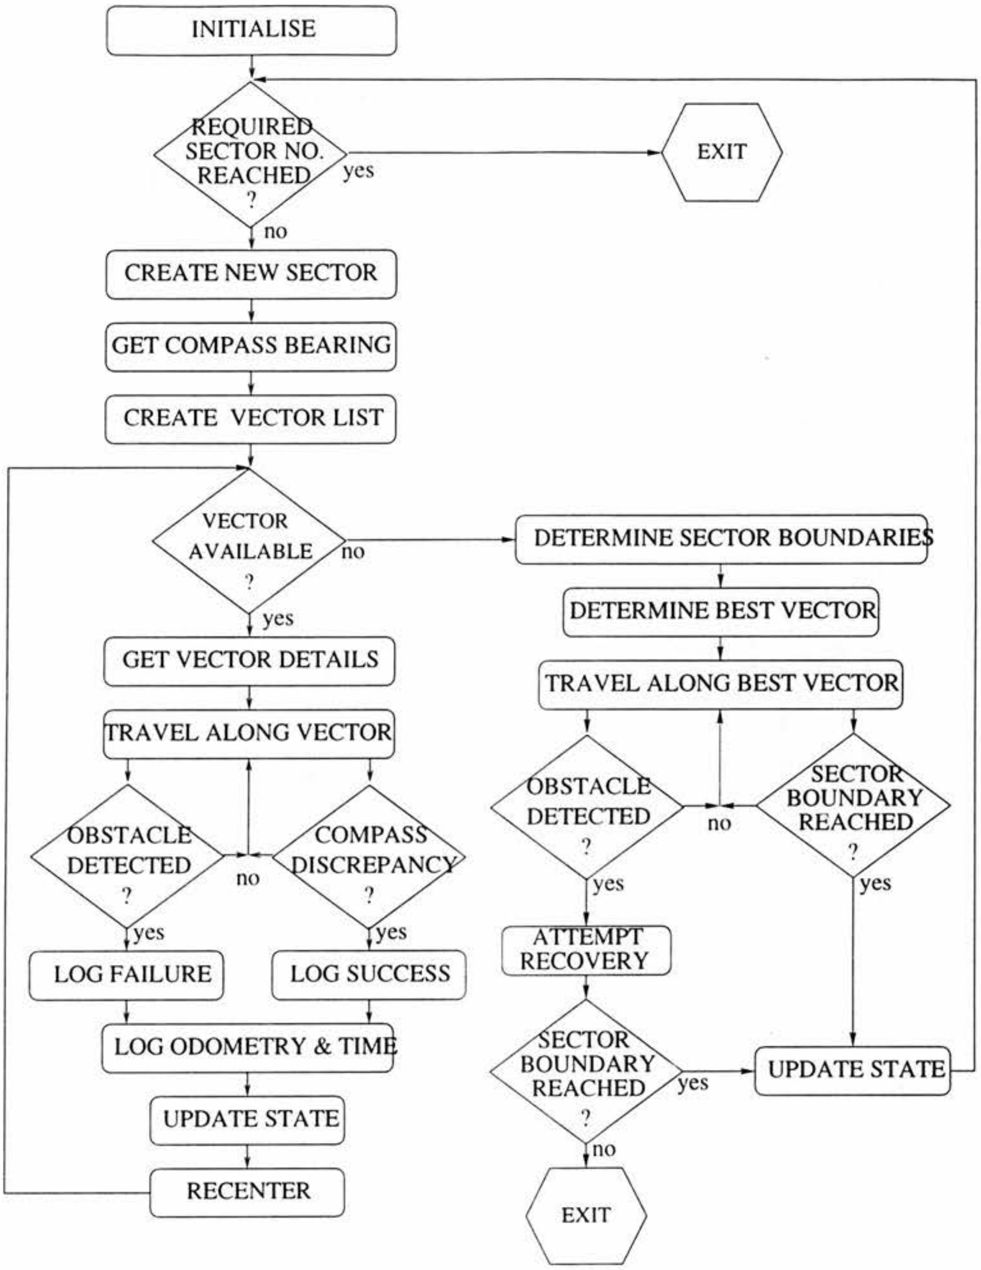

Figure 5.10: The final learning algorithm.

The algorithm

The general structure of the algorithm outlined above (see section 5.2.3) was retained with a number of modifications for this new approach to navigational learning. It was decided that the movement-related compass error would be incorporated in the modified algorithm. As before, the Nomad would begin with its initial orientation acting as a seed for a set of vectors to be derived and attempted, but now success would be defined as an (empirically-determined) unacceptable error build-up on the compass. Assuming that movement along a vector should result in stability of the compass reading, when the value returned by the magnetic compass attained a value differing by more than ± 15 degrees from the expected value `theta` this was taken as an indication that the Nomad had made some progress along the given vector and was thus taken as a criterion of successful movement (odometric and time-based data continued to be collected in order to ensure that some movement had, in fact, occurred and to enable arbitration between multiple `vectors` which resulted in compass error).

The behaviour `learn(number)[]` (described in section 4.3.5) was executed after `align` had exited successfully. This behaviour iterated a `number` of times, on each iteration a new `sector` was instantiated (corresponding to a discrete region of physical space) with a seed angle (`theta`) namely the initial bearing of the robot.

Each `sector` created a list of `potential-vectors` and instantiated each one as a `vector` object. Each `vector` was then executed in a random order (`vec::run()[time]`), with odometric, time and compass data logged at start and finish, along with success status. Upon meeting an obstacle, or the success criterion being met, the Nomad reversed to the origin of the sector using `vec::go-back()[0,1]`. Once a `sector` had exhausted its `potential-vectors` the most successful one was determined either from amongst those which had been successful (where compass error varies by ± 15% from `theta`) or, if no vectors had met the success criteria, by using a time-metric. Whichever `vector` permitted the longest amount of uninterrupted movement was chosen as the privileged `vector` for that `sector` (see `sec::get-best-vec()[]` and `sec::get-best-fail()[]` section 4.3.4).

The routine `sec::go-to-edge-state()[0,1]` then moved the robot to the edge of the current `sector` along the privileged `vector` using the routine `vec::go-to-edge()[0,1]`. This process could then be iterated a given `number` of times (see Figures 4.19 and 5.10).

It is important to note that at the edge of the `sector` (following implementation of `vec::go-to-edge()[0,1]`) along the `vector`(θ) the obtained compass reading is likely to differ substantially from `theta`. This was the problem with the initial attempt at axis learning (section 5.2.3). However this second implementation, by breaking up the niche into smaller units of space, determined either by obstacle layout *or* compass error, circumvents this problem. *Within* a `sector` the compass can be relied upon to give a stable reading. Essentially the compass is *recalibrated* at the transition points between `sectors` in that each possesses its own magnetic frame of reference (see Figure 5.11).

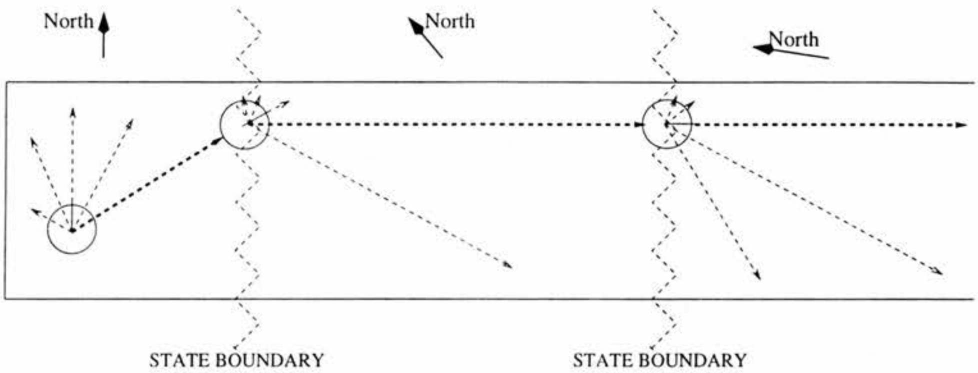

Figure 5.11: After three iterations: transition between sectors. Each sector possesses a unique magnetic frame of reference.

Experimental results

The outcome of the repeated iteration of this algorithm was movement along the free axis of the corridor area of the niche, which was broken down by self-determined internal criteria into discrete physical spaces (see Figures 5.12 & 5.13). The fixed starting position obtained through the orthotaxic mechanisms implemented by `align` meant that information about learned `sectors` could be written to memory and recalled at a later time — the Nomad possesses a *life-history* and can begin to develop a cumulative knowledge of its surroundings.

Once `learn(number)[]` had exited successfully the Nomad was positioned some distance away from home and could invoke `dead-rec(x,y,range)[]` in order to reach the vicinity of home, and then `align` to attain the exact origin position.

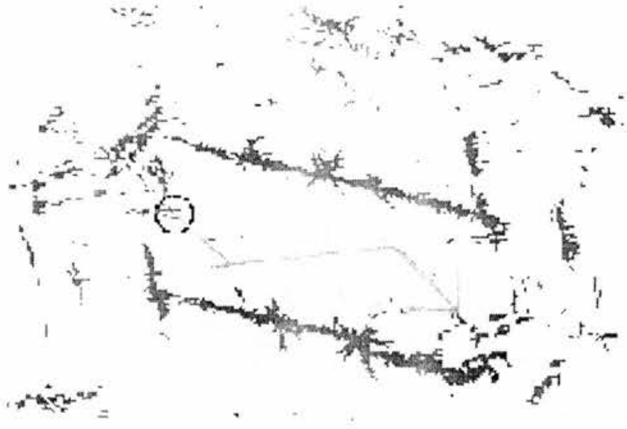

Figure 5.12: The Nomad's view of its noisy world.

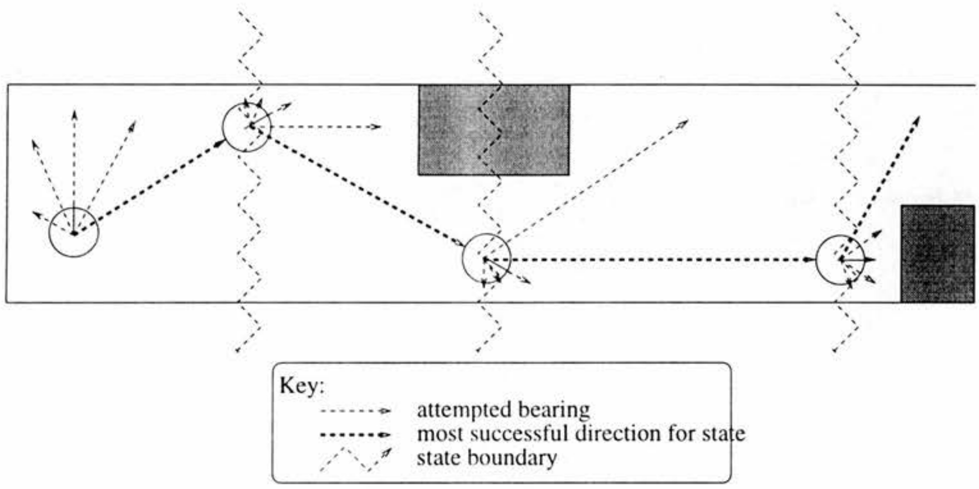

Figure 5.13: Fourth iteration: transitions between sectors.

5.3 Retracing routes

Learning a number of `sectors` and their corresponding `vectors` should ideally result in long-term benefit for the system. The Nomad is a large robot whose batteries run down after a period of 3–4 hours, therefore a long-term memory is essential (see section 4.2).

The routine `mem::save(filename,data-list)[0,1]` is invoked at the end of a run and uses the Perl module 'Storable'[3] to collapse `sector` and `vector` information, and write it to `filename`. At the beginning of a new run invocation of the command `mem::retrieve(filename)[0,reference]` loads this data into memory. The Nomad

[3] Freely available from the comprehensive Perl archive network CPAN: `ftp.funet.fi`.

can therefore access the data related to previous runs.

Assuming that the Nomad is placed near the home location and `align()[]` is implemented, the odometric and compass information relating to the stored `sectors` and `vectors` will remain valid. The task now is to use this information to navigate to a required `sector` without having to repeat the trial and error learning procedure.

5.3.1 Moving outward from home

Routines and behaviours recruited

Library functions	Behaviours
`mem::retrieve(filename)[0,reference]` (sec. 4.3.3) `sector` (sec. 4.3.4) `vector` (sec. 4.3.4)	`retrace` (section 4.3.5)

The algorithm

The behaviour `retrace` (section 4.3.5) queries the internal state variable `sector-list` for available `sectors` or, if none are present, attempts to load a previously learned list by invoking `retrieve(filename)[reference]` (section 4.3.5).

If the behaviour is called with the optional argument `sector-number` and there are sufficient `sectors` in the list, then the command `retrace(sector-number)[0,1]` will be invoked, otherwise the user is asked to input `sector-number` *via* the keyboard before the behaviour can commence.

Now the Nomad utilises routines from the `sector` and `vector` libraries to follow the previously learned vector route to the origin of `sector-number` (see Figure 5.14 and, for more detail, Figure 4.20). The Nomad therefore reaches the desired `sector` using the preferred `vector` within each learned `sector` on the way, without having to repeat the trial-and-error process.

5.3.2 Moving toward home

A vital competence for an active agent is the ability to get back home from a given `sector` of the niche. The way that this was finally implemented in the Nomad was

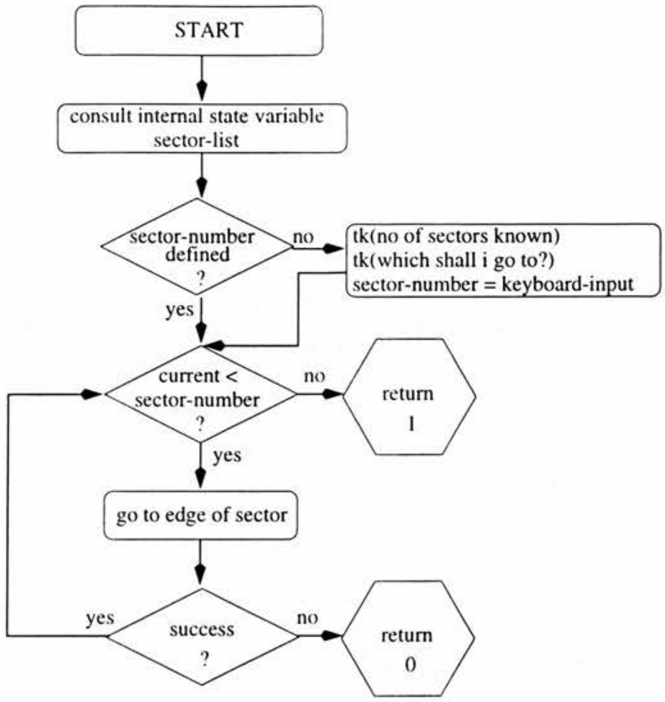

Figure 5.14: Retracing a vector route

simply through a process of dead-reckoning (`dead-rec`) directly home, although dead-reckoning through sector origins, and an alternative, vector-based, method were also examined (see Figure 5.15).

Vector-based

Routines and behaviours recruited

Library functions	Behaviours
`sector` (sec. 4.3.4)	
`vector` (sec. 4.3.4)	

Implementation

Initially, an attempt was made to use vectors to return, through each `sector`, to the home position (see Figure 5.15 B1 for a view of how this ideally appear), unfortunately there were two problems with this approach.

1. Within a given `sector` the most promising `vector` was used to get to the state boundary (`sec::go-to-state-edge()[0,1]`) and, once this boundary had been

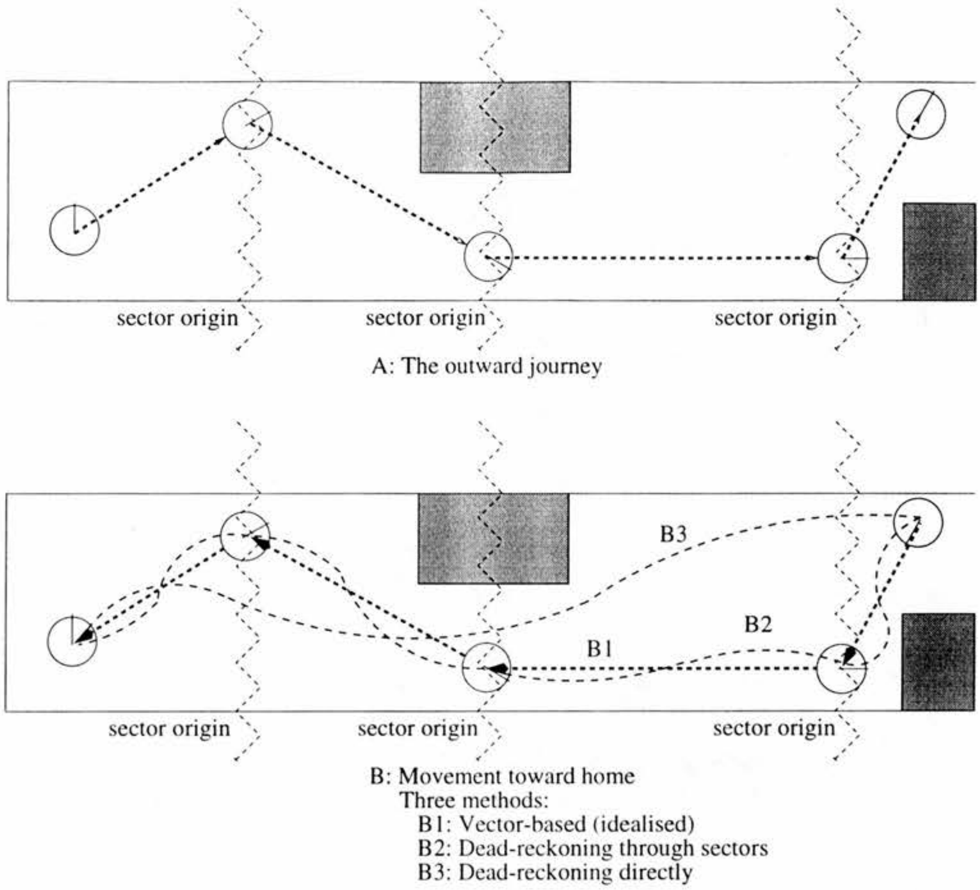

Figure 5.15: Three methods of moving toward home

reached, the value of the compass was read and stored as a private state variable `theta-edge`. This was because of the error build-up on the compass as discussed earlier — at the `sector` edge the value of the `compass` could vary by up to 15% from the initial value of `theta`. It was thought that this modified value might be used to return to the origin of a given `sector` on the return journey. However, movement on the new vector `theta-edge` did not correspond accurately enough to the odometric data gathered by the initial `vector` during the learning stage, with the result that return movement was highly error prone, and not sufficiently accurate to permit robust locomotion to home.

2. Due to this error, and because `sector` boundaries tend to occur relatively proximally to surfaces within the niche (this being one of the criteria, along with success, for termination of movement along a given `vector` — see `vec::run()[time]` section 4.3.4), return locomotion was often interrupted by these surfaces. The

Nomad was often unable to reach a given origin point even with a small angle of deviation from the original `vector`.

Dead-reckoning

Routines and behaviours recruited

Library functions	Behaviours
`sector` (sec. 4.3.4)	`dead-rec` (section 4.3.5)
`vector` (sec. 4.3.4)	

Implementation

Two methods of dead-reckoning were examined.

1. Dead-reckoning from `sector` origin to `sector` origin (see Figure 5.15 B2). This method moved the Nomad home through each of the sectors traversed on the outward journey using the routine `sec::nav-origin(0)[0,1]` — a dead-reckoning procedure using the library routine `go(x,y,range)[0,1]`. This procedure differs from `navigate` in that no account is taken of obstacles. The Nomad does not *'expect'* locomotion to be interrupted as it is essentially retracing a `vector` along which free movement should be possible. As with the compass-based method described above (section 5.3.2) the problem of origin points occurring near surfaces tended to prevent the routine from reaching the origin point of `sectors` and thus often derailed the entire procedure[4].

2. Dead-reckoning directly home from any given point (see Figure 5.15 B3). The behaviour `dead-rec` is invoked and steers the Nomad directly home *without* necessarily passing through the origin points of previous `sectors`. Again `dead-rec` as currently implemented takes no account of obstacles assuming that the run will be clear toward home. Replacing `dead-rec` with `navigate` would allow the Nomad to steer around obstacles on its homeward journey, although the consequence would be that novel obstacles would not be detected.

[4] Of course, providing error recovery procedures is relatively straightforward and would have allowed resolution of this problem.

5.4 Navigational derivations

The implementation described above forms part of an attempt to discover how much navigational behaviour can be achieved relying only on inbuilt compass information and short-range sensing in the absence of both geometric and topological maps. The implementation is grounded in our considerations of *economy* both in terms of design and as an internal arbitration metric. These considerations have resulted in a design approach which stresses installation of *design primitives* in the system, together with economy-based arbitration metrics. Interaction with the environment should, over the life-history of the system, result in improved performance constrained by economical considerations.

It is clear that successful corridor navigation might be achieved using more simple strategies such as wall following (Nehmzow & McGonigle, 1995, 1993), or by using long range sensing to determine areas of free space. However the advantage of the present strategy lies in its extendibility. The result of the learning phase is a segmentation of the Nomad's physical environment into discrete (and navigable/manageable) physical spaces each of which can be treated independently in terms of preferred `vector`, object density (and hence robot velocity), object identity etc. — `sectors` provide state, an *internal context* within which to interpret both signals and behaviours. State-based behaviour transitions, and the categorisation of physical space into areas corresponding to discrete internal, 'cognitive' states, provides a unique opportunity for shifting interpretative stance according to *both* internal and external context.

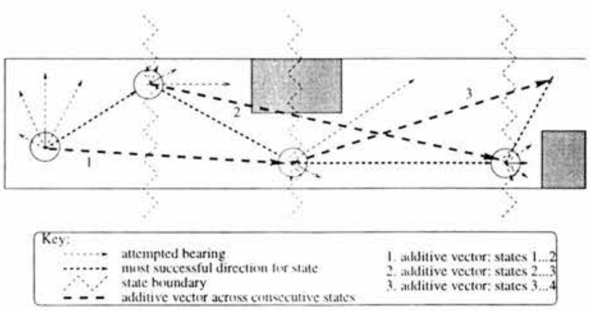

Figure 5.16: Vector addition across consecutive states: 1 & 3 would be successful extrapolations. 2 would fail.

The existential payoff from such a strategy is the automatic detection of environmental variance and invariance. The Nomad has *expectations* about its environment, in terms of the whereabouts of free space. Changing sensory stimulation within the same segment of physical space across runs signals novelty. Furthermore accurate and robust positioning of the Nomad within `sectors` provides a basis for the visual identification of landmarks and other objects from the same orientational perspective on different runs, circumventing the problem of multiple viewing angles. The state-based decomposition of physical space (reflected by discrete `sectors`) also helps to avoid the problems associated with 'perceptual aliasing' — the difficulty associated with distinguishing niche features which are very similar in sensor profile. This is a serious problem for indoor mobile robots (Nehmzow, 1995, for example) where the environment contains many near-identical features such as corners, intersections, doorways *etc.* Adopting a fundamentally state-based approach means that state is *prior* to landmark recognition, rather than subsequent to it as for many navigational implementations (Duckett & Nehmzow, 1999*a*, for example), and thus provides a method of distinguishing similar niche features grounded in *internal* context.

Integration of preferred vector across sectors leads to macro-environmental information becoming available (see Figure 5.16). In this way a progressively more global system of reference can be obtained, which could then serve to constrain future movement, interpretation of default, and of putative objects. With neither long range sensor data, nor human supervision, the Nomad develops a 'cognitive map' of the layout which reflects its history of exploration. This 'map' is analogous to the topological maps of other recent implementations (Duckett & Nehmzow, 1999*b*, for example) yet is achieved at far less computational cost.

One further important derivation of this approach concerns the installation of design primitives and economy metrics whose operation through the system's trial and error experience of its environment can lead to self-organisation of behaviour. Having demonstrated the feasibility of this approach with respect to essential low-level navigational behaviours, the next phase of research should be concerned with extending this process to other areas of system competence such as rational behaviour recombination with respect to goals. It is envisaged that this self-organisational process, constrained by arbitration mechanisms based on economy, will be applied next to self-selection of behaviours for task success — rather than installing a sequence of behaviours to achieve a goal, this sequence should be learnt by the system over time through an analogous trial and error process to the one described here.

Summary

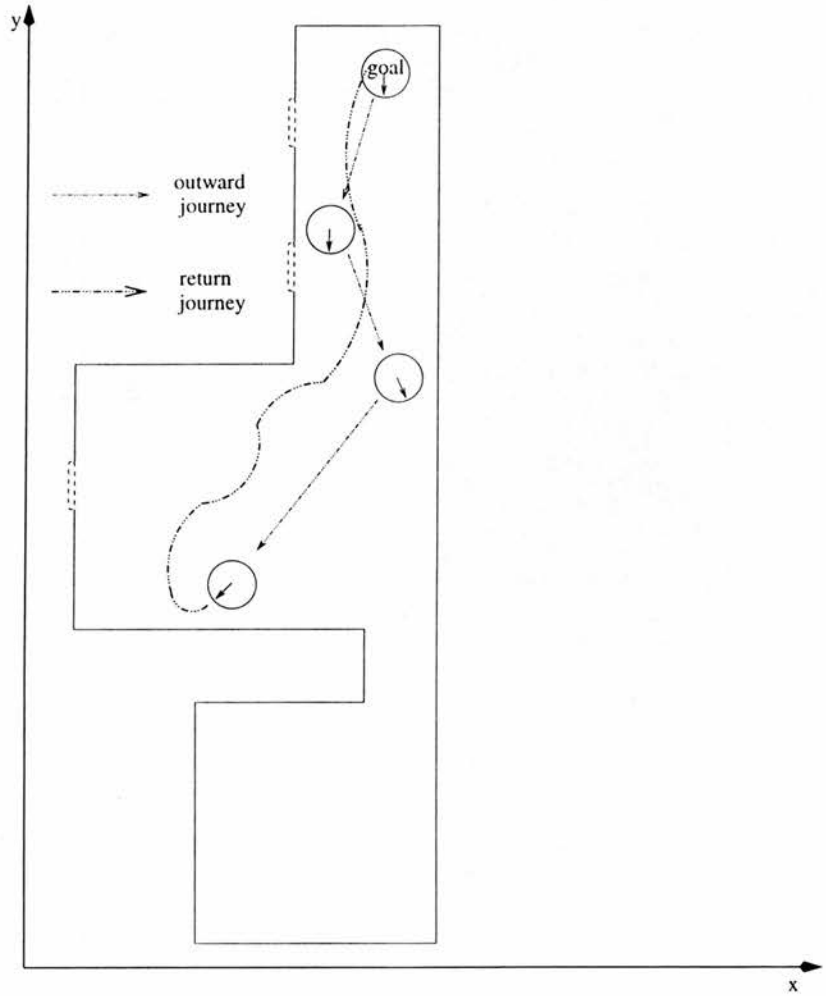

Figure 5.17: Going home

Combining alignment (`align`), learning (`learn`) and retracing (`retrace`) phases with direct dead-reckoning (`dead-rec`) home in a single task results in overall behaviour reminiscent of the navigation of the desert ant *Cataglyphis* (see Figure 5.17). Starting from a fixed position in space on each run, the Nomad first learns about distinct physical regions of space (`sectors`) and the privileged `vector` within each space. This data is written to memory and can be retrieved at any point. Retracing a learned route through the environment using `vector`-based information results in the Nomad following a canonical outward path from home. Use of canonical paths improves navigational accuracy (Gallistel, 1990) and is a feature of the navigation of many biological organisms[5]. Dead-reckoning toward home from a distant `sector` of the niche follows, the exact learned path is now ignored — instead a more direct route is adopted. Once near home an orthokinetic mechanism utilises known features of the niche to achieve accurate localisation.

In the absence of explicit design, the approach results in segmentation of external space reflected by internal `sectors` which provide vital contextual information, and a vector-based topological map of the system's environment achieved at little computational cost. Vector addition across states can lead over time to both increased economy of movement and recognition of macro-environmental features.

5.5 Related work

The problem of self-localisation of robot systems is often (Nehmzow, 1995, for example) divided between global localisation which requires a system to be capable of relocalisation under conditions of locative uncertainty, and position tracking, which assumes that the initial position of the system is known. As many methods of localisation for mobile robot systems have now been proposed (see Borenstien *et al.* 1996 for a recent review), this section will focus on implementations which feature an internal map-based representation of the environment with particular emphasis on those which incorporate autonomous map-construction. A number of exemplars will be described from each of the major approaches: metric and topological schemes, both installed and acquired. It is argued that although these methods have been shown to result in robust navigation, this has been achieved at relatively great expense — both in terms of hand-design of systems and run-time computational demands. Furthermore, although often successful, it will be suggested that these implementations are restricted to the navigational domain, providing little inspiration, or potential, for extension of methods to other domains of competence.

[5] The wood ant (*Formica rafa*) maintains stereotyped paths between food sources and the nest, for example.

5.5.1 Geometrical map-based systems

How accurate is the geometrical characterisation of a navigating system's representational scheme? Although humans clearly use geometrically-based representations their meaning is *acquired* over development. Over the course of ontogenesis an individual system's experience of space, through situated action within some environment, is convolved with *interpretation* of externalised forms of geometric representation (McGonigle, 1999). Such interpretation requires a theory of map understanding, which permits transfer of knowledge according to a conventionalised representational code allowing a system to benefit from the learning experience of others without having to directly learn the information contained within through direct experience. The geometrical maps which are an important feature of human navigational behaviour have evolved over collective human cultural experience — the medieval 'mappa mundi' are vastly inferior to the maps of today, which are themselves now being refined yet further with the aid of global positioning satellite data (Whitfield, 1994). Such maps are conventionalised to support *cognitive economy* on the part of the interpreter. To conflate such external symbolic representations, with the primary mode of spatial representation within systems, therefore, seems premature. Rather, geometrical map understanding is acquired over ontogenesis by a process which requires translation between primary internal, non-symbolic forms of representation and an externalised geometric code. Over the long term artificial systems should be able to acquire such understanding but it should not be *pre-installed* within the system.

Notwithstanding this argument, geometrically-based systems also suffer from a further serious problem — computational expense. Geometrical systems represent Euclidean space using a Cartesian (x,y,θ) coordinate system; all such information is represented whether immediately relevant or not.

Installed

Self-localisation involves map-based information such as landmark position being utilised for odometric recalibration. Typical Cartesian-based systems feature maps where landmark location is *pre-installed*. Lee (1995), for example, developed an implementation which recalibrates a position estimate based on current robot position relative to landmarks whose position was preinstalled within the system.

Burgard *et al.* (1998) have designed an effective museum[6] tour guide robot. A detailed map of the environment was preinstalled within the system based on painstaking hand-measurement of distances between environmental features. The system uses a probabilistic Markov localisation algorithm to update positional estimates over time.

Obviously this approach is not ideal in a number of respects:

- it is highly uneconomical — both in terms of the time spent gathering measurements by hand, and in the computational costs of self-localisation;
- the ontology of the system directly reflects that of the designer: both in terms of the geometric representation, and relevant landmark features — the meaning of the metrical representational schema and landmarks is *preinstalled* and does not reflect the situated action of the system.
- is is inflexible. Navigation will only be successful once the designers have hand-measured the environment. Additionally the system is inoperable in environments which lack the relevant features used for landmark recognition, for example the outdoors.
- it lacks autonomy. The system relies on its human designers both to provide a global frame of reference, and for its ontology of landmark features.

These problems are characteristic of many geometrical systems (Borenstien *et al.*, 1996) and have recently motivated two interlinked research directions. The first concerns autonomous recognition of landmarks based on an acquired feature set; the second requires autonomous map construction.

Acquired

Thrun (1998*a*), for example, describes a Bayesian landmark learning algorithm. The model relies on the assumption that the system operates in a partially observable

[6] Currently operative within the Deutsches Museum Bonn.

Markov environment (Chung, 1960) — one in which the only state is robot location — perceptual and control noise is independent of noise at previous points in time. The algorithm (BaLL) trains neural networks to extract low-dimensional feature representations (occupancy grids) of landmarks based on a Bayesian analysis of probabilistic localisation which provides the criteria for feature extraction — minimisation of localisation error — such that discriminatory features are optimised for locative accuracy. This approach adds autonomy, optimality and environmental flexibility to the system but is unintegrated within a larger navigational system.

Thrun *et al.* (1998) present an algorithm for landmark-based map acquisition with concurrent localisation where the problem of map building is characterised as a maximum likelihood estimation problem where both landmark location and robot position have to be estimated. The system has operated successfully in a cyclical environment of 60 by 25 metres and can manage both position tracking and global localisation. Map acquisition occurs through the process of teleoperation of the robot through the environment. The user selects 'significant places' (such as intersections, dead ends, and corners) and informs the robot that such a place has been reached and the system then learns the sensory signature of the location. Although an improvement over their previous implementation (Burgard *et al.*, 1998) in that a map does not have to be hand-crafted (a process which took several days) and pre-installed in the robot but can rather be acquired through teleoperation (taking instead a few hours), over-reliance on the designer's ontology of landmarks, and the computational costs of geometric representations remain unresolved.

Yamauchi *et al.* (1998) have developed a system which explores and builds maps of novel environments using a global occupancy grid. Again, problems with this approach are the computational costs of constructing a geometrical map of the environment and that this approach requires very accurate laser sensing and precise odometric information in order to construct an accurate map — it is unclear how effective this system might be given less accurate means of sensing.

5.5.2 Topological map-based systems

Resulting mainly from the huge computational costs associated with the use of geometric maps, and inspired by evidence of topological maps in human cortex (Knudsen, 1982; Churchland, 1986; Sparks & Nelson, 1987, for example), an alternative approach to map representation encodes the relative positions of landmarks rather than a global coordinate system. This approach obviously requires that the environment contains features which can serve as detectable landmarks for the system (either installed or acquired) and that the robot follows canonical paths between landmarks in order to disambiguate sensor signals and allow for reliable landmark recognition. Topological maps require much less computational power to construct and maintain than their geometric analogues thus providing the potential for representation of environments of much larger magnitude than those which can by navigated using only geometric representations (Duckett & Nehmzow, 1999*a*).

Installed

As for geometrically-based systems, many topological systems utilise Markov localisation which demands that a map of the environment is preinstalled in the system. Koenig & Simmons (1996), for example, assume that a topologically correct sketch of the environment is available to the system. They use probabilistic Markov localisation, specifically an extension of hidden Markov models to partially-observable decision processes where the robot maintains a probability distribution for a set of discrete locations which is refined over time. The major problem with this approach is again, that potential landmarks (and actions) are defined using a pre-given ontology of doors, intersections *etc.*

Yamauchi & Langley (1997) developed a method where each possible landmark location is represented by a local occupancy grid. Localisation of the system entails construction of a recognition grid from immediate sensor readings and use of a hill climbing procedure to search the space of potential transformations in order to find the best match with a set of stored grids. This system, although effective, demands high computational resources and therefore lacks real-time efficiency. Furthermore, it is ineffective within environments which feature high levels of perceptual aliasing.

Weiss & von Puttkamer (1995) use cross-correlation of laser scans in order to identify robot place and position within that place. In order to reduce computational load laser scans are reduced to histograms prior to matching. Angular histograms are initially convolved to overcome the problem of robot orientation and then x, y histograms are matched with stored representations in order to determine position. Although more rapid and less computationally expensive than the approach of Yamauchi & Langley (1997), it remains unclear whether effectiveness would be retained with the use of less detailed sensors and, again, system efficiency suffers in environments featuring high levels of perceptual aliasing.

Acquired

Shatkay & Kaelbling (1997) have generalised the Markov localisation approach of Koenig & Simmons (1996) to support mapping in the absence of prior information. Instead the system consults local geometric information to discriminate locations and thus localise, but a serious problem with this approach is that it fails to take cumulative rotational odometric error into account (Thrun *et al.*, 1998).

Lu & Milios (1997), together with Guttmann (1996), have proposed a method that matches laser data to partially constructed maps, utilising Kalman filters for positioning. This approach is incapable of representing ambiguities consequent upon perceptual aliasing and can only compensate for a small amount of odometric error. Furthermore, once again it remains unclear whether the approach would generalise to less data-rich sensors.

Duckett & Nehmzow (1999*a*,*b*) have developed an implementation for a Nomad 200 which features construction of a topological map augmented with metric information. This approach combines cross-correlation techniques for matching sonar data with learned sensory signatures, with a probabilistic algorithm for refining a position estimate over time using multiple Kalman filters. This algorithm delivers an updated estimate of both likely topographic location together with most likely Cartesian position. Although the system delivers robust navigation, problems associated with perceptual aliasing remains incompletely resolved (Duckett & Nehmzow, 1999*b*). Furthermore, exploration is unprincipled — the system moves randomly into areas of free space, unconstrained by economical principles. An additional problem is that, as the implementation lacks internal state, there is no way to reliably determine that multiple nodes of the representation reflect identical locations.

5.5.3 Comparison with the current implementation

The types of implementation described above often support robust navigation through the environment. The implementation described in this chapter, although not currently as advanced as some of the implementations described above, does, it is believed, incorporate a number of important principles absent from alternative approaches.

Ontology Many effective robotic navigational systems rely for their operation on a pre-given ontology of both map type and landmark features. Such systems rely on recalibration of odometry through recognition of landmarks which are either niche engineered such as bar-code reflectors (Everett *et al.*, 1994), or easily recognisable visual patterns (Borenstein, 1987), or niche features determined by the designer such as doors (Koenig & Simmons, 1996) or ceiling lights (King & Weiman, 1990) together with hand-crafted filtering algorithms which sift sensory data for the presence of landmarks.

In contrast the implementation described herein is provided with a design primitive whose iteration results in segmentation of physical space on the basis of system capacity (compass error) and obstacle distribution. The 'meaning' of the internal representation of a region of external space (a `sector`) directly reflects the system's situated activity within its environment. For the current area of operation of the system recalibration of odometry has not been found to be strictly necessary. When the size of the area to be explored is increased it is envisaged that odometric recalibration will occur first through orientation to a stored turret θ, followed by execution of an orthotaxic mechanism similar to that of `align` — no reliable detection nor disambiguation of landmarks will be required.

Perceptual aliasing Reliable determination of landmark identity in environments which feature a high degree of perceptual aliasing (the presence of hard to disambiguate landmarks) remains an unresolved problem for the majority of navigational implementations — especially those featuring topological maps (Weiss & von Puttkamer, 1995; Lu & Milios, 1997; Yamauchi & Langley, 1997, for example).

This problem arises because, for these systems, context is *subsequent* to landmark recognition — accurate determination of position is possible only once a niche feature has been reliably identified. The difficulty of effectively disambiguating near-identical landmarks has led to the development of probabilistic Markov algorithms which refine position estimates after the event. The implementation described herein currently possesses no methods for landmark recognition but addition of such procedures, such as occupancy grid methods (Moravec, 1988), is straightforward and is an obvious immediate extension to the system. The advantage of this implementation is that context is *prior* to landmark recognition — reflecting *internal* state. The important consequence of this design feature is that the problem of perceptual aliasing does not arise — within a given `sector` there will be *only one* possible landmark to be disambiguated. Odometric error can then be overcome either through recalibration or again through orthokinetic mechanisms.

Economy The implementation described herein stems from our explicit quest for economy of design, and for economical behaviours. It is now well accepted that geometric maps come at greater computational expense than their topological equivalents (Nehmzow, 1995, for example). Our implementation brings greater economy than characteristic of topological, or hybrid topological/geometrical systems. No complex algorithms are incorporated for filtering feature information from sensor data, no detailed or long range sensor information is required. Rather, through situated activity, the system develops a 'map' of its surrounding environment based purely on obstacle distribution through iteration of a simple algorithm.

Furthermore incorporating economy as a utility metric for arbitrating between behaviours means that the minimum number of `sectors` are generated consistent with relatively error-free navigation through any given environment — the result is minimisation of the amount of representation required.

Extendibility In terms of navigation, our implementation is easily extended to incorporate landmark recognition and object identification. It is also inherently flexible — successful operation does not depend on the presence of detectable niche features, nor indeed any surrounding objects. If operated in an environment devoid of objects, the system would segment physical space into regions whose size reflects only the capacity of the system to locomote within its compass error tolerance, resulting in exploration of the greatest amount of free space for the least computational cost.

The majority of navigational implementations are highly specialised. Their designers do not explicitly discuss the role of navigation within a broader range of competences possessed by the system as a whole. Our implementation is designed to support navigation as *one element* of a multitasking, multiply competent artificial system which has the potential to become progressively more competent over its life-span thanks to installed memory functions. As such its navigational behaviour admittedly cannot currently compete with the most advanced purely navigational systems now developed (Duckett & Nehmzow, 1999*b*, for example) — our *criterion of adequacy* demands only that it can support further competences in the logical hierarchy. Navigation forms only one element of the design, and is currently being used as a test bed for our range of design principles — faultless navigational behaviour is not currently our experimental target.

The most important area of extendibility of our implementation, absent from other navigational systems, is the application of our self-selection mechanism to other areas of competence. The next stage of development will involve applying the logical form of our self-selection algorithm (trial and error experimentation followed by arbitration and behaviour selection on the basis of economy) to other areas of system competence.

5.6 Summary: Navigation

An initial navigational competence was engineered for the Nomad which supported robust navigation to home from distant parts of the niche through the joint operation of an odometrically-based simple dead-reckoning algorithm and an orthokinetic alignment mechanism. State-based reinterpretation of signals was vital to this approach — allowing orientation to obstacles to supplant avoidance when near home. This portion of the research also demonstrated the capacity of the overall architecture to support stacked error recovery procedures.

The navigational learning implementation described in this chapter was designed to examine the operation of a number of design features fundamental to our synthetic approach. Critically, would we be able to design a navigational system based upon the operation over time of a small number of design primitives in conjunction with system *self-selection* on the basis of economy?

The final learning implementation described above demonstrates the feasibility of our synthetic approach.

- Repeated iteration of a simple algorithm resulted in segmentation of physical space into a number of discrete, internally represented `sectors` through situated trial and error activity. The meaning of these internal representations directly reflect the experience of the system within its niche; not the designer's preinterpretation of the world.

- Self-selection of privileged `vectors` was achieved through an inbuilt economy metric, guaranteeing that the most behaviour was obtained for the least computational and time costs.

- Finally the system supports robust and reliable navigation between sectors, both within runs and, due to installed memory functions, over time, providing the potential for lifelong learning.

Although clearly not as advanced as some specialised navigational implementations which have been developed recently, the system incorporates a number of important features absent from other implementations. Critically, the system provides the potential for extendibility in a number of important respects.

State-based interpretation of niche features is an important future development. The problem of perceptual aliasing of environmental features is circumvented since state is prior to landmarks rather than a derivation from them. With more detailed sensing algorithms, or addition of a visual layer, objects could be locatively encoded easing the burden of unique identification.

Odometric recalibration is envisaged to occur through orthotaxic mechanisms which do not necessarily rely on landmark identification.

Prospective control and expectancy are immediate derivations of niche segmentation, Each sector is associated with different 'expectations' of navigable distance; environmental variance can be easily identified by the system.

Abstraction of knowledge can be easily achieved by integrating `vectors` across `sectors`. In this way a more advanced vector-based map of the environment can be constructed.

Transfer of the logical structure of the approach — trial and error experimentation followed by self-selection on the basis of economy — to other domains of competence is an important future area of development.

Our design strategy has been applied to navigation using only very primitive sensing and minimal pragmatic representation. Self-selection of behaviours based on inbuilt economy metrics, in combination with a simple iterative algorithm results in segmentation of niche space into discretely represented `sectors` whose meaning derives from the situated locomotor activity of the system, and is not preinstalled by the designer. Having demonstrated the feasibility of our synthetic design stance within a navigational domain, the way lies open for its application to other areas of robotic competence.

Chapter 6

Resume, the future, and final conclusions

The last chapter described some experiments in self-organised navigation supported by the architecture described in chapter four, which extended the robotic implementations and was motivated by the biological characterisation described in chapter three. This final chapter strives to contextualise this research with respect to the alternative perspectives on intelligent systems, restate the synthetic strategy, and outline some promising future research directions.

Initially the argument is restated. Both traditional symbol-based and behaviour-based stances omit from their biological characterisations some important features of intelligent systems. Neither stance, nor their conjunction, seem capable of supporting the development of artificial intelligent systems. Dynamic and situated perspectives emphasise the transactional nature of intelligence and its development. In terms of design these stances stress the contextuality of adaptive behaviour and the importance of interactions both internal to a system, and between system and environment, in scaffolding intelligence. Unfortunately these stances under-specify artificial designs and the implementations inspired tend to be simple rather than complex.

Next the synthetic strategy adopted herein is restated — our understanding of intelligent systems must be based on both reverse engineering of biological systems, and attempts to construct artificial systems. In contrast with the majority of current reverse engineering approaches which tend to focus on the locomotor and perceptual behaviours of simple systems, McGonigle and Chalmers provide a detailed characterisation of *complex* biological systems. This characterisation of systems self-regulating over ontogenesis towards ever more economical information handling strategies suggests a number of design principles for artificial intelligent systems. Critically, such systems should be designed with sufficiently rich primitives, including both reactive behaviours, specialised inductive mechanisms, and arbitration criteria which through inter-system and system-environment interactions over ontogenesis can support the development of increasingly complex adaptive competences. This synthetic, *complex systems* stance is suggested to overcome the impasse between the static symbol-based characterisations of complex systems, and the unscalable traditional behaviour-based characterisations of simple systems.

Subsequently some promising future areas of research within robotics, suggested by our synthetic stance, are identified. The unifying theme of these areas is progressive adaptation over the life span. Learning to learn, error recovery and diagnosis, and action selection are all fundamental areas within which development is required. Finally, conclusions are drawn.

6.1 The argument

Both the cognitive sciences and artificial intelligence feature a tension between two stances: symbol-based and behaviour-based. Within the cognitive sciences the behaviour-based conceptualisation of systems gave way during the middle of the twentieth century to a symbol-based characterisation. Artificial intelligence has recently witnessed the reverse transition. Although often construed as mutually exclusive within psychology, when applied to the construction of artificial systems a number of limiting isomorphisms emerge.

Over the past two decades the symbol-based characterisation of complex biological systems has come under increasing attack (section 1.1.1). Its experimental emphasis on linguistic tasks serves to reinforce the underlying assumption that the core competence of complex biological systems is the logical manipulation of symbolic internal representations. The assumption that cognition can be abstracted away from its biological and environmental roots has led to the problem of symbol grounding: how to reconnect body and mind. Furthermore, cognitivist accounts are unable to explain the phylogenetic and ontogenetic origins of intelligence.

When applied to the development of artificial systems, in the form of classical robotics (section 1.1.2), the symbolic stance is derailed by two category errors: that the mind can be abstracted away from its biological instantiation; and that the complexity of behaviour is reflected in analogous internal control structures. The artificial systems which result are brittle and unreactive. The implicit assumption that an exhaustive symbolic model of the world is necessary leads to the adoption of toy-worlds and pre-interpreted domains of competence within which system behaviour is hand-crafted. The ontology of such systems is that of the designer.

Behaviourism, in contrast, eschews the explanatory role of internal control and representational structures and instead focuses on the operation of an optimised associative inductive mechanism in biological systems (section 1.3.1). As a result the experimental methodology of the paradigm centres on the relationship between arbitrary events within a small spatio-temporal space in simple systems, or reflexive subsystems of complex systems. Notwithstanding the accuracy of the account for the characterisation of simple systems it remains unclear how the progressive phylogenetic and ontogenetic adaptation of systems can be explained through the operation of, and modifications to, this optimised associative mechanism alone.

The analogous stance within robotics, reactive BBAI (section 1.3.2), attempted to move away from the pre-interpreted symbols of classical robotics by grounding behaviour directly in the world. However, the exclusive focus on behaviour, in the absence of internal control mechanisms and representation, meant that little remains to be grounded: such systems are semantically blind. The behaviour of such systems does not emerge through interactions as hoped, but rather is carefully engineered by the designer. Scaling of purely reactive systems by design has not been successful as a result of the difficulties inherent in hand-crafting 'emergence'. Learning in such systems is restricted to refinement of sensor-actuator links.

Hybridisation (section 1.5) initially appears to be a possible way to circumvent this impasse: traditional reflective components might effectively generate behaviours, with atomic components providing reactivity and robustness. Logically, however, hybridisation can only result in systems which remain hand-crafted and preinterpreted — now at both reflective and reactive levels.

Over the last two decades, dynamic systems and situated perspectives on intelligent systems have been proposed as potential means of resolving this impasse. The dynamic perspective (section 2.1) provides a rich metaphorical language for describing complex systems: ontological lower bounds, state spaces, trajectories, self-organisation *etc.* The computational approaches motivated by dynamicism — connectionism, ALIFE, and evolutionary robotics (section 2.3.2) — share a laudable emphasis on open-ended change in system organisation over time based on inter-system and system-environment interaction but lack explicit, contextualised knowledge and therefore the basis for adaptive internal control. Such systems remain essentially reactive.

The situated perspective on cognition (section 2.2) emphasises the coordination of perception and action dually constrained by physical embodiment and environment. For some theorists adaptive behaviour arises through direct coordination in the absence of internal representation; for others adaptive behaviour is mediated by contextual representations, grounded in direct experience of the world. Such considerations suggest that structural and niche constraints should be utilised to coordinate perception and action in artificial systems. Representation, if present at all, should be limited to non-manipulable, agent-centered information. Unfortunately, these considerations have not been easily translated into clear design prescriptions. Situated robotic systems (section 2.3.1) have demonstrated that augmentation of reactive behaviour-based systems with agent-centered representation can support contextually-appropriate behaviour yet these systems also remain essentially reactive.

Over the past decade researchers within robotics have coalesced into two main groups. One is essentially cognitivist in outlook, and includes classical, and more recently cognitive, robotics. Here researchers (Lespérance *et al.*, 1998; Thrun & O'Sullivan, 1998; Levesque & Reiter, 1998; Shanahan, 1998, for example) are attempting to augment the acontextual logical approach of classical robotics with special purpose algorithms and contextually-informed reasoning. Shanahan's (1998) 'reinvention' of Shakey, for example, uses circumscriptive event calculus (Shanahan, 1997) which views planning as an abductive task. Sensing, planning, and acting are interleaved; often temporal constraints are imposed on the time spent within each activity. These architectures remain essentially sense-plan-act in nature, but behaviour is generated from the partial results of the planner in order to provide reactivity. These systems, like those of classical robotics, are completely preinterpreted: symbols (doors, holds, inroom, for example) reflect the designer's ontological stance. Such cognitive robotic implementations are essentially traditional classical/behaviour-based hybrids with emphasis on the classical component. Researchers are interested in complex competences and thus find the classical cognitivist stance attractive. They hope that the addition of reactivity at the atomic level together with some element of contextual reasoning will allow them to avoid the problems of classical robotics.

Within the alternative, 'animat' approach (see Meyer & Guillot 1991, 1994 for reviews), researchers are designing traditional behaviour-based architectures augmented with deictic representations (Matarić, 1990, for example) or biologically-inspired 'motivational systems' (Halperin, 1991, for example), or are adopting an emergent stance which emphasises connectionist and evolutionary methods (Cliff *et al.*, 1993*b*, for example). These approaches remain at the reactive behavioural level and show little sign of making the transition from simple to complex artificial systems (see Nolfi 1998, for a review).

The central argument of this thesis is that this impasse should be resolved by a new *synthetic* design stance informed by the characterisation of *complex* biological systems, from within a dynamic perspective.

6.2 A new synthetic design stance

A number of researchers within robotics have advocated adoption of a synthetic strategy recently (Pfeifer, 1997, for example). Indeed the goals of artificial intelligence are often described as the construction of intelligent systems as a means of developing a better understanding of biological intelligence (Winston, 1984, for example). Many researchers within robotics are guided by Braitenburg's (1984) law of 'uphill analysis and downhill synthesis' which suggests adoption of a dialectic process of observation of biological systems, together with construction of their artificial counterparts. Currently the impasse within robotics stems from a tension between the construction of artificial systems whose competences are designed to reflect those of humans (cognitive robotics), and those whose competences reflect those of simple biological systems such as insects (the animat approach). Between these two extremes, as within human developmental and comparative psychology, we have an 'unbridgeable gap' (McGonigle & Chalmers, 1996).

The position advocated here is that a synthetic strategy will be effective only if conjoined with a *complex systems stance.* Although simple and complex systems share many design features there remain significant differences between the two which can only be understood by studying complex systems directly. There are qualitative, as well as quantitative, differences in the endowment of neurological machinery between systems of differing levels of complexity (Williamson *et al.*, 1993; Gazzaniga *et al.*, 1998, for example). These differences in lower bounds constrain the trajectories available to different systems, explaining why some remain fixed at purely behavioural levels of adaptation whilst others become progressively more *epistemically* adaptive over the course of ontogenesis. The most cognitively complex systems we see, humans, do not enter the world at their complex end-state, but rather *become complex* over development. This fact suggests that a *dynamic*, developmental perspective must form part of our synthetic position: both with respect to characterising biological systems and developing intelligent artificial systems. The dynamic perspective suggests that we should strive to characterise, and ultimately replicate, the ontogenetic process(es) which lead, from a core collection of design primitives, through inter-system, and system-environment transactions to increasingly adaptive systems, at both behavioural and epistemic levels, in response to environmental challenge.

A core feature of our synthetic stance is that the inspiration for the design of artificial systems comes from neither engineering nor computational perspectives but rather from biology: from the reverse engineering of complex biological systems. Our stance, therefore, falls midway between the 'insect' robotics of animat, evolutionary, and behaviour-based approaches, and classical or, more recently, cognitive robotics which shares a mischaracterised developmental end-state with cognitivism in general.

6.2.1 Traditional reverse engineering

Reverse engineering of biological systems is often claimed to be an important feature of behaviour-based and animat robotic approaches. These traditional behaviour-based approaches, strive to discover the principles underlying adaptive behaviours in their natural setting. Recent exemplary examples of such approaches follow:

Franceschini and colleagues (1991, 1993, 1997) have developed a robotic implementation capable of navigation through visuo-motor coordination mechanisms based on detailed study of the navigational ability of the fly (Franceschini *et al.*, 1996).

Ayers and colleagues are studying the neuroethology of invertebrate and lower vertebrate motor systems in an attempt to establish the adaptive mechanisms which underlie simple locomotor and action patterns in addition to more complex goal oriented behaviour (Ayers *et al.*, 1998). A major research focus of this group is the neuroethology of locomotory behaviour in lobster (*Homarus americanus*). Based on detailed study of the kinematics of walking and navigation behaviours, and the adaptation of lobsters to current and surge (Ayers & Davis, 1977; Swain *et al.*, 1995; Breithaupt & Ayers, 1996, for example), including recording from the motor system of freely behaving lobsters, the research group is developing biologically-based controllers for ambulatory lobster-based (Ayers, 1995, for example), and undulatory lamprey-based (Jalbert *et al.*, 1995, for example) robots.

Kirchner and colleagues are working in collaboration with Ayers *et. al.* in their attempt to reverse engineer the behaviour of the scorpion (SCORPIONIDAE) with particular emphasis on navigation and its capability for omnidirectional locomotion. Their observations of the adaptive behaviour of scorpion have led them to develop an autonomous system whose behaviour is directed by a sequencing controller which releases exteroceptive reflexes along with sequences of behaviour (Kirchner & Hertzberg, 1997).

Although these approaches are admirably grounded in biological observation most reverse engineering remains focussed on locomotor behaviours in simple systems. However, as previously argued, no convincing evidence that incrementation of such systems can lead to complexity has yet been provided in either psychology or robotics.

6.2.2 Reverse engineering complex systems

Evidence from the neurosciences (Gazzaniga *et al.*, 1998, for a review) strongly suggests that qualitative differences pertain between species in their complement of adaptive specialisations (Williamson *et al.*, 1993). Different species possess different inductive, and control, machinery optimised for use in differing internal and external contexts (Hinde & Stevenson-Hinde, 1973; Gallistel, 1990, 1995; Marler, 1991). Such considerations suggest that reverse engineering of simple biological systems, or reflexive subsystems thereof, will be insufficient in isolation to support the construction of complex artificial systems.

Until recently complex systems have been relatively neglected in biological analysis — due largely to a mischaracterised logico-deductive end-state which made elucidation of the causal determinants of behaviour impossible. We were left with the 'unbridgeable gap' (McGonigle & Chalmers, 1996) between traditional symbol-based and traditional behaviour-based stances. With the development of the new developmental and comparative paradigms of McGonigle and Chalmers, reverse engineering of complex biological systems becomes possible as never before. Accurate characterisations of complex biological systems, within dynamic and embodied perspectives, suggests a range of design principles for the construction of complex *artificial* systems.

McGonigle & Chalmers' (1977, 1984, 1998 for reviews) studies of cognitive growth in human and non-human primate suggest a radically different characterisation of intelligent biological systems than that provided by traditional cognitivism. Conjoining comparative and developmental perspectives has allowed them to circumvent the problem of the interdependence of language and thought, and of the origins of meaning for linguistic agents, which beset traditional cognitivist characterisations.

These experimental paradigms portray complex biological systems self-regulating towards maximally economic information-handling strategies in the face of incrementation of task difficulty. The "epistemic" agent is revealed to be on an open-ended growth trajectory, from a number of preinstalled 'design' primitives which form the ontological lower bound, towards progressively more powerful cognitive strategies through system-environment interaction and system self-organisation arbitrated in accordance with economy metrics. The serial control of behaviour does not depend on the possession of a linguistic competence nor on appropriately serially structured features of the environment, as implicitly assumed by traditional cognitive and behaviourist stances respectively, but rather reflects an hierarchical cognitive organisation at multiple levels of abstraction.

This novel characterisation is inherently embodied, situated, and dynamic but diverges from these perspectives, and from traditional reverse engineering approaches, in its focus on progressive *epistemic* adaptation, based on internal arbitrational mechanisms, across the lifespan. The resulting cognitive organisation however apparently "symbolic" derives from object-oriented transactions over ontogenesis. By grounding their characterisation in behavioural assays of system competence, McGonigle and Chalmers can characterise progressive adaptation of systems at *both* behavioural and epistemic levels. This new perspective provides a much-needed bridge between traditional cognitive accounts which focus on logico-deductive competences and traditional behaviour-based accounts which emphasise associationistic modification of tightly-coupled behaviours. The origin of complex, 'cognitive' competences in intelligent biological systems is no longer, therefore, the miracle it sometimes appears to be — we can observe the precursors of cognition in the variety of specialised inductive mechanisms which form the ontological lower bounds of systems, and witness the transformation from simple to complex over ontogenesis.

The central argument of this thesis is that adoption of this novel characterisation of intelligent systems within robotics will perform a similar function — providing a bridge between the ungrounded systems of classical, and cognitive, robotics and the unscalable systems of behaviour-based, and animat approaches.

6.2.3 Design prescriptions

McGonigle and Chalmers' biological characterisation suggests a number of design principles which must be incorporated within any architecture for artificial intelligence which aspires to *epistemic* ontogenetic extendibility — one of the central problems facing contemporary robotics (Kirsh, 1991, for example).

Such architectures should be *embodied*, with relevant structural features utilised to constrain induction whenever possible. Systems should be endowed with sensors which are sufficiently sensitive to deliver rich perceptual data; and actuators which enable the system to act on the world — action scaffolds cognitive development[1]. Systems should be *situated* in real, rather than toy, worlds thereby providing a rich, dynamic environment which challenges the system to become more complex.

The ontological lower bound of biological systems ultimately determines their adaptive potential over ontogenesis: the richness of both 'hardware' and 'software' differentiates systems which remain at primitive levels of adaptation from those capable of progressive epistemic adaptation. Biological systems possess a range of reflexes and instincts. Often these are system-preserving and are, therefore, especially critical at early developmental stages. There should be no objection to installing reactive system-preserving behaviours in our artificial systems. Neuroscience, and neuroethology, reveal that biological systems are endowed with *adaptive specialisations*, from associationistic mechanisms to more abstracted inductive devices, optimised for use in different contexts (Marler, 1991; Gallistel, 1995; Gazzaniga *et al.*, 1998). We should analogously instill a range of reflexive and inductive mechanisms, along with *internal* context, through state, within our artificial systems. McGonigle and Chalmers' characterisation reveals that self-organisation relies on internal arbitration mechanisms — especially those based on economy. Such mechanisms are especially vital for resource limited (in terms of both power and computation) artificial systems.

As with all complex biological systems, artificial systems should be endowed with *memory* in order to support progressive adaptation so that later behaviours rest on earlier achievements (McGonigle, 1991, 1995), and so that epistemic derivations of behaviour can emerge. Furthermore development should be staged (McGonigle & St Johnston, 1995) — it might be that staged development is a logical necessity for

[1] The current endowment of our robotic platform, the Nomad 200, omits any form of actuator which can change world state. Unlike the R2 which was capable of collecting and ordering stimuli (McGonigle & St Johnston, 1995), the Nomad can only move within an environment and was not, therefore, capable of equally striking behaviour. In order to demonstrate the complexity of internal processes in the current architecture the Nomad maintains a vocal commentary on its activities. Obviously this is not ideal. In the long term, the ability to act on the world must be provided.

intelligent systems. Developmental stages might serve to progressively limit induction, by decomposing problems into inductively tractable sub-problems, ensuring that the system advances through stages of 'proximal development' (Vygotsky, 1978). Our artificial systems might also benefit from staged growth through the provision of a series of inductive restrictions which constrain, and thereby enable, epistemic development (Luger & Stubblefield, 1998*b*, pp. 760–761).

Artificial systems should be *modular, hierarchical,* and *serial.* Modularity provides robustness, and allows a system to possess multiple encapsulated, and sometimes logically incompatible competences (Fodor, 1983; Sherry & Schacter, 1987; McGonigle, 1991). *Hierarchical* organisation underlies the ability of complex systems to recombine behaviours at multiple levels to achieve their adaptive goals and, furthermore, supports multiple control, induction, and monitoring mechanisms operating simultaneously at multiple time scales and levels of abstraction. The problem of serial control inheres to all actions of systems (Lashley, 1951). A *serial* architecture capable of accessing modular primitives provides the potential for the development of artificial systems whose behaviour follows a rational syntax and immediately confers contextuality upon a system. Action and perception can then be instigated and interpreted within specific contexts indexed by both internal and external state.

Finally our artificial systems should be truly *autonomous*; not merely flexless but capable of self-regulation, initially at behavioural, but later at epistemic levels. Such autonomy demands recombination of behaviours to achieve goals, in conjunction with contextual interpretation of behavioural success. Systems must, therefore, be endowed with, or acquire, arbitration criteria and be cognisant of behavioural success and default. *Self-organisation* through inter-system and epigenetic interactions over ontogenesis based on installed primitives and experience, should lead to progressive adaptation.

Adoption of these principles constitutes a third position, situated midway between the approaches of cognitive robotics, and animat researchers, yet providing a method of moving from simple to complex artificial systems both through design, and through system self-organisation.

6.3 The future of robotics

The research reported herein is based on a design stance which has motivated a formal functional architecture together with a number of robotic implementations within the Laboratory for Cognitive Neuroscience and Intelligent Systems at the University of Edinburgh — a stance which strives to steer a middle path between both traditional cognitivist and traditional behaviour-based approaches to robotics. Based on characterisations of the epistemic growth of complex biological systems (McGonigle & Chalmers, 1998*a*, for a review) the long term focus of research is the construction of complex artificial systems. McGonigle & Chalmers' biological characterisation concentrates on learning over ontogenesis motivated by incrementation of environmental complexity, which allows a system to develop progressively more efficient information-handling strategies and to become increasingly inductively powerful over ontogenesis.

This section focuses on future areas of development within robotics congruent with both our biological characterisation and the history of robotic implementations from our research group. The areas addressed focus on the core abilities of systems to 'learn to learn' and to self-organise in the light of experience over ontogenesis — convergent with one of the most fruitful areas of current machine learning and an area recently predicted to be the future of artificial intelligence (Mitchell, 1996).

6.3.1 A life-historical approach

Learning to learn demands an experiential history. Biological systems become progressively more adaptive by scaffolding competences on those already achieved in the face of continual environmental challenge (McGonigle & Chalmers, 1998*a*). Systems 'learn to learn' by exploiting previous successes in current challenges. A life historical approach is therefore a necessity for any artificial system which we desire to be capable of progressive epistemic adaptation (McGonigle, 1991).

Classical machine learning techniques required the designer to prespecify what is learned by a system and how (Lee & McGonigle, 1996; Luger & Stubblefield, 1998*b*). Recently reinforcement learning techniques have become more popular within robotics, but are slow and scale badly (Wyatt, 1995). Logic-based abduction has also become more popular (Shanahan, 1996, for example) but the restrictions of the logic-based approach result in a failure to capture the pragmatic roots of symbol semantics. Learning from experience, where the context and relevance of learned information is discovered by the system itself, remains the central problem of AI (Samuel, 1983).

More recent machine learning implementations have attempted to address some of these issues — learning to learn is a relatively recent departure. As with most traditional machine learning methods, general functions are induced from experience, although attempts are being made to develop algorithms that can change the way in which they generalise (Thrun & Pratt, 1998, for a collection of recent articles). Part of this endeavour involves a recognition of the importance of learning across the lifespan. Thrun (1998*b*), for example, addresses progressive system exposure to a series of tasks. In these cases the possibility of transfer across multiple tasks arises. Preliminary evidence in the domain of object recognition, indicates that resultant learning is more efficient than possible within traditional 'single-task' machine learning approaches.

Recent research directions focus on cross-task transfer as the basis of incrementation. CHILD (Ring, 1998), for example, is an agent capable of continual, incremental and hierarchical learning. Solution of complicated non-Markovian reinforcement learning tasks can be transferred to similar, but more complex tasks, which the system consequently solves more rapidly. Schmidhuber *et al.* (1998) describe a reinforcement learning system capable of altering the way it self-modifies over time, through utilisation of a "success-story" algorithm (SSA). The SSA uses backtracking to delete previous modifications which have not been found to be associated with accelerated learning. This approach, analogous to our implementation of internal arbitration based on an economy metric, was found to result in significantly accelerated learning in comparison with alternative methods.

The approach adopted herein shares some of the aims of these recent approaches, but relates most closely to contemporary cognitive and expert systems research that concentrate on a framework which assumes progressively more proficient interpretation of input based on a system's existing knowledge base and derived expectations (Jackendoff, 1983; Bruner, 1990; Stern & Luger, 1993, for example). Such approaches eschew the Tarskian (1944, 1956) static mapping of symbol to object instead adopting a more pragmatic, contextually-informed conception of meaning. Indeed, this perspective continues that of Pierce (1958), and later de Saussure (1974) and Grice (1975) where the meaning of signs is derived from their relation both to other signs and, critically, sign interpretation. This perspective suggests that the semantics of signs can be understood only in relation to their interpretative function and in the context of action within the environment. Clearly, if the changing interpretative stance of a system determines the meaning of signals and behaviours, then designing systems capable of contingent contextual interpretation is a necessity — artificial systems should therefore be capable of multiple task-achieving behaviours.

The following sections deal, first, with an abductive approach to error diagnosis and recovery — the primary focus of our long-term investigations of increased interpretative proficiency, and second with self-organisation of behaviours in the context of the problem of action selection.

6.3.2 Error recovery and diagnosis

We have seen that a core feature of the biological analysis discussed in section 3.1.3 is self-organisation through *internal arbitration.* Clearly such arbitration between competing behaviours requires system *cognisance* of behavioural success and default together with arbitration criteria and interpretative mechanisms. The synthetic agenda espoused herein suggests that intelligent artificial systems should be designed which incorporate these features and this must consequently be a key area of system development.

Within AI error is often regarded as the result of an imperfect implementation of an algorithm which is essentially faultless in human. The constrained domain of assembly robotics is the main area within robotics where the implications of error have been considered. Here two kinds of error are generally distinguished: errors avoidable by the system through fault tolerance mechanisms and those arising from inconsistency between a system's world model and current world state. The first type of error is most often handled by preventing the error situation occurring in the first place, for example Trevelyan & Nelson (1987), Längle & Lüth (1995), and Malcolm (1997) all discuss hybrid systems which utilise reactivity at the atomic level to provide fault tolerance and thereby circumvent the occurrence of this type of error. Errors resulting from inconsistency of world model and the current state of the world call for interpretation, diagnosis, and recovery (Srinivas, 1977).

Error cognisance, interpretation and recovery remains relatively unexplored within non-assembly robotics. An exception is Simmon's (1990) task control architecture (TCA) which consists of a set of task-specific computational processes (modules) which communicate with one another by passing messages through a central control module. This central module dynamically routes messages amongst the task modules. Tasks are constructed as hierarchical task trees which encode the parent-child relationships amongst messages. The system allows concurrent execution of steps in the task tree — both computational and physical processes. Monitoring and error handling procedures are then added after the code for handling normal situations is in place, but the architecture does not address low level control.

The most often cited example of error recovery within non-assembly robotics is ATLANTIS (Gat, 1991*a*,*b*) a hybrid, three-layer, navigational architecture based on the earlier reactive action packages (RAPs) of Firby (1987). The lowest level of the system controls the execution of action primitives; the second controls sequences of activities; the third runs asynchronously and performs time-consuming computation (i.e. planning). Action primitives are annotated with a list of resources they use, and a set of semaphores prevents two interfering primitives from being active simultaneously[2]. A task bin consists of the task schemas in which the system is currently engaged, each of which is a collection of methods for initiating, monitoring, and terminating behaviours, as well as for invoking other task schemas.

An explicit design feature of ATLANTIS is error cognisance and recovery. According to Gat (1991*b*):

> "Designing a robot from the ground up to fail cognizantly is absolutely crucial to robust intelligent behaviour." (Gat, 1991*b*, p. 36)

Gat's rationale is threefold: designing systems which fail explicitly simplifies the design

[2] Thereby providing a form of implicit behavioural syntax analogous to that supported by the initialisation and default conditions of the current architecture — see section 4.3.5.

process — it is easier to develop algorithms which fail and report that failure than it is to design those that never fail. Designing a system with cognisant failure and error recovery means that the designer does not have to explicitly anticipate all possible occurrences at the outset; furthermore, through combination of imperfect algorithms and error recovery overall system behaviour is far more reliable than any of its individual component actions (Gat, 1991*b*, p. 37).

Within ATLANTIS failure of an action primitive is reported to the invoking task. The sequencing layer then invokes recovery procedures. Recovery procedures include retrying the action, or choosing an alternative method from the task schema. If all attempted methods fail, this failure is reported once more to the invoking task, it is unclear what happens consequently. As instantiated, ATLANTIS possesses a pre-installed error typology, with hand-crafted error recovery procedures available for each type of failure.

State-based evaluation

The extendibility of Gat's approach, and the later development by (Gat & Dorais, 1994) in the absence of state-based decomposition (an issue he does not explicitly address with respect to error recovery) will be limited — just as classical learning approaches suffered from a lack of state-based decomposition. The storage of successful macro-operators in STRIPS (Fikes *et al.*, 1972; Nilsson, 1980), for example, resulted in a serious combinatorial problem — as the number of stored operators increased so did the time spent pattern matching in order to determine whether or not an operator could be applied (Luger & Stubblefield, 1998*b*).

Biological systems do not seem to suffer such penalties from acquiring new data or skills, and it is clearly important that our artificial systems can also benefit from learning. Decomposing system experience into distinct contexts, based on both internal and external state, is one method of circumventing this combinatorial problem (McGonigle, 1995). Induction is constrained by context — a system is enabled to learn by partitioning its experience, allowing *reinterpretation* of external signals together with the status of behaviours. The semantics of such a system no longer feature static associations between symbols and objects (Tarski, 1944, 1956) but rather permit variable interpretations of identical phenomena dependent on a range of contextual features.

Again the stance expounded herein strives to *re-embody* and *resituate* the behaviour of intelligent systems within a wider pragmatic context.

Furthermore, providing a *task grammar* enables a more principled approach towards error. As argued by McGonigle (1995) the syntax of most tasks results in error distributions skewed towards the consummatory end of a series of behaviours. The architecture described herein is motivated by an approach to error detection and abduction first detailed by McGonigle (1991). Here a task-grammar associated with internal state transitions allows a decomposition of serial behaviour into a number of discrete and critically, self-discriminable, states each of which is associated with its own repair procedure. Based on an earlier model of transitive choice behaviour (Harris & McGonigle, 1994), McGonigle (1995) proposed an abductive reasoning as rule stack model of error recovery (see figure 6.1).

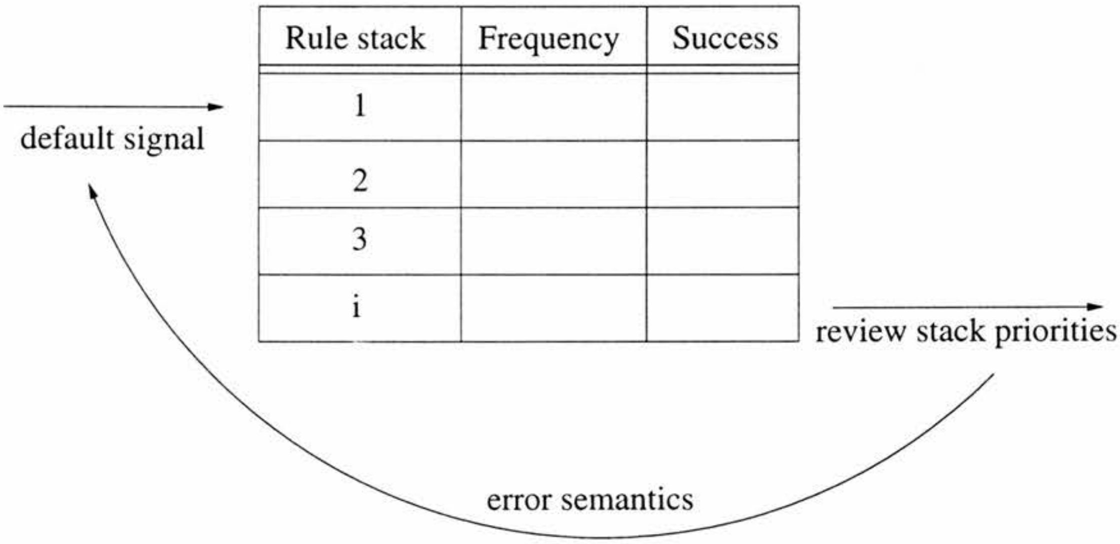

Figure 6.1: The rule-stack model of error recovery (after McGonigle, 1995).

The current architecture is designed to incorporate this rule stack approach whilst providing for the addition of interpretative mechanisms at a later stage of development. The complexity of both the niche (figure 4.3), and the Nomad's behaviours (section 4.3.5) ensures a rich error space. The architecture reported herein (chapter 4) bears some similarity to ATLANTIS. Here behaviours are associated with appropriate recovery procedures by the designer (see section 5.1.4) although an error typology is not provided explicitly. Behavioural default is associated with a set of recovery procedures, each of which maintains a record of the (internal) context in which it has occurred, the number of times it has been recruited, and its success rate.

Now armed with an architecture cognisant of error, the next phase of investigation should be targeted at more detailed error interpretation. Here a *state-based approach* is vital. Within the architecture described herein, each behaviour is provided with associated explanation templates, instantiated and maintained uniquely for different internal contexts (see figure 6.2). Error interpretation and recovery for the same behaviour can differ between different tasks, and even at different stages within the same task thereby providing context specificity. In this example, navigational error, *e.g.* due to timeout, might arise for different reasons dependent on the antecedent behaviour. By allowing contextual interpretation of behavioural default, timeout error due to temporary interruption, for example, could be contrasted with that consequent upon the Nomad's presence within a particular portion of the niche (the `sector` associated with `retrace(3)` for example). Through the provision of internal state both identical signals, and seemingly identical error types can be distinguished.

Currently our approach features hand-crafted priorities for different recovery procedures which can be modified in the light of experience. Based on successful transition to the consequent behaviour or task, the success of different recovery procedures within a given context can be assessed allowing a system to benefit from its experience over time. This demonstration of architectural support for error cognisance together with stacked recovery procedures, was clearly an essential precondition for our next phase of research.

It is suggested that providing *explanation templates* for behavioural default, together with an abductive inference procedure, will enable the system to learn its own error

TASK: ALIGN - AVOID - NAVIGATE - ALIGN

NAVIGATE (align - avoid -navigate - align)

ERROR	PROCEDURE	N	SUCCESS
obstacle	avoid		
timeout	create new goal		
	follow wall		
odometry	call for help		
power	call for help		

TASK: ALIGN - RETRACE(3) - NAVIGATE - ALIGN

NAVIGATE(align - retrace(3) - navigate - align)

ERROR	PROCEDURE	N	SUCCESS
obstacle	avoid		
timeout	create new goal		
	follow wall		
odometry	call for help		
power	call for help		

Figure 6.2: Explanation templates for two different tasks.

typology, and determine effective recovery procedures for each type, over the course of its life-history. Clearly, allowing systems to self-determine error occurrence and recruit appropriate recovery procedures is an important future area of research.

The future: abductive error interpretation

A critical future area of system development is learning of an error typology through abductive mechanisms. The abduction of error types is made tractable by the segmentation of experience into discrete contexts in turn made possible by encapsulated tasks, behaviours, and action primitives. In the long term it is hoped that the use of explanation templates together with state-transition analyses will enable a system to learn the *semantics* of default. By keeping a log of errors by frequency and mapping these according to both niche space and the grammatical characteristics of the task in hand,

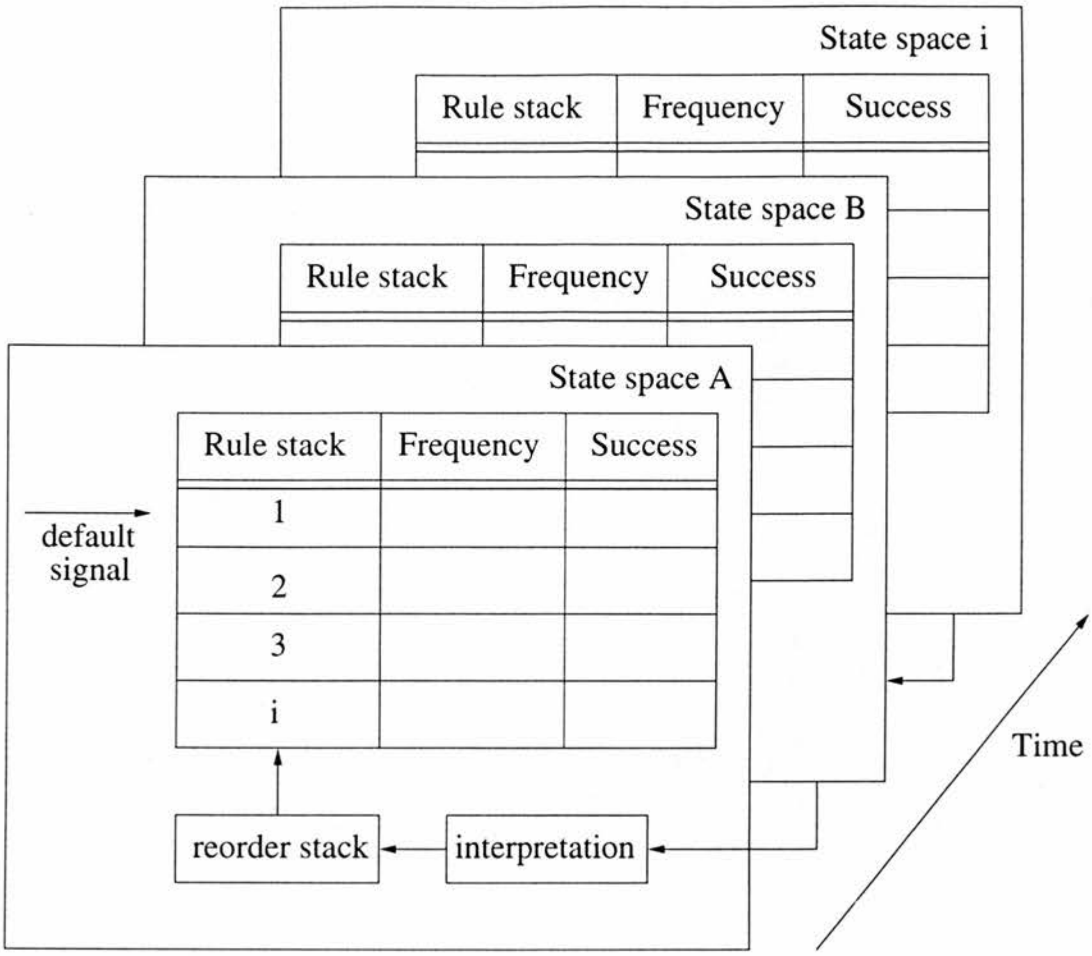

Figure 6.3: State-transition analysis (after McGonigle, 1995).

the original set of recovery procedures can not only be reprioritised but, furthermore, a new set can be established.

Within AI diagnosis has been investigated primarily only within the field of expert systems. Early such systems such as MYCIN, for example, represented domain knowledge and problem solving strategies as a set of relational facts and production rules. Rule chaining was used to search through the space of potential solutions. A serious weakness in the production rule approach, however, is that the intermingling of theoretical domain knowledge with heuristic search strategies led to only weak semantics (Clancey, 1985) and also to brittleness (Luger & Stubblefield, 1998*b*); furthermore, the system had to be provided with explicit representations of search strategies and heuristics for evaluation of evidence by domain experts resulting in high knowledge acquisition requirements (Luger & Stern, 1999). These problems led to the development of systems whose representational scheme more clearly distinguished domain knowledge and heuristics, and which were less reliant on explicit expert knowledge for search heuristics and evidence evaluation.

Later systems, such as CASNET (Weiss & Kulikowski, 1979) utilised an explicit model of hierarchical causal relations. This system featured a 'casual-associational' network, a type of semantic network which represents dynamic processes by the causal relationships amongst nodes. This representational scheme had three interconnected levels — classification, observation, and pathophysiology. A complete casual path through the network from start node to terminal node reflects a complete disease process. The activation of the network spreads in an analogous fashion to weight propagation algorithms (Luger & Stern, 1999), with current network state driving hypothesis construction through test selection. Although an improvement over earlier systems, causality in CASNET remained at a superficial level and the structure of the network was hand-crafted by the designer. ABEL (Patil, 1981) extended hierarchical causal models by providing much richer, and contextually-informed, causal links.

The architecture described herein relates most closely to recent models of schema-based abductive interpretation, based on analysis of the reasoning process of human experts. Whereas logic-based accounts rest on a Tarskian commitment to a fixed relationship between sign and signified, and tend to assume that the explicandum is fixed at the outset of the inductive process, abduction is the inductive process of constructing explanatory hypotheses dynamically (Pierce, 1958) where the current state of knowledge regarding the problem drives knowledge acquisition and evaluative processes.

Stern & Luger (1993) have developed a very fruitful approach to error diagnosis. Their system is designed for semiconductor failure analysis and is informed by detailed observation and interviews with human experts. Human experts seemed to rely only on a limited number of explanation patterns reflecting the causal patterns ('failure mechanisms') they had experienced. The initial hypothesis formation process involved the use of heuristics to select and order possible appropriate failure mechanisms, followed by elaboration of causal hypotheses over the investigative process. The entire process seemed to be motivated by the use of schemas, abstract causal patterns, to interpret evidence and focus abductive reasoning. This approach is much more congruent with the one envisaged here in its pragmatism and dynamism — current explanatory hypotheses influence the explanatory causal patterns chosen. Furthermore, the current architecture allows nesting of error analysis — recovery procedures can, themselves, be treated as behaviours and therefore subjected to similar error analyses — as do explanation schema models. Causal processes can be specified in terms of other schemas (Stern & Luger, 1993). Luger and Stern (1999) are currently investigating the possibility of extending their abductive approach with Bayesian techniques in a robotics domain — this is an extremely promising future research direction.

6.3.3 Action selection and multitasking

A second crucial area of future development concerns multitasking artificial systems. The ability for a single system to perform multiple tasks is not merely important with respect to the interpretative utility of multiple contexts, but also with regard to issues of design, control flow, and scaling.

For resource and time limited biological systems the ability to 'do the right thing at the right time' (Maes, 1989) is critical. The problem of maximising fitness, fundamental to all biological systems (Dawkins, 1989), decomposes into many sub-requirements necessitating the possession by systems of effective mechanisms for choosing which activities to engage in and at what time. As the world is of course only partially knowable, such 'action selection mechanisms' can neither be completely rational nor optimal but must be robust and effective (Simon, 1955). Biological systems must interleave high priority system-preserving behaviours such as eating, drinking and avoiding predation with more slow-burning behaviours such as reproduction. Action selection must therefore be both goal oriented and responsive to the current situation, reflecting both internal and external contextual features. Biological adaptation involves multiple coexisting competences operating on multiple timescales with great potential for both conflict and cooperation.

Humphrys (1997) distinguishes between the 'w-problem' — choice at 'system' level of which goal-directed behaviour to pursue — and the 'q-problem' — which activity, or subsystem, to execute in pursuit of the chosen goal. A large number of mechanisms have now been proposed to account for action selection in both biological and artificial

simple	*vs.*	complex
single task	*vs.*	multiple task
constrained	*vs.*	unconstrained

Table 6.1: A proposed classification scheme for artificial architectures.

systems, some of these address only one of these problems whilst others strive to address both. Generally, traditional hierarchical and serial models of action selection are top-down featuring choice of goal followed by choice of activity. More recently, parallel models tend to intermix the two problems.

We suggest[3] that architectures can be characterised on three dimensions (see table 6.1): *Simple vs. complex* refers to the complexity of competences exhibited by the architecture; *single vs. multiple task* reflects whether the executed action of the system is informed by multiple goals or by only one goal at a time; *constrained vs. unconstrained* reflects whether the designer has specified (usually hierarchical or serial) constraints on the choice of final behaviour or their relationships to one another.

Currently the major debate within action selection is between models which feature some designer-imposed constraints on the ordering of system behaviour (*i.e.* they are hand-crafted), and those which rely on free competition between behavioural modules. A brief description of some of these different architectures follows, before the next section examines recent suggestions that the limits of hand-crafted control have been reached.

Constrained systems feature hand-crafted conditional relationships between successive behaviours or constraints on action selection.

Single, simple systems such as those of cybernetics (Grey Walter, 1950, 1951, for example). Here the system's behaviour is in pursuit of a single goal. System behaviour arises from hand-crafted responses to external stimuli.

Single, complex systems such as those of traditional AI, reactive planning, and hybridised classical/behaviour-based systems. Some systems (Kaelbling, 1993) feature only a single task leaving only the q-problem to be

[3] Tentatively!

addressed whilst others, such as Gat (1991*a*) resolve the w-problem (which task to pursue) before choosing which actions to execute in its pursuit (the q-problem).

Hierarchical models feature behaviours or action modules ranked in some pre-defined order of priority. Control flow is unidirectional — from high-level to low-level modules. A 'behavioural common path' (McFarland, 1974) acts as a bottleneck to ensure that only one activity can be expressed at a given time. Many hierarchical architectures also feature constraints on the serial ordering of behaviours through annotation with pre- and post-conditions. Termination of a behaviour updates state, allowing the activation of further behaviours whose pre-conditions are met thereby providing pre-installed routes through the space of possible module combinations (Gat, 1991*a*, for example).

Within ethology Tinbergen's (1950, 1951), Lorenz' (1973), and Baerends' (1976) models all feature hierarchical decision structures with activation spreading downward through the hierarchy. Both Tinbergen and Baerends' models feature inhibition within layers such that only one activation path can be active at any time thus ensuring that the w-problem and q-problem are separated.

Multiple, simple systems such as subsumption (Brooks, 1986*a*), Maes' (1989) spreading activation networks, Rosenblatt and Payton's free-flow hierarchies (1989), and Tyrrell's (1993) extension of the latter.

These systems feature constantly active behavioural modules where the outcome of competition is pre-determined to some extent by constraints imposed by the designer [4].

Unconstrained systems have become more fashionable recently due partly to the popularity of emergent views of action and cognition and also reflecting dissat-

[4] Although subsumption is often portrayed as a parallel distributed system (Brooks, 1986*a*, 1991*c*) in reality such architectures are best characterised as pre-structured hierarchies (McGonigle & Humphrys, 1998); Maes' spreading activation networks feature hand-crafted serial conditions; Rosenblatt and Payton's DAMN features hand-crafted utility functions; Tyrrell's extension features hand-crafted penalties for temporal delays and reward uncertainty allowing the system to optimise behaviour (Humphrys, 1997, for further details).

isfaction with the degree of hand-crafting necessary for constrained single and multiple-task models. These systems feature multiple parallel flows of control each in competition with one another. The w-problem is either unaddressed (for single task systems) or intermixed with the q-problem.

Single, simple systems such as those of evolutionary robotics (see section 2.3.2). These systems pursue one simple goal. The behaviour of the system emerges from the evolutionary (and sometimes also learning) process(es) without explicit constraint, other than task and fitness specifications, from the designer.

Multiple, simple systems such as that of Humphrys (1997). These systems feature multiple behavioural modules constantly active in parallel. The designer does not impose hierarchical or serial constraints on which behaviours can be executed at given times.

Humphrys (1996*a*, 1997) describes a multiple mind architecture, tested in a simulated environment, which features multiple parallel control flows each competing with one another. The designer does not impose constraints on behavioural choice but rather provides the potential for *self-organised* action selection. Multiple behaviours are instantiated and henceforth compete with one another for control of the system, control is gained once the strength of a behaviour (derived jointly from the expected reward for the behaviour if obeyed *vs.* the penalty of the behaviour *not* gaining control) exceeds that of the behaviour currently in control. Control of the system is thus *dynamically determined* by the strengths of all behaviours at a given time. By logging the reward functions of behaviours the system can self-determine unexpected compromises and conflicts between behaviours. Humphrys' system uses a 'minimise the worst unhappiness' strategy of exploiting reinforcement learning for action selection.

The limits of hand-crafted control?

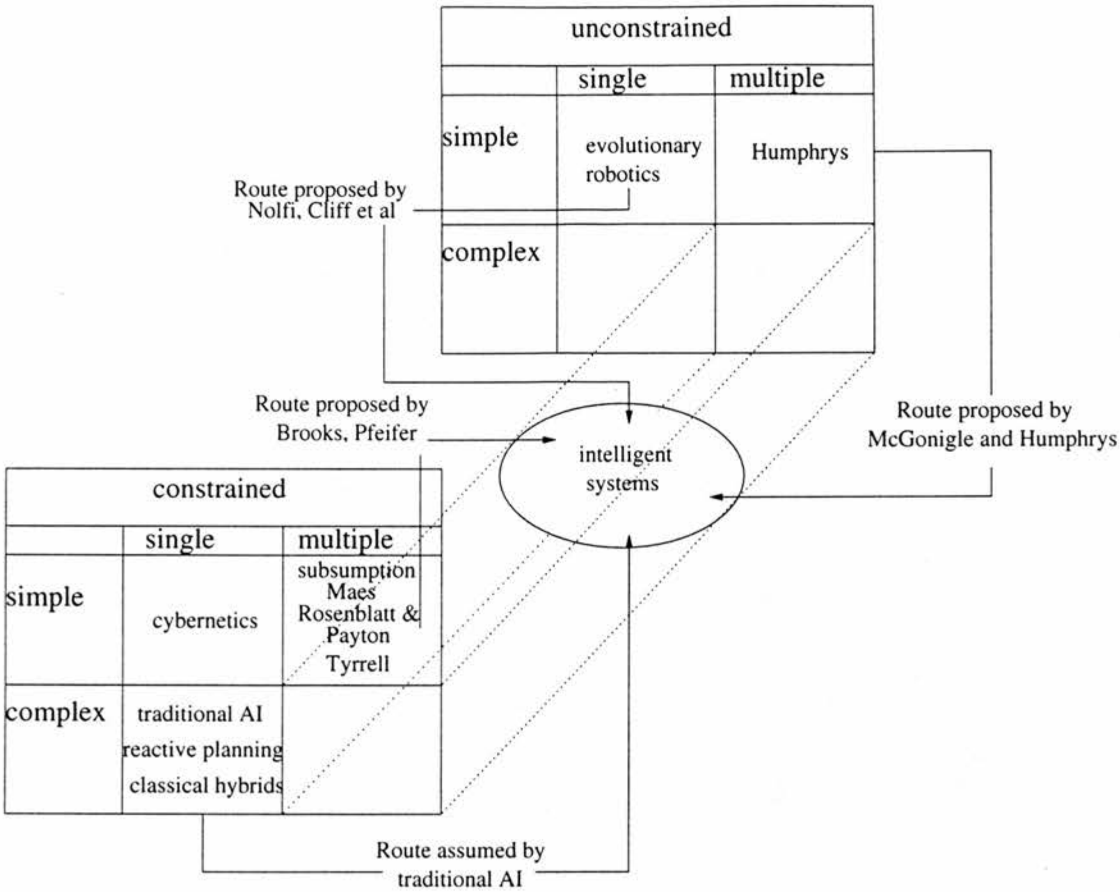

Figure 6.4: Possible routes towards complex artificial systems.

The biological systems which are the targets of artificial intelligence fall into a category unaddressed above — their behaviour is *complex* and informed by *multiple* goals. It is constrained to some degree (by embodiment, the environment, and the ontological lower bounds of the system (see section 3.1.2), yet it is also derived from self-organisational processes over ontogenesis (see section 3.1.3). How are we to move from the artificial systems of today — the majority of which are constrained, multiple, simple, or constrained, single, complex architectures — towards artificial systems which exhibit complex behaviours informed by multiple goals whilst constrained by system primitives, yet which are capable of self-organisation over development? What possible routes (see figure 6.4 and table 6.2) for the scaling of systems can we envisage?

From	History	Prospects
(A) constrained-single-simple (Grey Walter, Weiner)	Abandoned	
(B) constrained-single-complex (Traditional AI)	Unsuccessful (so far)	cognitive robotics?
(C) constrained-multiple-simple (Brooks, Maes, Pfeifer)	Unsuccessful (so far)	distributed BBAI?
(D) unconstrained-single-simple (Cliff *et. al.*, Nolfi)	Unsuccessful (so far)	Evolutionary robotics with learning?
(E) unconstrained-multiple-simple (McGonigle & Humphrys)	Unattempted	self-organised action selection?

Table 6.2: A preliminary assessment of potential routes towards complex artificial systems: their advocates, history, and prospects.

Route A from constrained, single-task, simple systems to complex artificial systems. This route essentially died out, excepting the thought experiments of Braitenburg (1984), with the abandonment of cybernetics. Its prospects are, therefore, not good.

Route B from constrained, single-task, complex systems to complex artificial systems. The preferred route of traditional AI and contemporary cognitive robotics. Although single task, constrained (hierarchical and serial) models like those outlined above succeed in their task of augmenting internal control with reactivity, arbitration between behaviours, and ultimately action selection, remains hand-crafted by the designer. Indeed such hand-crafting is explicitly adopted as a means of reducing the combinatorial problems of control as well as providing context sensitivity and goal-directedness (Bryson, In press). This focus, of the majority of robotics research, on single goal systems has recently been suggested to be a fundamental strategic error (McGonigle & Humphrys, 1998; McGonigle, 1998). Research on single complex tasks seems to be undertaken with an implicit assumption that these competences can later be added together to support complex behaviour informed by multiple goals. Possibly no such path exists — McGonigle (1998) suggests that multi-goal systems raise control issues which are not simple additions of single-task skills, nor simple forms of scheduling or optimisation but rather that, as Dennett (1978) argues, we need to conceptualise the entire adaptive system from the outset, and address a developmental progression from multiple simple tasks to multiple complex tasks. The flow of control for such systems might be so complex that it is impossible to hand-craft — suggesting that self-organised action selection might therefore be the *only* way of developing multi-tasking artificial systems.

Route C from constrained, multiple-task, simple systems to complex artificial systems. The preferred route of the behaviour-based, and more recent animat, camps within contemporary AI. The prospects for this route are not good. These approaches hope that 'emergent functionality' will result in ever-more complex behaviours. As argued in section 1.3.2, such emergence is really hand-crafted — McGonigle & Humphrys' (1998) critique applies equally to this approach — if hand-crafted, yet the interactional complexity of our target systems is intractable, then how can scaling take place? Over the past 15 years of research no indications have been provided that this route is viable (Kirsh, 1991; Brooks, 1997); these systems remain at approximately the level attained by the late 1980s.

Route D from unconstrained, single-task, simple systems to complex artificial systems. Evolutionary roboticists hope to follow this route[5]. However, progress is currently slow and it appears that initial optimism might have been unwarranted — the evolved implementations resulting from this approach remain at simple levels of adaptation and, even with the addition of learning mechanisms, do not seem to be scaling-up (see Nolfi, 1998, and section 2.3.2).

Route E from unconstrained,multiple-task, simple systems to complex artificial systems. A relatively unexplored area of research. Preliminary simulated experiments suggest that self-organised mechanisms are sufficient to support interleaved simple adaptive behaviours (Humphrys, 1996*b*, 1997). More research is required in order to determine whether this approach will be successful when the complexity of the required behaviours is incremented.

[5] "Insects first, people later" (Cliff, 1991*a*).

The future — self-organised action selection

Clearly the debate between pre-constrained (hierarchical and serial) *vs.* less constrained, more dynamic, parallel models of action selection will continue. Indeed, within our own research group, both approaches are advocated. Joanna Bryson is developing an architecture (EDMUND, Bryson 1999) which provides some constraint, is informed by multiple tasks, yet extends purely reactive behaviour-based architectures by: providing the potential for context and expectation by giving individual activities memory; encapsulating such memory but allowing some transfer of information between activities in order to allow for action on the basis of overall internal context; and pre-ordering behaviours in sequences, parallel activities, and prioritised sets in an attempt to combine reactivity with the ability to perform complex sequential behaviours.

Conversely, Brendan McGonigle and Mark Humphrys propose to extend Humphrys' work on reinforcement learning based action selection in simulation to the Nomad (McGonigle & Humphrys, 1998). They intend to implement a multiple-task, unconstrained, complex competence system through self-organised action selection.

It is suggested that the architecture described herein might be a good test-bed for determining whether, in fact, the limits of hand-crafted control have been reached by comparing a range of approaches to action selection.

Constrained, single task, complex competence The architecture is both hierarchically and sequentially organised. Currently tasks are defined by the designer: once a task has been placed on the task queue its execution results in sequential behaviours constrained by initialisation conditions. Assuming no unrecoverable default the sequence of behaviours will run until completion. Within behaviours the ordering of individual activities is pre-determined. Furthermore, currently no opportunity exists for a switch of task based on environmental conditions: the architecture is interruptible *within* behaviours but not *across* tasks. The w-problem (which task to execute) is resolved prior to resolution of the q-problem (which behaviour to execute).

Partially constrained, multiple task, complex competence The proposed research of McGonigle & Humphrys' (1998) can be supported by this architecture. Multiple action primitives are now in place[6], the addition of a self-organisational action selection mechanism, would allow assessment of the prospects of route E above: from unconstrained, multiple-task, simple competence systems towards complex artificial systems. McGonigle and Humphrys intend to extend Humphrys' (1997) system to a scenario where the system must address three goals, at multiple time scales, simultaneously. Tasks will be interruptible between sequential behaviours such that control of the system can be opportunistic, yet within tasks the sequential characteristics of the architecture will remain intact. Through interleaving task-level control of the system, whilst retaining seriality, McGonigle & Humphrys hope to move toward artificial intelligent systems.

Partially constrained, multiple task, complex competence The favoured approach of the author, that suggested by the biological analysis of McGonigle & Chalmers (1998*a*), and that proposed by McGonigle (1991, 1998). Complex biological systems are both constrained and unconstrained. This is, of course, the age-old issue of nativism *vs.* empiricism. The approach adopted throughout this thesis is *constructivist* in nature: constraints are imposed by embodiment and the ontological lower bounds of the system. Over ontogenesis, self-organisation of competences can occur, constrained by these ontological lower bounds, but relatively unconstrained in the terminology used above — not, that is, hand-crafted.

The current architecture might support examination of the *degree of constraint* required. Hierarchical and serial constraints can be imposed at multiple levels, and a self-organisational mechanism, whether that proposed by McGonigle & Humphrys (1998) or an internally-arbitrated trial-and-error mechanism, inspired by the biological characterisation supplied by McGonigle & Chalmers, analogous to that applied to navigation (section 5.2), can be used in order to support change over ontogenesis.

Now pre-installed arbitration procedures based on economy, together with action and behavioural primitives constrained to varying degrees might, in conjunc-

[6] Disregard the behaviours of the system for the time being.

tion with an abductive mechanism dependent on contextual interpretation of success and default, enable the system, over its life-history, to learn in which internal and external contexts actions are likely to succeed, and provide a possible route towards complex artificial systems. In this scenario (utilising McGonigle & St Johnstons' (1995) keyboard metaphor of serial control) the system will itself learn permissible 'keys' in which its behaviours can be played, and those 'melodies' which are both adaptive with respect to its multiple goals, and amenable to its designer.

6.3.4 Comparison between architectures

Recently there have been calls for a more quantitative analysis of robotic implementations (Duckett & Nehmzow, 1997, for example) with a view to determining the relative efficacy of different design strategies. However such calls are somewhat premature for three main reasons.

- Fruitful interpretation of such analyses depends upon an explicit and coherent mapping between theory and model construction which has yet to be routinely achieved within robotics. The review of traditional stances presented in chapter 1 shows that both feature imperfect or limited characterisations of intelligent biological systems that are transcribed into artificial systems which are formally isomorphic across stances in spite of their differing core theoretical assumptions. The more novel stances, dynamicism and situated cognition (described in chapter 2) enrich our understanding of intelligent systems through emphasis on self-organisation from core design primitives over ontogeny (dynamicism), the importance of context to behaviour and representation (situated cognition), and the transactional and emergent nature of intelligence (both) yet neither provide clear prescriptions, in isolation, for the design of complex artificial systems and are used to augment, rather than to supplant, the more traditional kinds of systems.
- The inaccuracy, and limitations, of the majority of characterisations of complex biological systems results in a second problem: the absence of consensual

benchmarks of achievement within the field. Across the different perspectives on intelligent systems, a variety of key features are regarded as essential or desirable and these become the targets of model construction. Such benchmarks range from the ability to maintain, manipulate and reason about symbols for the purposes of plan construction (classical and cognitive robotics) to the ability to react quickly and appropriately to dynamic external circumstances through reflexes (BB and animat robotics), some contextual knowledge (situated robotics), direct perceptuo-motor coupling to the world (connectionist, evolutionary, and animat robotics), or evolved controllers (evolutionary robotics).

Contemporary robotic implementations remain entrenched within two main camps: the symbolic representational/reasoning and the reactive/associative neither of which are based on accurate and comprehensive characterisations of complex biological systems. Our inspiration for the design of intelligent systems must surely be those which have evolved on our planet (see McGonigle, 1998, for further details) and, indeed, the perspectives discussed in chapter 3, from ethology, neuroscience and the conjoined comparative/developmental programme of McGonigle & Chalmers, provide a much richer view of situated biological intelligence. Such research suggests a number of grounded, and therefore non-arbitrary, core architectural features of intelligent systems that provide a number of benchmarks for artificial systems.

As Lashley (1951) observed, the problem of serial order is not peculiarly linguistic but rather inheres to all adaptive behaviour. The capacity for serial output of behaviour requires hierarchical and modular construction. Hierarchical organisation underlies the ability of systems to recombine behaviours to meet their adaptive needs (McGonigle & St Johnston, 1995) and, furthermore, supports multiple concurrent mechanisms operating at multiple time scales and levels of abstraction. Modularity confers robustness and permits a system to possess multiple encapsulated, and potentially logically incompatible, competences simultaneously (Fodor, 1983; Sherry & Schacter, 1987; McGonigle, 1991). Serial behaviour demands a repertoire of task-achieving behaviours together with robust and effective (Simon, 1955) action selection mechanisms which allow a system to "do the right thing at the right time" (Maes, 1989). Since biological systems must interleave high and low priority behaviours at different time scales such action selection mechanisms must be both goal-oriented and contextually appropriate. A further benefit of seriality is the conferment of contextuality upon a system thereby providing the potential for interpretation and reinterpretation of external phenomena, together with behavioural success and failure, modulated by internal state.

Intelligent systems are non-trivially autonomous — able to self-select and self regulate on the basis of both in-built and learned utility metrics such as economy, initially at behavioural, and later, epistemic levels. Systems must, therefore, be cognisant of behavioural success and default without which there can be little scope for progressive adaptation and epistemic growth over ontogenesis. Such progressive adaptation necessitates memory so that later competences do not merely recapitulate their predecessors, but extend and generalise them (Fodor & Pylyshyn, 1988), and support epistemic derivations of earlier behaviours.

Biological systems possess a range of adaptive specialisations subserved by functionally distinct neuronal regions (Gazzaniga *et al.*, 1998) which constrain the ontogenetic potential of systems thereby differentiating the simple from the complex. Ontogenetic development, constrained by a set of core perceptual, behavioural and cognitive primitives — the ontological lower bounds of the system — occurs through self-organisational processes resting on inter-system and system-environment interactions. Artificial systems too, should be extendible both through design and self-organisation.

These are the kinds of metrics which have been used to assess artificial systems throughout this thesis, and those which motivated the robotic architecture presented in chapter 4.

- Finally, quantitative analysis requires a variety of architectures which are sufficiently complex to make the exercise worthwhile. Currently the majority of implementations remain capable of only single simple behaviours, or single planning tasks with research focussed on refinement of such capabilities rather than the development of richer and more versatile systems. The key benchmarks of achievement outlined in the preceding paragraphs are not the target of the majority of robotics research and are only minimally present in a very small number of implementations.

Once the discipline of robotics has reached the stage where it features multiple architectures which incorporate many, or all, of these benchmarks, quantitative analysis can begin and we will be able to compare the kinds of engineering solutions which can support them and determine whether, as suggested by McGonigle (1987):

> "The number of solutions to the problem of complexity may be finite and may, indeed, reduce to one."

Until this point has been attained the focus of argument and research must be at the paradigmatic rather than implementational level. For this reason the thesis falls not within the Popperian (1963) falsificatory tradition, but rather adheres more closely to the analyses of Lakatos (1978) which depicts the growth of scientific knowledge within research programmes (which often rest on untested, and untestable, assumptions[7]) and Kuhn (1962) which suggests that the transitions between competing research programmes ('paradigm shifts') are often not based on falsificatory evidence.

The aim of the research presented herein is to suggest that the time has now come for a paradigm shift within both cognitive science and robotics. It is suggested that neither of the traditional approaches to intelligent systems, symbol-based nor behaviour-based, can support the development of complex artificial systems. Both of these stances can be viewed as degenerating research programmes — in terms of both their characterisation of intelligent biological systems and their application to the construction of artificial systems. This analysis is supported by evidence from biology which suggests that the core assumptions of both stances involve mischaracterisation of the indexical competences of complex systems — symbols and rules or reflexes and associations.

Currently there exist a variety of architectures motivated by competing views of adaptivity and intelligence, and many skilled programmers and engineers capable of developing successful artificial systems within given problem domains. However, it appears that neither of the traditional stances, nor more recent approaches inspired by dynamicism or situated theories, can support the development of complex artificial systems in

[7] Rather such assumptions stand or fall with the research program itself.

addition to simple ones. The problems of control, seriality, multi-tasking, learning *etc.*, tend to be addressed separately within the majority of implementations. Quantitative analysis of system performance in any one of these areas, or for individual competences, neglects a fundamental issue: which of our competing paradigms is capable of both adequately characterising biological systems *and* can support the construction of artificial systems which exhibit behaviour of similar robustness and complexity? We are currently only beginning to address this question.

This thesis has attempted to outline some of the major competing paradigms within both cognitive science and robotics in order to make a preliminary assessment of their efficacy. The conclusion of the thesis is that a new synthetic paradigm is required, motivated by the characterisation of *complex* biological systems from a dynamic perspective. The implementations derived from this stance are currently in their infancy but it is believed that this approach might be capable of supporting incrementation of systems both by design and through system self-organisation with greater extendibility than traditional approaches.

6.4 Final conclusions

This thesis has argued that the stalemate in robotics arising from cognitive and behavioural characterisations of intelligent systems might best be resolved by adoption of a new design stance. It is suggested that our understanding of intelligent systems will be optimally enhanced by a *synthetic* strategy of biological characterisation in conjunction with engineering of artificial systems. It is suggested, furthermore, that the synthetic strategy will be successful only if reverse engineering is targeted at *complex* systems for these possess design features qualitatively different from those of simple systems.

Although currently artificial intelligence has a lot more to learn from psychology than *vice versa*, intelligent systems can only be understood through their reverse engineering followed by rigorous testing of the resulting designs in complex environments. An architecture inspired by McGonigle and Chalmers' characterisation of the epistemic growth of complex biological systems was presented. It is suggested that this architecture can support scaling of the system both by design and, ultimately, through self-organisation. Some experiments in self-organised navigation are also reported which illustrate the feasibility of the synthetic intelligence stance.

Finally, a *dynamic perspective* is advocated which emphasises the transactional nature of the development of intelligence. Just as biological systems develop over ontogenesis, our artificial systems should be provided with both rich design primitives, and complex environment, which can support the self-organisation of competences over the lifespan. It is suggested that self-organisation from system lower bounds is the future of robotics — we cannot install complexity in artificial systems, but must rather provide the preconditions for its growth.

Bibliography

Ackley, D. H. & M. L. Littman (1991). 'Interactions between learning and evolution.' In C. G. Langton, J. D. Farmer, S. Rasmussen, & C. E. Taylor (editors), *Artificial Life II*. Addison-Wesley, Reading, MA, USA.

Adams, C. D. & A. Dickinson (1981). 'Instrumental responding following reinforcer devaluation.' *Quarterly Journal of Experimental Psychology*, 33B:109–121.

Agnew, N. M., K. M. Ford, & P. J. Hayes (1993). 'Expertise in context: personally constructed, socially elected, and reality-relevant?' *International Journal of Expert Systems*, 7(1).

Agre, P. E. (1988). *The dynamic structure of everyday life*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA.

Agre, P. E. & D. Chapman (1987). 'Pengi: An implementation of a theory of activity.' In *Proceedings of the sixth annual meeting of the american association for artificial intelligence*, pages 268–272, Seattle, Washington, USA. Morgan Kaufmann Publishers. also appeared in Luger (1995).

Agre, P. E. & D. Chapman (1990*a*). 'What are plans for?' In P. Maes (editor), *Designing autonomous agents: Theory and practice from biology to engineering and back*, pages 17–34. MIT Press, Cambridge, MA.

Agre, P. E. & D. Chapman (1990*b*). 'What are plans for?' In P. Maes (editor), *Designing autonomous agents: theory and practice from biology to engineering and back*, pages 17–34. MIT Press, Cambridge, MA, USA.

Aldridge, J. W., K. C. Berridge, M. Herman, & L. Zimmer (1993). 'Neuronal coding of serial order: syntax of grooming in the neostriatum.' *Psychological Science*, 4:391–395.

Anderson, J. R. (1990). *The adaptive character of thought*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Anokhin, P. K. (1978). *Beitrage zur allgemeinen Theorie des funktionellen Systems*. Fischer, Jena, Germany.

Aristotle (4th Century BC/1961). *De anima*. Oxford University Press, Oxford, England. Translated by D W Hamlyn.

Aristotle (4th Century BC/1990). *The works of Aristotle*. Great books of the western world. Encyclopaedia Britannica, London, England. Two volumes.

Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman and Hall, London, England.

Avenarius, R. H. L. (1890). *Kritik der reinen Erfahrung*, volume 2. O.R. Reisland, Leipzig, Germany.

Ayer, A. J. (1947). *Thinking and meaning*. Lewis, London, England. Inaugural lecture.

Ayer, A. J. (1971). *The foundations of empirical knowledge*. Macmillan, London, England.

Ayers, J. (1995). 'A reactive ambulatory robot architecture for operation in current and surge.' In *Proceedings of the Autonomous Vehicles in Mine Countermeasures Symposium*, pages 15–31.

Ayers, J., P. Zavracky, N. McGruer, D. Massa, V. Vorus, R. Mukherjee, & S. Currie (1998). 'A modular behavioral-based architecture for biomimetic autonomous underwater robots.' In *Proceedings of the Autonomous Vehicles in Mine Countermeasures symposium*.

Ayers, J. L. & W. J. Davis (1977). 'Neuronal control of locomotion in the lobster, homarus americanus I. motor programs for forward and backward walking.' *Journal of Comparative Physiology A*, 115:1–27.

Baerends, G. P. (1976). 'The functional organization of behaviour.' *Animal Behaviour*, 24:726–738.

Ballard, D. H. (1989). 'Reference frames for active vision.' In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1635–1641, Detroit, Michigan.

Ballard, D. H. (1991). 'Animate vision.' *Artificial Intelligence*, 48:57–86.

Ballard, D. H. (1993). 'Sub-symbolic modelling of hand-eye coordination.' In A. Broadbent (editor), *Topics in the philosophy of biology*. Blackwell, London, England.

Ballard, D. H., M. M. Hayhoe, P. K. Pook, & R. P. N. Rao (1997). 'Deictic codes for the embodiment of cognition.' *Behavioural and Brain Science*, 20(4).

Barricelli, N. A. (1962). 'Numerical testing of evolution theoires: Part 1.' *Acta Biotheoretica*, 16.

Barricelli, N. A. (1979). 'On the origin and evolution of the genetic code: 2. origin of the genetic code as a primordial collector language: The pairing-release hypothesis.' *BioSystems*.

Barsalou, L. W. (1993). 'Flexibility, structure, and linguistic vagary in concepts: manifestations of a compositional system of perceptual symbols.' In A. C. Collins, S. E. Gathercole, & M. A. Conway (editors), *Theories of memories*, pages 29–101. Lawrence Erlbaum Associates, London, England.

Barsalou, L. W. (1999). 'Perceptual symbol systems.' *Behavioural and Brain Science*, 22:577–609.

Bartlett, F. C. (1932). *Remembering: a study in experimental and social psychology*. Cambridge University Press, Cambridge, England.

Bartlett, F. C. (1958). *Thinking: an experimental and social study*. Allen and Unwin, London, England.

Barto, A. G., S. J. Bradtke, & S. P. Singh (1995). 'Learning to act using real-time dynamic programming.' *Artificial Intelligence*, 60:105–122.

Barton, S. (1994). 'Chaos, self-organization, and psychology.' *American Psychologist*, 49(1):5–14.

Barwise, J. (1987). 'Unburdening the language of thought.' *Mind and Language*, 2:82–96.

Bateson, P. P. G. (1990). 'Is imprinting a special case?' *Philosophical Transactions of the Royal Society of London. B*, 203:125–131.

Bateson, P. P. G. & H. Klopfer (editors) (1982). *Perspectives in ethology*, volume 5: Ontogeny. Plenum Press, New York, NY, USA.

Bechtel, W. & A. Abrahmsen (1991). *Connectionism and the mind: an introduction to prallel processing*. Basil Blackwell, Cambridge, MA, USA.

Beer, R. D. (1995*a*). 'Computational and dynamical languages for autonomous agents.' In R. Port & T. van Gelder (editors), *Mind as motion: explorations in the dynamics of cognition*, pages 121–147. MIT Press, Cambridge, MA, USA.

Beer, R. D. (1995*b*). 'A dynamical systems perspective on agent-environment interaction.' *Artificial Intelligence*, 72:173–215.

Beer, R. D. & J. C. Gallagher (1992). 'Evolving dynamic neural networks for adaptive behaviour.' *Adaptive Behaviour*, 1:91–122.

Belew, R. K. (1992). 'Un ruolo per la scienza dei calcolatori nella vita artificiale.' *Sistemi Intelligenti*, 2.

Bergson, H. (1944). *Creative evolution*. Random House, New York, NY, USA. Trans. Arthur Mitchell.

Berridge, K. C. & I. Q. Wishaw (1992). 'Cortex, striatum and cerebellum: control of serial order in a grooming sequence.' *Experimental Brain Research*, 90:275–290.

Bertenthal, B. I. & J. Pinto (1993). 'Complementary processes in the perception and production of human movements.' In L. B. Smith & E. Thelen (editors), *A dynamic systems approach to development: applications*. MIT Press, Cambridge, MA, USA.

Bethe, A. (1931). 'Plastizität und zentrenlehre.' In *Handbuch der normalen und patologischen Physiologie*, pages 1175–1220. Verlag von Julius Springer, Berlin, Germany.

Bickhard, M. H. & D. M. Richie (1983). *On the nature of representation: a case study of James Gibson's theory of perception*. Praeger Publishers, New York, NY, USA.

Birnbaum, L. (1991). 'Rigor mortis: a response to Nilsson's 'Logic and and artificial intelligence'.' *Artificial Intelligence*, 47:57–77.

Bolles, R. C. (1970). 'Species-specific defence reactions and avoidance learning.' *Psychological Review*, 77:32–48.

Bonner, J. T. (1988). *The evolution of complexity by means of natural selection*. Princeton University Press, Princeton, NJ, USA.

Booker, L., D. E. Goldberg, & J. H. Holland (1989). 'Classifier systems and genetic algorithms.' *Artificial Intelligence*, 40:235–282.

Bookman, M. A. (1978). 'Sensitivity of the homing pigeon to an earth-strength magnetic field.' In K. Schmidt-Koenig & W. T. Keeton (editors), *Animal Migration, Navigation and Homing*, Proceedings in Life Sciences, pages 127–134. Springer-Verlag, Berlin.

Boole, G. (1854/1951). *An investigation into the laws of thought*. Dover, New York, NY, USA. Originally published 1951.

Borenstein, J. (1987). *The nursing robot system*. Ph.D. thesis, Technion, Haifa, Israel.

Borenstien, J., H. R. Everett, & L. Feng (1996). *Navigating mobile robots: systems and techniques*. A K Peters Ltd, Wellesley, MA, USA.

Braitenburg, V. (1984). *Vehicles: experiments in synthetic psychology*. MIT Press, Cambridge, MA, USA. Reprinted 1986 as MIT Press paperback.

Breithaupt, T. & J. Ayers (1996). 'Visualization and quantitative analysis of underwater biological flow fields using suspended particles.' *Marine Behavior and Physiology*.

Bremner, G., A. Slater, & G. Butterworth (editors) (1997). *Infant development: recent advances*. Psychology Press, Hove, England.

Bretano, F. (1874). *Psychologie vom empirischen Standpunkt*. Verlag von Felix Meiner, Leipzig, Germany.

Brooks, R. A. (1985). 'A robust layered control system for a mobile robot.' Memo 864, MIT AI Laboratory, Cambridge, MA, USA.

Brooks, R. A. (1986*a*). 'Achieving artificial intelligence through the building of robots.' Internal Memo 899, MIT AI Laboratory, Cambridge, MA, USA.

Brooks, R. A. (1986*b*). 'A robust layered control system for a mobile robot.' *IEEE Journal of Robotics and Automation*, 2(1):14–23.

Brooks, R. A. (1989). 'A robot that walks; emergent behaviour from a carefully evolved network.' *Neural Computation*, 1:253–262.

Brooks, R. A. (1991*a*). 'How to build complete creatures rather than isolated cognitive simulators.' In K. VanLehn (editor), *Architectures for intelligence*, pages 225–239. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Brooks, R. A. (1991*b*). 'Intelligence without reason.' In *Proceedings of IJCAI-91*, pages 569–595.

Brooks, R. A. (1991*c*). 'Intelligence without representation.' *Artificial Intelligence*, 47:139–159.

Brooks, R. A. (1997). 'From earwigs to humans.' *Robotics and Autonomous Systems*, 20(2–4):291 -304.

Bruner, J. S. (1990). *Acts of meaning.* The Jerusalem-Harvard lectures. Harvard University Press, Cambridge, MA, USA; London, England.

Bryant, P. E. (1974). *Perception and understanding in young children.* Methuen, London, England.

Bryant, P. E. & T. Trabasso (1971). 'Transitive inferences and memory in young children.' *Nature*, 232:456–458.

Bryson, J. J. (1999). *The study of sequential and hierarchical organisation of behaviour via artificial mechanisms of action selection.* Master's thesis, University of Edinburgh, Edinburgh, Scotland.

Bryson, J. J. (In press). 'Hierarchy and sequence vs. full parallelism in reactive action selection.' In *Cognitive Science.*

Burgard, W., A. B. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, & S. Thrun (1998). 'The interactive museum tour-guide robot.' In *Proceedings of the AAAI thirteenth national conference on artificial intelligence.*

Burgess, N. & G. Hitch (1992). 'Towards a network model of the articulatory loop.' *Journal of Memory and Language*, 31:429–460.

Campbell, D. T. (1974). 'Evolutionary epistemology.' In P. A. Schlipp (editor), *The philosophy of Karl Popper*, pages 413–463. Open Court Publishers, La Salle, ILL, USA.

Capek, M. (1987). 'The philosophical significance of piaget's researches on the genesis of the concept of time.' In A. Shimony & D. Nails (editors), *Naturalistic Epistemology. A Symposium of Two Decades*, pages 91–111. D Reidel Publishing Company, Dordrecht, Holland.

Carpenter, R. H. S. (1984). *Neurophysiology.* Edward Arnold Ltd., London, England.

Cartwright, B. A. & T. S. Collett (1982). 'How honey bees use landmarks to guide their return to a food source.' *Nature*, 295:560–564.

Cartwright, B. A. & T. S. Collett (1983). 'Landmark learning in bees: Experiments and models.' *Journal of Comparative Physiology*, 151:521–543.

Casti, J. L. (1979). *Connectivity, complexity and catastrophe.* John Wiley, New York, NY, USA.

Chalmers, M. & B. O. McGonigle (1984). 'Are children any more logical than monkeys on the five-term series problem?' *Journal of Experimental and Child Psychology.*

Chalmers, M. & B. O. McGonigle (1997). 'Capturing dynamic structuralism in the laboratory.' In L. Smith (editor), *Piaget, Vygotsky and beyond*. Routledge, London, England.

Chapman, D. (1990). 'Vision, instruction and action.' Technical report TR–1085, Massachusetts Institute of Technology Artificial Intelligence Lab, Cambridge, MA, USA.

Chapman, D. & P. Agre (1987). 'Abstract reasoning as emergent from concrete activity.' In M. P. Georgeff & A. L. Lansky (editors), *Reasoning about actions and plans*. Morgan Kaufmann, Los Altos, CA, USA. Proceedings of the 1986 workshop: June 30 – July 2, 1986, Timberline, Oregon.

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

Chomsky, N. (1957). *Syntactic structures*. Janua linguarum. Series minor; 4. Mouton, The Hague, Netherlands.

Christianini, N. (1995). 'Evolution and learning: An epistemological perspective.' *Axiomathes*, (3).

Chung, K. (1960). *Markov chains with stationary transition probabilities*. Springer Verlag, Berlin, Germany.

Churchland, P. S. (1986). *Neurophilosophy: toward a unified science of the mind/brain*. MIT Press, Cambridge, MA, USA.

Clancey, W. (1991*a*). 'Book review of Israel Rosenfield, 'The invention of memory: a new view of the brain'.' *Artificial Intelligence*, 50:241–284.

Clancey, W. J. (1985). 'Heuristic classification.' *Artificial Intelligence*, 27:289–350.

Clancey, W. J. (1987). 'Review of Winograd and Flores' 'understanding computers and cognition'.' *Artificial Intelligence*, 31(2):232–250.

Clancey, W. J. (1991*b*). 'The frame of reference problem in the design of intelligent machines.' In K. VanLehn (editor), *Architectures for intelligence*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Clancey, W. J. (1993). 'Situated action: a neuropsychological interpretation (response to Vera and Simon).' *Cognitive Science*, 17(1):87–107.

Clancey, W. J. (1995). 'A boy scout, toto, and a bird: how situated cognition is different from situated robotics.' In L. Steels & R. A. Brooks (editors), *The artificial life route to artificial intelligence: building situated embodied agents*, pages 227–236. Lawrence Erlbaum Associates, Hillsdale, NJ, USA. Prepared for the 1991 NATO Workshop on Emergence, Situatedness, Subsumption, and Symbol Grounding.

Clark, A. (1989). *Microcognition: philosophy, cognitive science and parallel distributed processing*. MIT Press, Cambridge, MA, USA.

Clark, A. (1994). 'Representationalism refreshed?' Talk presented at the University of Sussex, Brighton, England, 14th June.

Clark, A. (1997). *Being there: putting brain, body, and world together again.* MIT Press, Cambridge, MA, USA.

Clark, A. & D. J. Chalmers (1998). 'The extended mind.' *Analysis*, 58(1):7–19.

Clark, A. & A. Karmiloff-Smith (1993). 'The cognizer's innards.' *Mind and Language*, 8(4):540–547.

Cliff, D. (1991*a*). 'The computational hoverfly.' In J.-A. Meyer & S. W. Wilson (editors), *From animals to animats: Proceedings of the first international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Cliff, D. (1991*b*). 'Computational neuroethology: a provisional manifesto.' In J.-A. Meyer & S. W. Wilson (editors), *From animals to animats: Proceedings of the first international conference on the simulation of adaptive behaviour*, London, England. MIT Press/Bradford Books.

Cliff, D., P. Husbands, & I. Harvey (1993*a*). 'Analysis of evolved sensory motor controllers.' In *Proceedings of the second european conference on artificial life*, pages 192–204.

Cliff, D. T. & G. F. Miller (1996). 'Co-evolution of pursuit and evasion II: simulation methods and results.' In P. Maes, M. Mataric, J. A. Meyer, J. Pollack, R. H, & S. W. Wilson (editors), *From animals to animats 4. Proceedings of the fourth international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Cliff, F., I. Harvey, & P. Husbands (1993*b*). 'Explorations in evolutionary robotics.' *Adaptive Behavior*, 2:73–110.

Cohen, H. (1988). 'How to draw three people in a botanical garden.' In *Proceedings of the seventh national conference on artificial intelligence*, pages 846–855.

Colombo, M., A. E. Eickhoff, & C. G. Gross (1993). 'The effects of inferior temporal and dorsolateral frontal lesions on serial-order behavior and visual imagery in monkeys.' *Cognitive Brain Research*, 1:211–217.

Connell, J. H. (1989). 'A colony architecture for an artificial creature.' Technical report 1151, Massachusetts Institute of Technology Artificial Intelligence Lab, Cambridge, MA, USA.

Costall, A. (1994). 'On neonatal competence: sleepless nights for representational theorists?' In P. van Geert & L. P. Mos (editors), *Annal of theoretical psychology*, volume 7. Plenum Press, New York, USA.

Craik, K. J. W. (1943). *The nature of explanation.* Cambridge University Press, Cambridge, England.

Craik, K. J. W. (1966). *The nature of psychology: a selection of papers, essays and other writings.* Cambridge University Press, Cambridge, England.

Damasio, A. R. (1989). 'Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition.' *Cognition*, 33:25–62.

Darwin, C. (1859). *On the origin of species by means of natural selection, or, the preservation of favoured races in the struggle for life*. John Murray, London, England.

Darwin, C. (1871). *The descent of man and selection in relation to sex*. John Murray, London, England.

Davies, M. (1989). 'Connectionism, modularity, and tacit knowledge.' *British Journal for the Philosophy of Science*, 40:541–555.

Davies, M. (1991). 'Concepts, connectionism, and the language of thought.' In W. Ramsey, S. P. Stich, & D. E. Rumelhart (editors), *Philosophy and connectionist theory*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Dawkins, R. (1989). *The selfish gene*. Oxford University Press, Oxford, England, second edition.

Dawkins, R. & J. R. Krebs (1979). 'Arms races between and within species.' In *Proceedings of the Royal Society of London B*, volume 205, pages 489–511.

de Lillo, C. (1994). *The logic of memory search in non-human primates (Cebus apella)*. Unpubl. phd thesis, University of Edinburgh, Edinburgh, Scotland.

de Lorenzana, J. M. A. & J. M. Ward (1987). 'On evolutionary systems.' *Behavioural Science*, 32.

de Saussure, F. (1974). *Course in general linguistics*. Owen, London, England, revised edition. Translated by Jonathan Culler.

Dennett, D. C. (1978). 'Why not the whole iguana?' *Behavioural and Brain Science*, 1:103–104.

Deutsch, J. A. (1960). *The structural basis of behaviour*. Cambridge University Press, Cambridge, England.

Dewey, J. (1896). 'The reflex arc concept in psychology.' *Psychological Review*, III:357–70.

Dickinson, A. R. (1997). *Hierarchical organisation in serial search tasks by* Cebus apella *monkeys*. Ph.D. thesis, University of Edinburgh, Edinburgh, Scotland.

Dobzhansky, B., T ans Spassky (1947). 'Evolutionary changes in laboratory cultures of d. pseudoobscura.' *Evolution*, 1:191–216.

Donnett, J. & B. O. McGonigle (1991). 'Evolving speed control in mobile robots: from blindness to kinetic vision.' In *Proceedings of the Vision Interface Conference*, Calgary.

Dorigo, M. & M. Colombetti (1994). 'Robot shaping: developing autonomous agents through learning.' *Artificial Intelligence*, 71:321–370.

Dorigo, M. & M. Colombetti (1997). *Robot shaping: an experiment in behaviour engineering.* MIT Press/Bradford Books, Cambridge, MA, USA.

Dretske, F. J. (1988). *Explaining behaviour: reasons in a world of causes.* MIT Press, Cambridge, MA, USA.

Dreyfus, H. L. (1972). *What computers can't do: a critique of artificial reason.* Harper and Row, New York, NY, USA.

Dreyfus, H. L. (1992). *What computers still can't do: a critique of artificial reason.* MIT Press, Cambridge, MA, USA.

Dreyfus, H. L. (1996). 'The current relevance of Merleau-Ponty's phenomenology of embodiment.' *The Electronic Journal of Analytic Philosophy*, 4.

Duckett, T. & U. Nehmzow (1997). 'Quantitative analysis of mobile robot localisation systems.' Technical report UMCS–97–9–1, Department of Computer Science, University of Manchester, England. also in *Proceedings TIMR UK 97* British conference on autonomous mobile robotics and "Scientfic methods in mobile robotics", Manchester, 1997.

Duckett, T. & U. Nehmzow (1999*a*). 'Knowing your place in real world environments.' In *Proceedings EUROBOT'99, 3rd European workshop on advanced mobile robots.* Zurich, Switzerland, September 6–8, 1999.

Duckett, T. & U. Nehmzow (1999*b*). 'Self-localisation and autonomous navigation by a mobile robot.' Technical report UMCS–99–3–1, Department of Computer Science, University of Manchester, England. also in *Proceedings TIMR 99 "Towards intelligent mobile robots"* Bristol, England.

Dyson, G. B. (1998). *Darwin among the machines. The evolution of global intelligence..* Perseus Books, Reading, MA, USA.

Eliasmith, C. (1996). 'The third contender: a critical examination of the dynamicist theory of cognition.' *Journal of Philosophical Psychology*, 9(4):441–463.

Elman, J. L. (1990). 'Finding structure in time.' *Cognitive Science*, 14:179–211.

Emlen, S. T. (1972). 'The ontogenetic development of orientation capabilities.' In S. R. Galler, K. Schmidt-Koenig, G. J. Jacobs, & R. E. Belleville (editors), *Animal Orientation and Navigation*, NASA SP-262, pages 191–210. US Govt. Printing Office, Washington, DC, USA.

Epicurus (4th Century BC/1994). *The Epicurus reader: selected writings and testimonia.* Hackett, Indianapolis, IN, USA. Translated by B Inwood and L P Gerson.

Estes, W. K. (1972). 'An associative basis for coding and organisation in memory.' In A. W. Melton & E. Martin (editors), *Coding processes in memory.* Winston, Washington, DC, USA.

Etienne, A. S., S. Sitbon, C. Dahn-Hurni, & R. Maurer (1994). 'Golden hamsters on the eight-arm maze in light and darkness.' *Quarterly Journal of Experimental Psychology*, 47B:401–425.

Everett, H., D. Gage, G. Gilbreth, R. Laird, & R. Smurlo (1994). 'Real-world issues in warehouse navigation.' In *Proceedings of the SPIE conference on mobile robots IX*, volume 2352, Boston, MA, USA.

Fikes, R. E., P. E. Hart, & N. J. Nilsson (1972). 'Learning and executing generalised robot plans.' *Artificial Intelligence Journal*, 3:251–288.

Fikes, R. E. & N. J. Nilsson (1971). 'STRIPS: A new approach to the application of theorem proving to artificial intelligence.' *Artificial Intelligence*, 1(2):189–208.

Firby, R. J. (1987). 'An investigation into reactive planning in complex domains.' In *Proceedings, AAAI-87*, pages 202–206.

Firby, R. J. (1989). *Adaptive execution in dynamic domains*. Ph.D. thesis, Department of Computer Science, Yale University, New Haven, CT, USA.

Flook, J. P. & B. O. McGonigle (1977). 'Serial adaptation to conflicting prismatic rearrangement effects in monkey and man.' *Perception*, 6:15–29.

Floreano, D. & F. Mondada (1994). 'Automatic creation of an autonomous agent: genetic evolution of a neural-network driven robot.' In D. Cliff, P. Husbands, J. A. Meyer, & S. W. Wilson (editors), *From animals to animats 3. Proceedings of the third international conference on the simulation of adaptive behaviour*, London, England. MIT Press/Bradford Books.

Floreano, D. & F. Mondada (1996). 'Evolution of plastic neurocontrollers for situated agents.' In P. Maes, M. Mataric, J. Meyer, J. Pollack, & S. W. Wilson (editors), *From animals to animats 4. Proceedings of the fourth international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Fodor, J. A. (1975). *The language of thought*. Thomas Y Cromwell, New York, NY, USA.

Fodor, J. A. (1983). *The modularity of mind: an essay on faculty psychology*. MIT Press, Cambridge, MA, USA.

Fodor, J. A. & Z. Pylyshyn (1988). 'Connectionism and cognitive architecture: a critical analysis.' *Cognition*, 28:3–71.

Franceschini, N. (1997). 'Combined optical, neuroanatomical, electrophysiological and behavioral studies on signal processing in the fly compound eye.' In C. Taddei-Ferretti (editor), *Biocybernetics of vision: integrative mechanisms and cognitive processes*. World Scientific, London, England.

Franceschini, N., J. M. Pichon, & C. Blances (1993). 'From insect vision to robot vision.' In H. B. Barlow, J. P. Frisby, A. Horridge, & M. Jeeves (editors), *Natural and artificial low-level seeing systems*, pages 31–42. Oxford Scientific Publications, Clarendon Press, Oxford, England.

Franceschini, N., J. M. Pichon, & C. Blances (1996). 'From fly vision to robot vision: bionics of signal processing.' In W. Nachtigall & A. Wisser (editors), *BIONA reports X*, pages 47–60. G Fischer, Stuttgart, Germany.

Franceschini, N., J. M. Pichon, & C. Blanes (1991). 'Real time visuomotor control: from fli es to robots.' In *Proceedings of the international confernence on advanced robotics*, Pisa, Italy.

Freeman, W. J. & C. A. Skarda (1990). 'Representations: who needs them?' In J. L. McGaugh, J. L. Weinberger, & G. Lynch (editors), *Brain organization and memory: cells, systems and circuits*, pages 375–380. Guilford Press, New York, USA.

Gainotti, G., M. C. Silveri, A. Daniele, & L. Giustolisi (1995). 'Neuroanatomical correlates of category-specific semantic disorders: a critical survey.' *Memory*, 3:247–264.

Gallant, J. L., C. E. Connor, H. Drury, & D. C. van Essen (1995). 'Neural responses in monkey visual cortex during free viewing of natural scences.' *Invest. Opthalml. Vis. Sci.*, 36:1052.

Gallistel, C. R. (1980). *The organisation of action: a new synthesis*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Gallistel, C. R. (1995). 'The replacement of general-purpose theories with adaptive specializations.' In M. S. Gazzaniga (editor), *The cognitive neurosciences*, pages 1255–1225. MIT Press, Cambridge, MA, USA.

Gallistel, R. (1990). *The organization of learning*. MIT Press, Cambridge, MA, USA.

Garcia, J. & R. A. Koelling (1966). 'Relation of cue to consequences in avoidance learning.' *Psychonomic Science*, 4:123–124.

Garner, W. R. (1962). *Uncertainty and structure as psychological concepts*. John Wiley, New York, NY, USA.

Gat, E. (1991*a*). 'Integrating planning and reacting in a heterogenous asynchronous architecture for mobile robots.' *SIGART Bulletin*, 2:17–74.

Gat, E. (1991*b*). *Reliable goal-directed reactive control of autonomous mobile robots*. Ph.D. thesis, Virginia Polytechnic Institute and State University, Viginia, VA, USA.

Gat, E. & G. Dorais (1994). 'Robot navigation by conditional sequencing.' In *Proceedings of the international conference on robotics and automation (ICRA)*.

Gazzaniga, M. S., R. B. Ivry, & G. R. Mangun (1998). *Cognitive neuroscience. The biology of the mind*. W W Norton & Co, London, England.

Geib, C. (1992). 'The intentional planning system: ItPlanS.' Technical report MS–CIS–92–73, Department of Computer and Information Science, University of Pennsylvania, PA, USA.

Geib, C., L. Levison, & M. Moore (1994). 'SodaJack: an architecture for agents that search for and manipulate objects.' Technical report MS–CIS–94–16, Department of Computer and Information Science, University of Pennsylvania, PA, USA.

Geschwind, N. & W. Levitsky (1968). 'Human brain: left-right asymmetries in temporal speech region.' *Science*, 161:186–187.

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin, Boston, MA, USA.

Giralt, G., R. Chatila, & M. Vaisset (1984). 'An integrated navigation and motion control system for multisensory robots.' In *Robotics research 1*, pages 191–214. MIT Press, Cambridge, MA, USA.

Glaser, W. R. (1992). 'Picture naming.' *Cognition*, 42:61–105.

Glasgow, J. I. (1993). 'The imagery debate revisited: a computational perspective.' *Computational Intelligence*, 9:309–333.

Glasspool, D. W. (1985). 'Competitive queueing and the articulatory loop: an extended network model.' In J. Levy, D. Bairaktaris, J. Bullinaria, & D. Cairns (editors), *Connectionist models of memory and language*. UCL Press, London, England.

Globus, G. G. (1992). 'Toward a noncomputational cognitive neuroscience.' *Journal of Cognitive Neuroscience*, 4(4):299–310.

Gödel, K. (1990). *Collected works*, volume 2: Publications 1938–1974. Clarendon Press, Oxford, England.

Goodwin, B. C. (1994). *How the leopard changed its spots: the evolution of complexity*. Weidenfeld and Nicolson, London, England.

Gould, J. L. (1980). 'The case for magnetic-field sensitivity in birds and bees (such as it is).' *American Scientist*, 68:256–267.

Grey Walter, W. (1950). 'An imitation of life.' *Scientific American*, 182(5):42–45.

Grey Walter, W. (1951). 'A machine that learns.' *Scientific American*, 185(5):60–63.

Grice, H. P. (1975). 'Logic and conversation.' In P. Cole & J. L. Morgan (editors), *Studies in syntax*, volume 3. Academic Press, New York, NY, USA.

Grossberg, S. (1978). 'Behavioral contrast in short term memory: serial binary memory models or parallel continuous memory models?' *Journal of Mathematical Psychology*, 17:199–219.

Guthrie, E. R. (1952). *The psychology of learning*. Harper, New York, NY, USA.

Guttmann, J.-S. (1996). *Vergleich von Algorithmen zur Selbstlokalisierung eines mobilen Roboters*. Master's thesis, University of Ulm, Ulm, Germany.

Haldane, J. S. (1917). *Organism and environment as illustrated by the physiology of breathing*. Yale University Press, New Haven, CT, USA.

Halford, G. S. (1971). 'Acquisition of conservation through learning a consistent classificatory system for quantities.' *Australian Journal of Psychology*, 23:151–159.

Halford, G. S. (1993). *Childrens' understanding. The development of mental models*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Halford, G. S., S. B. Smith, J. C. Dickson, M. T. Maybery, M. E. Kelly, J. D. Bain, & J. E. M. Stewart (1995). 'Modelling the development of reasoning strategies: the roles of analogy, knowledge, and capacity.' In T. Simon & G. S. Halford (editors), *Developing cognitive competence: new approaches to cognitive modelling*. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Halford, G. S., W. H. Wilson, & S. Phillips (1998). 'Processing capacity defined by relational complexity: implications for comparative, developmental and cognitive psychology.' *Behavioural and Brain Science*, 21(6):803–831.

Halperin, J. R. P. (1991). 'Machine motivation.' In J. Meyer & S. Wilson (editors), *From Animals to Animats: Proceedings of the First International Conference on the Simulation of Adaptive Behaviour*. MIT Press, Cambridge, MA.

Harlow, H. (1949). 'The formation of learning sets.' *Psychological Review*, 56:51–65.

Harnad, S. (1990*a*). 'Lost in the hermeneutic hall of mirrors.' *Journal of Experimental and Theoretical Artificial Intelligence*, 2:321–327.

Harnad, S. (1990*b*). 'The symbol grounding problem.' *Physica D*, 42:335–346.

Harnad, S. (1992). 'Connecting object to symbol in modelling cognition.' In A. Clark & R. Lutz (editors), *Connectionism in context*, pages 75–90. Springer-Verlag.

Harnad, S. (1993). 'Problems, problems: the frame problem as a aymptom of the symbol grounding problem.' *Psycoloquy*, 4(34).

Harnad, S. (1995). 'Grounding symbolic capacity in robotic capacity.' In L. Steels & R. A. Brooks (editors), *The 'artificial life' route to 'artificial intelligence'*, pages 276–286. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Harris, M. R. & B. O. McGonigle (1994). 'Modelling transitive inference.' *Quarterly Journal of Experimental Psychology*, 47B(3):319–348.

Harvey, I. (1992). 'Untimed and misrepresented: connectionism and the computer metaphor.' Cognitive Science Research Paper 245, Cognitive Science Department, University of Sussex.

Harvey, I., P. Husbands, & D. Cliff (1993). 'Issues in evolutionary robotics.' In *From animals to animats 2: Proceedings of the second international conference on the simulation of adaptive behaviour*, pages 364–373, London, England. Bradford Books.

Harvey, I., P. Husbands, D. Cliff, A. Thompson, & N. Jakobi (1997). 'Evolutionary robotics: the sussex approach.' *Robotics and Autonomous Systems*, 20:205–224.

Harvey, I., H. P, & D. Cliff (1994). 'Seeing the light: artificial evolution, real vision.' In D. Cliff, P. Husbands, J. A. Meyer, & S. W. Wilson (editors), *From animals to animats 3. Proceedings of the third international conference on simulation of adaptive behaviour*, Cambridge, MA, London, England. MIT PressBradford Books.

Haugeland, J. (1985). *Artificial intelligence. The very idea.* MIT Press, Cambridge, MA, USA.

Helms-Tillery, S. I., M. Flanders, & J. F. Soechting (1991). 'A coordinate system for the synthesis of visual and kinesthetic information.' *Journal of Neuroscience*, 11(3):770–778.

Henson, R. N. A., D. G. Norris, M. P. A. Page, & A. D. Baddeley (1996). 'Unchained memory: error patterns rule out chaining models of immediate serial recall.' *Quarterly Journal of Experimental Psychology*, 49A:80–115.

Hesse, M. (1988). 'Theories, family resemblances and analogy.' In *Analogical Reasoning*, pages 317–340. Kluwer Academic Publishers.

Heylighen, F. (1991). 'Cognitive levels of evolution.' In F. Geyer (editor), *The cybernetics of complex systems*, pages 75–91. Intersystems, Salinas, CA, USA.

Hinde, R. A. (1966). *Animal behaviour: A synthesis of ethology and comparative psychology*. McGraw-Hill, New York, NY, USA.

Hinde, R. A. (1982). *Ethology*. Fontana Masterguides. Fontana Press, London, England.

Hinde, R. A. & J. Stevenson-Hinde (editors) (1973). *Constraints on learning: limitations and predispositions*. Academic Press, London, England. Based on a conference held in St John's College, Cambridge, April 4th–7th, 1972.

Holland, J. H. (1975). *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor, Michigan, USA.

Holyoak, K. J. & P. Thagard (1995). *Mental leaps*. MIT Press, Cambridge, MA, USA.

Horswill, I. (1995). 'Visual routines and visual search.' In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, Montreal.

Horswill, I. (1997). 'Visual architecture and cognitive architecture.' *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2 and 3):277–293.

Houghton, G. (1990). 'The problem of serial order: a neural network model of sequence learning and recall.' In R. Dale, C. Mellish, & M. Zock (editors), *Current research in natural language generation*. Academic Press, London, England.

Houghton, G. & T. Hartley (1995). 'Parallel models of serial behaviour: Lashley revisited.' *Psyche*, 2(25). Available: http://psyche.cs.monash.edu.au/v2/psyche-2-25-houghton.html.

Hull, C. L. (1934). 'The concept of the habit-family hierarchy and maze learning.' *Psychological Review*, 41.

Hull, C. L. (1945). 'The place of innate individual and species differences in a natural science theory of behaviour.' *Psychological Review*, 52:55–60.

Hull, C. L. (1952). *A behaviour system*. Yale University Press, New Haven, CT, USA.

Humphrys, M. (1996*a*). 'Action selection in a hypothetical house robot: using those RL numbers.' In P. G. Anderson & K. Warwick (editors), *Proceedings of the first international ICSC symposia on intelligent industrial automation (IIA-96) and soft computing*, pages 216–222, Reading, England. ICSC Academic Press. March 26–28.

Humphrys, M. (1996*b*). 'Action selection methods using reinforcement learning.' In P. Maes, M. Matarić, J.-A. Meyer, J. Pollack, & S. W. Wilson (editors), *From animals to animats 4: Proceedings of the fourth international conference on the simulation of adaptive behaviour*, pages 135–144, Cambridge, MA, USA. MIT Press.

Humphrys, M. (1997). *Action selection methods using reinforcement learning*. Ph.D. thesis, University of Cambridge, Cambridge, England.

Husbands, P., I. Harvey, & D. Cliff (1995). 'Circle in the round: state space attractors for evolved sighted robots.' *Robotics and Autonomous Systems*, 15:83–106.

Hutchins, E. (1995). *Cognition in the wild*. MIT Press, Cambridge, MA, USA.

Inhelder, B. (1990). 'Introduction.' In D. Leiser & C. Gillieron (editors), *Cognitive science and genetic epistemology*. Plenum Press, New York, USA.

Inhelder, B. & J. Piaget (1964). *The early growth of logic in the child*. Routledge and Kegan Paul, London, England.

Jackendoff, R. (1983). *Semantics and cognition*, volume 8 of *Current studies in linguistics*. MIT Press, Cambridge, MA, USA.

Jalbert, J., S. Kashin, & J. Ayers (1995). 'A biologically-based undulatory lamprey-like AUV.' In *Proceedings of the Autonomous Vehicles in Mine Countermeasures symposium*, pages 39–52.

James, W. (1981). *The principles of psychology*. Harvard University Press, Harvard, MA, USA.

Janlert, L. E. (1987). 'Modelling change — the frame problem.' In Z. W. Pylyshyn (editor), *The robot's dilemma. The frame problem in artificial intelligence*. Ablex, Norwood, NJ, USA.

Jeannerod, M. (1988). *The neural and behavioural organisation of goal-directed movements*. Clarendon Press, Oxford, England.

Johnson-Laird, P. N. (1983). *Mental models*. Harvard University Press, Cambridge, MA, USA.

Jordan, M. I. (1986). 'Serial order: a parallel distributed approach.' ICI report 8604, Institute for Cognitive Science, University of California, San Diego, CA, USA.

Kaelbling, L. P. (1991). 'A situated automata approach to the design of embedded agents.' *SIGART Bulletin*, 2(4):85–88.

Kaelbling, L. P. (1993). *Learning in embedded systems*. MIT Press/Bradford Books, Cambridge, MA, USA/London, England.

Kaelbling, L. P. & S. J. Rosenschein (1990). 'Action and planning in embedded agents.' In P. Maes (editor), *Designing autonomous agents: theory and practice from biology to engineering and back*, pages 35–48. MIT Press, Cambridge, MA, USA.

Kahneman, D., P. Slovic, & A. Tversky (editors) (1982). *Judgement under uncertainty: heuristics and biases*. Cambridge University Press, Cambridge, England.

Kahneman, D. & A. Tversky (1973). 'On the psychology of prediction.' *Psychological Review*, 80:237–51.

Kandel, E. R., J. H. Schwartz, & T. M. Jessell (1991). *Principles of neural science*. Elsevier, London, England, third edition.

Karmiloff-Smith, A. (1983). 'A new abstract code or the new possibility of multiple codes?' *Behavioural and Brain Science*, 6(1):149–150.

Karmiloff-Smith, A. (1992). *Beyond modularity: a developmental perspective on cognitive science*. MIT Press, Cambridge, MA, USA.

Kauffman, S. A. (1993). *The origins of order: self-organisation and selection in evolution*. Oxford University Press, Oxford, England.

King, S. & C. Weiman (1990). 'Helpmate autonomous mobile robot navigation system.' In *Proceedings of the SPIE conference on mobile robotics*, volume 2352, pages 190–198, Boston, MA, USA.

Kirchner, F. & J. Hertzberg (1997). 'A prototype study of an autonomous robot platform for sewerage system maintenance.' In G. Bekey (editor), *Autonomous Robots*, volume 4, pages 319–333. Kluwer Academic Publishers, Amsterdam, Holland.

Kirsh, D. (1991). 'Today the earwig, tomorrow man?' *Artificial Intelligence Journal*, 47:161–184.

Knudsen, E. I. (1982). 'Auditory and visual maps of space in the optic tectum of the owl.' *Journal of Neuroscience*, 2:1177–1194.

Kodjabachian, J. & J. A. Meyer (1998). 'Evolution and development of modular control architectures for 1-D locomotion in six-legged animats.' *Connection Science*, 10(3–4).

Koenig, S. & R. Simmons (1996). 'Passive distance learning for robot navigation.' In *Proceedings of the thirteenth international conference on machine learning*, pages 266–274, San Mateo, CA, USA. Morgan Kauffman.

Köhler, W. (1927). 'Zum problem der regulation.' *Wilhelm Roux Archiv fur Entwicklungsmechanik der Organismen*, 112:315–322. Reprinted in Köhler (1971). Trans. Mary Henle and Erich Goldmeier.

Köhler, W. (1971). *Selected Works*.

Kohonen, T. (1982). 'Self-organized formation of topologically correct feature maps.' *Biological Cybernetics*, 43:59–69.

Koza, J. R. (1990). *Genetic programming: On the programming of computers by means of natural selection*. Complex adaptive systems. MIT Press, Cambridge, MA, USA.

Kramer, G. (1950). 'Eine neue methode zur enforschung der zugorientierung und die bisher damit erzielten ergebnisse.' *Proc. A. Intern. Ornithol. Congr. Uppsala*, pages 269–280.

Kravitz, E. A. (1988). 'Hormonal control of behaviour: amines and the biasing of behavioral output in lobsters.' *Science*, 241:1775–1781.

Kripke, S. (1982). *Wittgenstein on rules and private language*. Blackwell, Oxford, England.

Kripke, S. A. (1980). *Naming and necessity*. Harvard University Press, Cambridge, MA, USA.

Kuhn, T. (1962). *The structure of scientific revolutions*. Cambridge University Press, Cambridge, England.

Laird, J. E. & P. S. Rosenbloom (1990). 'Integrating execution, planning, and learning in SOAR for external environments.' In *Proceedings of the eighth national conference on artificial intelligence*, pages 1022–1029, Menlo Park, CA, USA. The AAAI Press.

Lakatos, I. (1978). *The methodology of scientific research programmes*, volume 1 of *Philosophical papers*. Cambridge University Press, Cambridge, England.

Lakoff, G. & M. Johnson (1980). *Metaphors we live by*. University of Chicago Press, Chicago, IL, USA.

Längle, T. & T. C. Lüth (1995). 'Decentralized control of distributed intelligent robots and subsystems.' In *AIRTC Symposium on artificial intelligence in real-time control*, pages 447–452. Oct. 3–5, Valencia, Spain.

Langton, C. G. (editor) (1987). *Proceedings of Artificial Life*. Addison-Wesley. appeared 1989.

Langton, C. G. (1995). *Artificial life: An overview*. MIT Press, Cambridge, MA, USA.

Lashley, K. S. (1929). *Brain mechanisms and intelligence: a quantitative study of injuries to the brain*. Behaviour research fund monographs. University of Chicago Press, Chicago, IL, USA.

Lashley, K. S. (1951). 'The problems of serial order in behaviour.' In L. A. Jeffries (editor), *Cerebral mechanisms in behaviour*, pages 112–136. John Wiley, New York, NY, USA.

Lawrenson, W. & P. E. Bryant (1972). 'Absolute and relative codes in young children.' *Journal of Child Psychology and Psychiatry*, 13:25–35.

Le Gare, M. (1987). 'The use of general system theory as metatheory for developing and evaluating theories in the neurosciences.' *Behavioural Science*, 32:106–120.

Lee, D. C. (1995). *The map-building and exploration strategies of a simple, sonar-equipped mobile robot: an experimental, qualitative evaluation*. Ph.D. thesis, University College London, London, England.

Lee, M. & B. O. McGonigle (1996). *Cognitive modelling and the design of artificially intelligent systems*. Final Report MRC/JCI Project G9110835.

Lenat, D. B. & E. A. Feigenbaum (1991). 'On the thresholds of knowledge.' *Artificial Intelligence*, 47:185–250.

Lespérance, Y., K. Tam, & M. Jenkin (1998). 'Reactivity in a logic-based programming framework.' In *Working papers of the AAI symposium on cognitive robotics*.

Levesque, H. & R. Reiter (1998). 'High-level robotic control: beyond planning. position paper.' In *Working papers of the AAI symposium on cognitive robotics*.

Levison, L. (1996). *Connecting planning and acting via object-specific reasoning*. Ph.D. thesis, Computer and Information Science Department, University of Pennsylvania School of Engineering and Applied Science, PA, USA.

Lewis, R., S. Huffman, B. John, J. Laird, J. Lehman, A. Newell, P. Rosenbloom, T. Simon, & S. Tessler (1990). 'Soar as a unified theory of cognition.' In *Proceedings of the thirteenth annual conference of the Cognitive Science Society*, pages 1035–1042, Pittsburgh, MA, USA. Cognitive Science Society Incorporated.

Lifschitz, V. (1984). 'Some results on circumscription.' In *AAAI 1984*.

Linney, J. (1991). 'Review of reactive planning literature.' Technical report 560, Deparment of Computer Science, Queen Mary and Westfield College, London, England.

Longo, G. O. (1994). 'From golem to Gödel and back.' *Nuova Civilta delle Macchine*, (12):28–38.

Lorenz, K. (1935). 'Der kumpan in der umwelt des vogels.' *J. Ornithol.*, 83:137–213.

Lorenz, K. (1973). *Behind the mirror*. Harcourt, Brace, Jovanovich, New York.

Lu, F. & E. Milios (1997). 'Globally consistent range scan alignment for environment mapping.' *Autonomous Robots*, 4:333–349.

Luger, G. F. (1994). *Cognitive science. The science of intelligent systems*. Academic Press, San Diego, CA, USA.

Luger, G. F. (1995). *Computation and intelligence: Collected readings*. AAAI Press, MIT Press, Menlo Park, California.

Luger, G. F. & C. Stern (1999). 'Nsf grant proposal, summer 1999, extending bayesian techniques to abductive error interpretation.'

Luger, G. F. & W. A. Stubblefield (1998*a*). *Artificial intelligence and the design of expert systems*. Benjamin/Cummings, California, first edition.

Luger, G. F. & W. A. Stubblefield (1998*b*). *Artificial intelligence. Structures and strategies for problem solving*. Addison Wesley, Longman Inc., Harlow, England.

Lund, H. H., J. Hallam, & W. P. Lee (1997). 'Evolving robot morphology.' In *Proceedings of the IEEE 4th international conference on evolutionary computation*. IEEE Press.

Lund, J., H H an Hallam (1996). 'Sufficient neurocontrollers can be surprisingly simple.' Technical report 824, Department of Artificial Intelligence, University of Edinburgh, Edinburgh, Scotland.

Lycan, W. (1993). 'A deductive argument for the representational theory of thinking.' *Mind and Language*, 8(3):404–422.

Macfarlane, D. A. (1930). 'The role of kinesthesis in maze learning.' *University of California Publications in Psychology*, 4:277–305.

Mach, E. (1959). *The analysis of sensations, and the relation of the physical to the psychical*. Dover Publications, New York, NY, USA. Translation of: Die Analyse der Empfindungen und das Verhaltnis der Physischen zum Psychischen.

Mackay, D. G. (1987). *The organization of perception and action*. Springer Verlag, New York, NY, USA.

Mackintosh, N. J. (1983). *Conditioning and associative learning*. Oxford University Press, Oxford, England.

Macphail, E. M. (1982). *Brain and intelligence in vertebrates*. Oxford science publications. Clarendon Press, Oxford, England.

Maes, P. (1989). 'How to do the right thing.' *Connection Science*, 1:291–323.

Maes, P. (1990). *Designing automomous agents: Theory and practice from biology to engineering and back*. MIT Press, Cambridge, MA, USA.

Maes, P. & R. A. Brooks (1990). 'Learning to coordinate behaviours.' In *AAAI-90*, pages 796–802, Boston, MA, USA.

Maier, N. R. F. & T. C. Schneirla (1935). *Principles of animal psychology*. McGraw-Hill, New York, NY, USA.

Malcolm, C., T. Smithers, & J. Hallam (1989). 'An emerging paradigm in robot architecture.' In T. Kanade & F. C. A. Groen (editors), *Intelligent Autonomous Systems*, volume 2.

Malcolm, C. A. (1987). *Planning and performing the robotic assembly of soma cube constructions*. MSc dissertation, University of Edinburgh, Edinburgh, Scotland.

Malcolm, C. A. (1997). 'A hybrid behavioural/knowledge-based approach to robotic assembly.' Research paper RP875, Department of Artificial Intelligence. University of Edinburgh, Edinburgh, Scotland.

Malcolm, C. A. (1998). 'The SOMASS system: a hybrid symbolic and behaviour-based system to plan and execute assemblies by robot.' Research paper RP899, Department of Artificial Intelligence. University of Edinburgh, Edinburgh, Scotland.

Marler, P. (1991). 'Song-learning behaviour: the interface with neuroethology.' *Trends in Neuroscience*, 14:199–206.

Marr, D. (1982). *Vision: a computational investigation into the human representation and processing of visual information*. W.H. Freeman, San Francisco, CA, USA.

Matarić, M. J. (1990). 'Environment learning using a distributed representation.' In *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*, pages 402–406.

Matarić, M. J. (1991). 'Behavioral synergy without explicit integration.' *SIGART Bulletin*, 2(4):130–133.

Matarić, M. J. (1992). 'Behavior-based control: main properties and implications.' In *Procedings, IEEE international conference on robotics and automation, Workshop on architectures for intelligent control systems*, pages 46–54, Nice, France. IEEE.

Matarić, M. J. & R. a. Brooks (1990). 'Learning a distributed map representation based on navigation behaviours.' USA-Japan Symposium on Flexible Automation.

Matarić, M. J. & D. Cliff (1996). 'Challenges in evolving controllers for physical robots.' *Robotics and autonomous systems*, 1:67–83.

Maturana, H. R. (1983). 'What is it to see?' *¿Que es ver?*, 16:255–269.

Maturana, H. R. & F. J. Varela (1980). *Autopoiesis and cognition: The realization of the living*. Boston Studies in the Philosophy of Science, Vol. 42. D Reidel Publishing Company, Dordrecht, Holland. With a preface to 'Autopoiesis' by Sir Stafford Beer.

McCarthy, J. (1963). 'Situations, actions and causal laws.' Memo 2, Stanford Artificial Intelligence project, Stanford, CA, USA.

McCarthy, J. (1980). 'Circumscription – a form of nonmonotonic reasoning.' *Artificial Intelligence*, 13.

McCarthy, J. (1986). 'Applications of circumscription to formalizing common sense knowledge.' *Artificial Intelligence*, 28:89–116.

McCarthy, J. & P. J. Hayes (1969). 'Some philosophical problems from the standpoint of artificial intelligence.' In B. Meltzer & D. Michie (editors), *Machine Intelligence, 4*, pages 463–502.

McClelland, J. L., D. E. Rumelhart, & the PDP research group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*, volume 1. MIT Press, Cambridge, MA, USA.

McCormick, P. (1991). *AARON's code: meta-art, artificial intelligence, and the work of Harold Cohen*. W H Freeman and Company, New York, NY, USA.

McCulloch, W. S. & W. Pitts (1943). 'A logical calculus of the ideas immanent in nervous activity.' *Bulletin of Mathematical Biophysics*, 5:115–137.

McDermott, D. (1987). 'A critique of pure reason.' *Computational Intelligence*, 3:151–160.

McDermott, D. V. (1978). 'Planning and acting!' *Connection Science Journal, Special Issue on Hybrid Systems*, 2(2):71–109.

McFarland, D. (1974). 'Time-sharing as a behavioural phenomena.' In D. S. Lehrman, J. S. Rosenblatt, R. A. Hinde, & E. Shaw (editors), *Advances in the study of behaviour*, volume 4. Academic Press, New York, NY, USA.

McFarland, D. (1999). *Animal behaviour*. Longman, Harlow, England, third edition.

McGonigle, B. O. (1987). 'Non-verbal thinking by animals.' *Nature*, 325:110–112.

McGonigle, B. O. (1991). 'Incrementing intelligent systems by design.' In J. Meyer & S. Wilson (editors), *From Animals to Animats, Proceedings of the First International Conference on the simulation of adaptive behaviour*. MIT Press., London, England.

McGonigle, B. O. (1995). *Giving a robot a life history*. Talk and video demonstration at the JCI Conference, Royal College of Pathologists, London.

McGonigle, B. O. (1998). 'Autonomy in the making: Getting robots to control themselves.' In O. Holland & D. McFarland (editors), *Artificial ethology*. Oxford University Press, Oxford, England. Invited contribution. International symposium, Lanzarote, 1998, sponsored by OUP.

McGonigle, B. O. (1999). 'Review of Liben — development of map understanding.'

McGonigle, B. O. & M. Chalmers (1977). 'Are monkeys logical?' *Nature*, 267:694–696.

McGonigle, B. O. & M. Chalmers (1984). 'The selective impact of question form and input mode on the symbolic distance effect in children.' *Journal of Experimental Child Psychology*, 37:525–554.

McGonigle, B. O. & M. Chalmers (1992). 'Monkeys are rational!' *Quarterly Journal of Experimental Psychology*, 45B(3):189–228.

McGonigle, B. O. & M. Chalmers (1993). 'An experimental analysis of ordering skills in children.' Esrc project grant final report, British Library, London, England.

McGonigle, B. O. & M. Chalmers (1994). 'Monkeys are rational!' *Quarterly Journal of Experimental Psychology*, 47B(3):189–228.

McGonigle, B. O. & M. Chalmers (1996). 'The ontology of order.' In L. Smith (editor), *Critical readings on Piaget*, pages 279–311. Routledge, London, England.

McGonigle, B. O. & M. Chalmers (1998*a*). 'Rationality as optimised cognitive self regulation.' In M. Oaksford & N. Chater (editors), *Rational models of cognition*. Oxford University Press, Oxford, England.

McGonigle, B. O. & M. Chalmers (1998*b*). 'Rationality as optimised cognitive self regulation.' In M. Oaksford & N. Chater (editors), *Rational models of cognition*. Oxford University Press, Oxford, England.

McGonigle, B. O. & M. Chalmers (In press*a*). 'Cognition in monkeys and men.' In S. Fountain (editor), *Cognition in relation to the biomedical sciences*. Sage.

McGonigle, B. O. & M. Chalmers (In press*b*). *The growth of intelligence in complex systems*. MIT Press, Cambridge, Mass. London, England.

McGonigle, B. O., C. De Lillo, & A. J. Dickinson (1992). 'Serial order induced search in children and monkeys.' Paper presented at the 5th European Conference on Developmental Psychology, Seville, Spain. 3–7 September.

McGonigle, B. O., C. De Lillo, & A. J. Dickinson (1994). 'Classification to order: a comparative analysis of categorical seriation in monkey and man.' Paper presented XVth Congress of the International Primatological Society, Bali, Indonesia, 12–16th August.

McGonigle, B. O. & M. Humphrys (1998). 'Action selection in a multi-goal mobile robot by self-organising automatic internal competition.'

McGonigle, B. O. & V. Jaswal (1993). 'Categorical seriation in children: a preliminary report.' Research report, Laboratory for Cognitive Neuroscience and intelligent systems, University of Edinburgh, Edinburgh, Scotland.

McGonigle, B. O. & B. Jones (1975). 'The perception of linear gestalten by rat and monkey: sensory sensitivity or the perception of structure?' *Perception*, 4:419–429.

McGonigle, B. O. & B. Jones (1978). 'Levels of stimulus processing by the squirrel monkey: relative and absolute judgements compared.' *Perception*, 7:635–659.

McGonigle, B. O. & J. Ravenscroft (pers. comm.). 'Emergent transitivity and seriation in *cebus apella*.' Ongoing research program at the Laboratory for Cognitive Neuroscience and Intelligent Systems, University of Edinburgh, Scotland.

McGonigle, B. O. & B. St Johnston (1995). 'Applying staged cognitive growth principles to the design of autonomous robots.' Internal memo, Intelligent Systems Laboratory, Department of Psychology, University of Edinburgh.

Menzies, S. (1996). 'Assessing responses to situated cognition.' In *Proceedings of the 10th knowledge acquiisition workshop for knowledge-based systems*, Banff, Canada.

Merleau-Ponty, M. (1962). *Phenomenology of perception*. Routledge, London, England. Orignially publ. 1945.

Meyer, J.-A. & A. Guillot (1991). 'Simulation of adaptive behaviour in animats: review and prospect.' In J.-A. Meyer & S. W. Wilson (editors), *From animals to animats: Proceedings of the first international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Meyer, J.-A. & A. Guillot (1994). 'From SAB90 to SAB94: four years of animat research.' In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (editors), *From animals to animats 3: Proceedings of the third international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Michalski, R. (1986). In *Machine learning: An artificial intelligence approach*, volume 2. Morgan-Kauffman.

Miller, G. A. (1956). 'The magical number seven, plus or minus two: some limits on our capacity for processing information.' *Psychological Review*, 63:81–97.

Miller, G. A. (1966). *Psychology: the science of mental life*. Penguin, Harmondsworth, England.

Miller, G. A., E. Galanter, & K. H. Pribham (1960). *Plans and the structure of behaviour*. Holt, Rinehart and Winston, New York, NY, USA.

Miller, G. F. & P. M. Todd (1990). 'Exploring adaptive agency I: theory and methods for simulating the evolution of learning.' In D. S. Touretzky, J. L. Elman, T. J. Sejnowski, & G. E. Hinton (editors), *Proceedings of the 1990 connectionist models summer school*, San Mateo, CA, USA. Morgan Kauffman.

Milner, A. D. & M. A. Goodale (1995). *The visual brain in action*. Oxford University Press, Oxford, England.

Minsky, M. (1986). *The society of mind*. Simon and Schuster, New York, NY, USA.

Minsky, M. & S. Papert (1969). *Perceptrons*. MIT Press, Cambridge, MA, USA.

Mishkin, M. & H. L. Petri (1984). 'Memories and habits: Some implications for the analysis of learning and retention.' In I. R. Squire & N. Butters (editors), *Neuropsychology of memory*. Guilford Press, New York, NY, USA.

Mitchell, T. (1996). *What we have learned about learning*. Keynote address, AAAI.

Mithen, S. (1996). *The prehistory of the mind*. Thames & Hudson, London, England.

Mittelstaedt, H. & M. L. Mittelstaedt (1982). 'Homing by path integration.' In F. Papi & H. G. Wallraff (editors), *Avian navigation*, pages 290–297. Springer-Verlag, Berlin.

Moravec, H. P. (1982). 'The Stanford CART and the CMU rover.' In *Procedings of the IEEE*, volume 71, pages 872–884.

Moravec, H. P. (1988). 'Sensor fusion in certainty grids for mobile robots.' *AI Magazine*, 9:61–74.

Morgan, C. L. (1894). *An introduction to comparative psychology*. Scott, London, England.

Moyer, R. S. (1973). 'Comparing objects in memory: evidence suggesting an internal psychophysics.' *Perception and Psychophysics*, 13(2):180–184.

Nehmzow, U. (1991). *Experiments in competence aquisition for autonomous mobile robots*. Ph.D. thesis, University of Edinburgh, Edinburgh, Scotland.

Nehmzow, U. (1995). 'Animal and robot navigation.' *Robotics and Autonomous Systems*, 15(1–2):71–81.

Nehmzow, U. & B. O. McGonigle (1992). *Demonstration of the Edinburgh R2 showing the first experimental verification of the differential light compass*. NATO Spring school on biology and technology of intelligent autonoumous agents, Trento, Italy.

Nehmzow, U. & B. O. McGonigle (1993). 'Navigation mobiler roboter mittels differential-lichtkompass.' *Hybride und intergrierte Ansatze zur Raumrepraentation und uhre Andewung*. Berlin. German journal for Artificial Intelligence, "KI". ISBN 0933-1875.

Nehmzow, U. & B. O. McGonigle (1995). 'Achieving rapid adaptations in robots by means of external tuition.' In D. Cliff, P. Husbands, J. Arcady-Meyer, & S. W. Wilson (editors), *From Animals to Animats 3. Proceedings of the Third International Conference on Simulation of Adaptive Behaviour*. MIT Press, Cambridge, MA, USA.

Nemhzow, U., T. Smithers, & B. O. McGonigle (1993). 'Increasing behavioural repertoire in a mobile robot.' In J. Meyer, H. Roitblat, & S. W. Wilson (editors), *From Animals to Animats 2. Proceedings of the Second International Conference on Simulation of Adaptive Behaviour*. MIT Press, Cambridge, Massachusetts; London, England.

Newell, A. (1990). *Unified theories of cognition*. Harvard University Press, Cambridge, MA, USA.

Newell, A., J. C. Shaw, & H. Simon (1957). 'Empirical explorations with the Logic Theory machine.' In *Proceedings of the western joint computer conference*, volume 15, pages 218–329.

Newell, A., J. C. Shaw, & H. Simon (1959). 'A general problem solving program for a computer.' *Computers and Automation*, 8(7):10–16.

Newell, A., J. C. Shaw, & H. A. Simon (1958). 'Elements of a theory of human problem solving.' *Psychological Review*, 65:151–166.

Newell, A. & H. Simon (1976). 'Computer science as empirical enquiry: symbols and search.' *Communications of the ACM*, 19(3):113–126.

Nilsson, N. J. (1980). *Principles of artificial intelligence*. Tioga, Palo Alto, CA, USA.

Nilsson, N. J. (1984). 'Shakey the robot.' Technical Report 323, Stanford Research Institute AI Centre, Stanford, CA, USA.

Nolfi, S. (1997). 'Evolving non-trivial behavior on autonomous robots: adaptation is more powerful than decomposition and integration.' In T. Gorni (editor), *Evolutionary robotics: from intelligent robots to artificial life*. AAI Books, Kanata, Canada.

Nolfi, S. (1998). 'Evolutionary robotics: exploting the full power of self-organization.' *Connection Science*, 10:3–4.

Nolfi, S., D. Floreani, O. Miglino, & F. Mondada (1994). 'How to evolve autonomous robots: different approaches in evolutionary robotics.' In R. A. Brooks & P. Maes (editors), *Proceedings of the fourth international conference on artificial life*, Cambridge, MA, USA. MIT Press.

Nolfi, S. & D. Floreano (In press). 'Co-evolving predators and prey robots: do 'arms races' arise in artificial evolution.' In *Artificial Life*.

Nolfi, S. & O. Miglino (In press). 'Studying the emergence of grounded representations: exploiting the power and the limits of sensory-motor coordination.' In G. Mulhauser (editor), *Evolving consciousness*, Advances in consciousness research series. John Benjamins Press, Amsterdam, Holland.

Nolfi, S. & D. Parisi (1997). 'Learning to adapt to changing environments in evolving neural networks.' *Adaptive Behavior*, 1:99–105.

Oaksford, M. & N. Chater (1998*a*). 'An introduction to rational models of cognition.' In M. Oaksford & N. Chater (editors), *Rational models of cognition*, chapter 1, pages 1–18. Oxford University Press, London, England.

Oaksford, M. & N. Chater (editors) (1998*b*). *Rational models of cognition.* Oxford University Press, London, England.

Paivio, A. (1975). 'Perceptual comparisons through the mind's eye.' *Memory and Cognition*, 3(6):635–647.

Paivio, A. (1986). *Mental representations: a dual coding approach.* Oxford University Press, Oxford, England.

Passingham, R. E. (1982). *The human primate.* W H Freeman, Oxford, England.

Patil, R. (1981). 'Causal representation of patient illness for electrolyte and acid-base diagnosis.' Technical report MIT/LCS/TR–267, Massachusetts Institute of Technology, Cambridge, MA, USA.

Pellizzer, G., P. Sargent, & A. P. Georgopoulos (1994). 'Motor cortex and visuomotor memory scanning.' In *Abstracts, 24th Annual Meeting, Society for Neuroscience*, page 983 (403.12), Miami Beach, FL, USA.

Pfeifer, R. (1996). 'Symbols, patterns, and behaviour: towards a new understanding of intelligence.' AI lab report, Department of Computer Science, University of Zurich, Zurich, Switerland.

Pfeifer, R. (1997). 'Building "fungus eaters": design principles of autonomous agents.' *From animals to animats 4. Proceedings of the fourth international conference on the simulation of adaptive behaviour.*

Piaget, J. (1950). *The psychology of intelligence.* Routledge and Kegan Paul, London, England. Translated by M Piercy and D E Berlyne, originally published 1947.

Piaget, J. (1953). *The origins of intelligence in the child.* Routledge and Kegan Paul, London, England.

Piaget, J. (1955). *The child's construction of reality.* Routledge and Kegan Paul, London, England.

Piaget, J. (1970). *Genetic epistemology.* Columbia University Press, New York, NY, USA.

Piaget, J. (1971). *Biology and knowledge.* Edinburgh University Press, Edinburgh, Scotland.

Pierce, C. S. (1958). *Collected papers 1931–1958.* Harvard University Press, Cambridge, MA, USA.

Plato (1926). *Cratylus; Parmenides; Greater Hippias; Lesser Hippias.* Heinemann, London, England.

Plotkin, H. C. (1988). 'Behaviour and evolution.' In H. C. Plotkin (editor), *The role of behaviour in evolution.* MIT Press, Cambridge, MA, USA.

Poincaré, J. H. (1946). *The foundations of science: Science and hypothesis, the value of science, science and method.* Science and education, Vol. 1. The Science Press, Lancaster, PA. Translated by G. B. Halsted.

Polyani, M. (1976). 'Life's irreducible structure.' In M. Grene & E. Mendelsohn (editors), *Topics in the Philosophy of Biology*. D Reidel Publishing Company, Boston, MA, USA.

Popper, K. H. (1963). *Conjectures and refutations: The growth of scientific knowledge*. Routledge and Kegan Paul, London, England.

Popper, K. R. (1959). *The logic of scientific discovery*. Hutchinson, London, England.

Popper, K. R. (1972). *Objective knowledge: an evolutionary approach*. Clarendon Press, Oxford, England.

Popper, K. R. (1985). 'Thought and experience and evolutionary epistemolgy.' In *¿Atti del convegno Che cos'e' il pensiero?*, Accademia Nazionale dei Lincei. Unita' dell'essere.

Potts, G. (1972). 'Information processing strategies used in the encoding of linear orderings.' *Journal of Verbal Learning and Verbal Behaviour*, 11:727–740.

Potts, G. R., W. S. Banks, S. Kosslyn, R. S. Moyer, C. Riley, & K. H. Smith (1978). 'Comparative judgements.' In N. J. Castellan & F. Restle (editors), *Cognitive theory*, volume 3. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Prem, E. (1997). 'Epistemic autonomy in models of living systems.' In P. Husbands & I. Harvey (editors), *Fourth european conference on artificial life*, pages 2–9, Cambridge, MA, USA. MIT Press. Held July 28–31, Brighton, England.

Premack, D. (1983). 'The codes of man and beasts.' *Behavioural and Brain Science*, 6:125–167.

Prigogine, I. (1980). *From being to becoming: time and complexity in the physical sciences*. W H Freeman, San Francisco, CA, USA.

Putnam, H. (1975). 'Philosophy and our mental life.' In H. Putnam (editor), *Mind, language and reality: philosophical papers*, volume 2. Cambridge University Press, Cambridge, England.

Pylyshyn, Z. W. (editor) (1987). *The robot's dilemma. The frame problem in artificial intelligence*. Ablex, Norwood, NJ, USA.

Quillian, R. (1968). 'Semantic memory.' In M. Minsky (editor), *Semantic information processing*. MIT Press, Cambridge, MA, USA.

Reeke, G. N. & G. M. Edelman (1988). 'Real brains and artificial intelligence.' *Daedalus*, 117(1). 'Artificial intelligence' issue.

Rescorla, R. A. (1971). 'Variations in the effectivenss of reinforcement and nonreinforcement following prior inhibitory conditioning.' *Learning and Motivation*, 2:113–123.

Rey, G. (1995). 'A not 'merely empirical' argument for a language of thought.' In J. Tomberlin (editor), *Philosophical perspectives*, volume 9, pages 201–222.

Ring, M. B. (1998). 'CHILD: a first step towards continual learning.' In S. Thrun & L. Y. Pratt (editors), *Learning to learn*. Kluwer Academic Publishers.

Roberts, L. G. (1963). 'Machine perception of three-dimensional solids.' Technical report 315, MIT Lincoln Laboratory, Cambridge, MA, USA.

Robertson, S. S., A. H. Cohen, & R. G. Mayer-Kess (1993). 'Behavioural chaos: beyond the metaphor.' In L. B. Smith & E. Thelen (editors), *A dynamic systems approach to development: applications*, pages 120–150. MIT Press, Cambridge, MA.

Roitblat, R. (1991). 'Mechanism and process in animal behaviour: models of animals, animals as models.' In D. Cliff, P. Husbands, J. Meyer, & S. W. Wilson (editors), *From animals to animats 3: Proceedings of the third international conference on the simulation of adaptive behaviour*, pages 12–21, Brighton, England.

Romanes, G. J. (1882). *Animal Intelligence*. Kegan, Paul, Trench, London, England.

Rommetveit, R. (1987). 'Meaning, context, and control: convergent trends and controversial issues in current social-scientific research on human cognition and communication.' *Inquiry*, 30:77–79.

Rosenbaum, D. A. (1991). *Human motor control*. Academic Press, San Diego, USA.

Rosenblatt, K. & D. Payton (1989). 'A fine-grained alternative to the subsumption architecture for mobile robot control.' In *Proceedings of the IEEE/INNS International Joint Conference on Neural Networks*.

Rosenfield, I. (1988). *The invention of memory: a new view of the brain*. Basic Books, New York, NY, USA.

Rosenschein, S. & L. P. Kaelbling (1986). 'The synthesis of digital machines with provable epistemic properties.' In J. V. Halpern (editor), *Proceedings of the 1986 conference on theoretical aspects of reasoning about knowledge*, pages 83–98, San Mateo, CA, USA. Morgan Kauffmann.

Rumelhart, D. E., G. E. Hinton, & R. J. Williams (1986). 'Learning representation by back-propagation of erros.' *Nature*, 323:533–536.

Russell, B. (1956/88). *Logic and knowledge: essays 1901–1950*. Unwin Hyman, London, England.

Rutkowska, J. C. (1993). *The computational infant: looking for developmental cognitive science*. Harvester Wheatsheaf, Hemel Hempstead, England.

Rutkowska, J. C. (1994). 'Emergent functionality in human infants.' In D. Cliff, P. Husbands, J.-A. Meyer, & S. W. Wilson (editors), *From animals to animats 3: Proceedings of the third international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Rutkowska, J. C. (1997). 'Reassing Piaget's theory of sensorimotor intelligence: a view from cognitive science.' Cognitive science research paper CSRP–369, School of Cognitive and Computing Sciences, University of Sussex, Brighton, England. Also appeared in (Bremner *et al.*, 1997).

Samuel, A. L. (1983). 'AI: where it has been and where it is going.' In *IJCAI'83*. Keynote address.

Schmajuk, N. A. & A. D. Thierne (1992). 'Purposive behaviour and cognitive mapping: a neural network model.' *Biological Cybernetics*, 6:165–174.

Schmidhuber, J., J. Zhao, & N. N. Schraudolph (1998). 'Reinforcement learning with self-modifying policies.' In S. Thrun & L. Y. Pratt (editors), *Learning to learn*. Kluwer Academic Publishers.

Schmidt-Koenig, K. (1979). *Avian Orientation and Navigation*. Academic Press, London, England.

Schoppers, M. J. (1987). 'Universal plans for reactive robots in unpredictable domains.' In *Proceedings, IJCAI-87*, pages 1039–1046, Menlo Park, CA, USA.

Searle, J. R. (1980*a*). 'Intrinsic intentionality.' *Behavioural and Brain Science*, 3:450–457.

Searle, J. R. (1980*b*). 'Minds, brains, and programs.' *Behavioural and Brain Science*, 3:417–457.

Searle, J. R. (1992). *The Rediscovery of the Mind*. MIT Press, London, England.

Seidenberg, M. & J. L. McClelland (1989). 'A distributed developmental model of word recognition and naming.' *Psychological Review*, 98:523–568.

Seifert, L. S. (1997). 'Activating representations in permanent memory: different benefits for pictures and words.' *Journal of Experimental Psychology: Learning, Memory and Cogntion*, 23:1106–1121.

Seligman, M. E. P. (1970). 'On the generality of the laws of learning.' *Psychological Review*, 77:406–418.

Shallice, T. (1972). 'Dual functions of consciousness.' *Psychological Review*, 79:383–393.

Shallice, T., D. Glasspool, & G. Houghton (1995). 'Can neuropsychological evidence inform connectionist modelling? analyses of spelling.' *Language and Cognitive Processes*, 10:195–225.

Shanahan, M. P. (1996). 'Robotics and the common sense informatic situation.' In *Working notes for "Common sense 96", the third symposium on logical formulations of common sense*.

Shanahan, M. P. (1997). *Solving the frame problem, a mathematical investigation of the common sense law of intertia*. MIT Press, Cambridge, MA, USA.

Shanahan, M. P. (1998). 'Reinventing Shakey.' In *Working Papers of the AAI symposium on cognitive robotics*.

Shannon, B. (1987). 'The non-abstractness of mental representations.' *New Ideas in Psychology*, 5:117–126.

Shannon, C. E. & W. Weaver (1949). *The mathematical theory of communication.* University of Illinois Press, Urbana, IL, USA.

Shatkay, H. & L. Kaelbling (1997). 'Learning topological maps with weak local odometric information.' In *Proceedings of the fifteenth international joint conference on artificial intelligence*, pages 920–927, San Mateo, CA, USA. Morgan Kauffman.

Sherry, D. F. & D. L. Schacter (1987). 'The evolution of multiple memory systems.' *Psychological Review*, 94:439–454.

Shimony, A. (1987). 'Introduction.' In A. Shimony & D. Nails (editors), *Naturalistic Epistemology. A Symposium of Two Decades*, pages 1–13. D Reidel Publishing Company, Dordrecht, Holland.

Shorey, H. H. (1976). *Animal communication by pheromones.* Academic Press, New York, NY, USA.

Simmons, R. (1990). 'Concurrent planning and execution for a walking robot.' Technical report CMU–RI–TR–90–16, Robotics Institute, Carnegie Mellon University.

Simmons, R. (1991). 'Coordinating planning, perception and action for mobile robots.' *SIGART Bulletin*, 2:156–159.

Simon, H. (1955). 'A behavioral model of rational choice.' *Quarterly Journal of Economics*, 69:99–118.

Simon, H. (1962). 'The architecture of complexity.' *Proceedings of the American Philosophical Society*, 26:467–482.

Skarda, C. A. & W. J. Freeman (1987). 'How brains make chaos in order to make sense of the world.' *Behavioural and Brain Science*, 10:161–195.

Skinner, B. F. (1938). *The behavior of organisms: an experimental analysis.* Appleton Century, New York, NY, USA.

Skinner, B. F. (1953). *Science and human behaviour.* MacMillan, New York, NY, USA.

Skinner, B. F. (1957). *Verbal behaviour.* The Century psychology series. Appleton-Century-Crofts, New York, NY, USA.

Sloman, A. (1998). 'The "semantics" of evolution: trajectories and trade-offs in design space and niche space.' In *IBERAMIA 98*. Springer-Verlag. Invited talk for IBERAMIA'98, Lisbos, October 1998.

Small, W. S. (1901). 'Experimental study of the mental processes of the rat ii.' *American Journal of Psychology*, 12:206–239.

Smith, B. C. (1991). 'The owl and the electronic encyclopedia.' *Artificial Intelligence*, 47:251–288.

Smith, L. (1993). *Necessary knowledge: Piagetian perspectives on constructivism.* Lawrence Erlbaum Associates, Hove, England.

Smolensky, P. (1988). 'On the proper treatment of connectionism.' *Behavioural and Brain Science*, 11:1–74.

Snyder, L. H. & R. A. Andersen (1994). 'Effects of vestibular and neck proprioceptive signals on visual responses in posterior parietal cortex.' In *Abstracts, 24th Annual Meeting, Society for Neuroscience*, page 1278 (525.1), Miami Beach, FL, USA.

Sokal, R. R. (1974). 'Classification: purposes, principles, progress, prospects.' *Science*, 185:1115–1123.

Sparks, D. L. & J. S. Nelson (1987). 'Sensory and motor maps in the mammalian superior colliculus.' *Trends in Neuroscience*, 10:312–317.

Spence, K. W. (1949). 'Theoretical interpretations of learning.' In S. S. Stevens (editor), *Handbook of experimental psychology*, pages 690–730. John Wiley, London, England.

Spencer, H. (1855). *The principles of psychology*. Longman, Brown, Green and Longmans, London, England.

Srinivas, S. (1977). *Error recovery in robot systems*. Ph.D. thesis, California Institute of Technology, California, USA.

Steels, L. (1991). 'Towards a theory of emergent functionality.' In J.-A. Meyer & S. W. Wilson (editors), *From animals to animats: Proceedings of the first international conference on the simulation of adaptive behaviour*, Cambridge, MA, USA. MIT Press.

Stern, C. & G. F. Luger (1993). 'Abduction and abstraction in diagnosis: a schema-based account.' In Feltovich, Ford, & Hoffman (editors), *Expertise in context*. MIT Press, Cambridge, MA, USA.

Strawson, G. (1994). *Mental reality*. MIT Press, Cambridge, MA, USA.

Stricanne, B., J. Xing, P. Mazzoni, & R. A. Andersen (1994). 'Response of LIP neurons to auditory targets for saccadic eye movements: a distributed coding for sensorimotor transformation.' In *Abstracts, 24th Annual Meeting, Society for Neuroscience*, page 143 (65.1), Miami Beach, FL, USA.

Suchman, L. (1987). *Plans and situated actions*. Cambridge University Press, New York, NY, USA.

Swain, G. P., J. Ayers, & M. E. Selzer (1995). 'Metamorphosis of spinal-projecting neurons in thge brain of the sea lamprey during transformation of the larva to adult.' *J. Comp. Neurol.*, 362:453–467.

Swenson, R. & M. T. Turvey (1991). 'Thermodynamics reasons for perception-action cycles.' *Ecological Psychology*, 3:317–348.

Tagaris, G., S.-G. Kim, R. Menon, J. Strupp, P. Andersen, K. Ugurbil, & Georgopoulos (1994). 'High field (4 tesla) functional MRI of mental rotation.' In *Abstracts, 24th Annual Meeting, Society for Neuroscience*, page 353 (152.10), Miami Beach, Fl, USA.

Tarski, A. (1944). 'The semantic conception of truth and the foundations of semantics.' *Philosophical and Phenomenological Research*, 4:341–376.

Tarski, A. (1956). *Logic, semantics, metamathematics*. Oxford University Press, London, England.

Terrace, H. & B. O. McGonigle (1994). 'Serial learning in pigeon, monkey and child.' In R. Gallistel (editor), *Psychological Sciences*. In press.

Thagard, P. (1992). *Conceptual revolutions*. Princeton University Press, Princeton, NJ, USA.

Thelen, E. & L. B. Smith (1994). *A dynamic systems approach to the development of cognition and action*. Bradford Books, London, England.

Thorndike, E. L. (1898). 'Animal intelligence: and experimental study of the associative processes in animals.' *Psychological Review*. Monograph supplement no. 8.

Thornton, C. (1997). 'A no-strings representation theory for adaptive researchers.' Cognitive science research paper CSRP–451, School of Cognitive and Computing Sciences, University of Sussex, Brighton, England.

Thrun, S. (1998*a*). 'Bayesian landmark learning for mobile robot localization.' *Machine Learning*.

Thrun, S. (1998*b*). 'Lifelong learning algorithms.' In S. Thrun & L. Y. Pratt (editors), *Leaning to learn*. Kluwer Academic Publishers.

Thrun, S., W. Burgard, & D. Fox (1998). 'A probabilistic approach to concurrent mapping and localization for mobile robots.' *Machine Leaning* and *Autonomous Robots* (joint issue), 31(5):1–25.

Thrun, S. & J. O'Sullivan (1998). 'Clustering learning tasks and the selective cross-transfer of knowledge.' In S. Thrun & L. Y. Pratt (editors), *Learning to learn*. Kluwer Academic Publishers, Amsterdam, Holland.

Thrun, S. & L. Y. Pratt (editors) (1998). *Learning to learn*. Kluwer Academic Publishers.

Tinbergen, N. (1950). 'The hierarchical organisation of nervous mechanisms underlying instinctive behaviour.' *Symposia of the Society for Experimental Biology*, 4:305–312.

Tinbergen, N. (1951). *The study of instinct*. Clarendon Press, Oxford, England.

Tinbergen, N. (1952). 'The curious behaviour of the stickleback.' Bound in pamphlet volume with title "Scientific American offprints", San Fransisco, CA, USA.

Tinbergen, N. (1960). 'The evolution of behaviour in gulls.' Bound in pamphlet volume with title "Scinetific American offprints", San Fransisco, CA, USA.

Tolman, E. C. (1932). *Purposive behaviour in animals and men*. Appleton–Century–Crofts, New York, NY, USA.

Tolman, E. C. (1948). 'Cognitive maps in rats and men.' *Psychological Review*, 55(4):189–208.

Tolman, E. C. & C. H. Horcik (1930). "Insight' in rats.' *University of California Publications in Psychology*, 4:215–232.

Trabasso, T., C. A. Riley, & E. G. Wilson (1975). 'The representation of linear and spatial strategies in reasoning: A developmental study.' In R. Falmagne (editor), *Reasoning: representation and process in children and adults*, pages 201–229. Lawrence Erlbaum Associates, Hillsdale, NJ, USA.

Trevelyan, J. P. & M. Nelson (1987). 'Adaptive robot control incorporating automatic error recovery.' In *Proceedings of the International conference on advanced robotics*.

Turing, A. (1948/1970). 'Intelligent machinery.' In B. Meltzer & D. Michie (editors), *Machine intelligence 5*, pages 3–23. American Elsevier Publishing, New York, NY, USA.

Turing, A. A. (1950). 'Computing machinery and intelligence.' *Mind*, 59:433–460.

Turvey, M. T. & C. Carello (1981). 'Cognition: the view from ecological realism.' *Cognition*, 10:313–321.

Tversky, A. (1969). 'Intransitivity of preferences.' *Psychological Review*, 76:31–48.

Tyrrell, T. (1993). 'The use of hierarchies for action selection.' *Adaptive Behavior*, 1(4):387–420.

Tyrrell, T. (1994). 'An evaluation of Mae's bottom-up mechanism for behavior selection.' *Adaptive Behaviour*, 2(4):307–348.

Uexküll, J. v. (1864/1944). 'Environment (umwelt) and inner world of animals.' In G. M. Burghardt (editor), *The foundations of comparative ethology*, pages 222–245. Van Nostrand Reinhold, New York, NY, USA. Translated by C J Mellor and D Grove. Reprinted from von Uexküll (1934).

Ullman, S. & W. Richards (1984). *Image understanding*. Ablex, Norwood, NJ, USA.

Urzelai, J., D. Floreano, M. Dorigo, & M. Colombetti (1998). 'Incremental robot shaping.' *Connection Science*, 10(3–4).

van Brakel, J. (1992). 'The complete description of the frame problem.' *Psycholoquy*, 3(60).

van Brakel, J. (1993). 'Unjustified coherence.' *Psycholoquy*, 4(23).

van Geert, P. (1994). *Dynamic systems of development: change between complexity and chaos*. Harvester Wheatsheaf, New York, USA.

van Gelder, T. (1993). 'What might cognition be if not computation?' Research report 75, Cognitive Sciences, Indiana University.

van Gelder, T. (1995). 'What might cognition be if not computation?' *Journal of Philosophy*, 91:345–381.

van Gelder, T. (1998). 'The dynamical hypothesis in cognitive science.' *Behavioural and Brain Science*, 21(5):616–665.

van Gelder, T. & R. Port (1995). 'It's about time: an overview of the dynamical approach to cognition.' In *Mind as motion: Explorations in the dynamics of cognition*. MIT Press, Cambridge, MA.

Varela, F. J. (1979). *Principles of biological autonomy*. The North Holland series in general systems research; Volume 2. North Holland.

Varela, F. J. (1989). *Connaitre - Les sciences cognitives: tendances et perspectives*. Editions du Seuil, Paris, France.

Varela, F. J., E. Thompson, & E. Rosch (1993). *The embodied mind: cognitive science and human experience*. MIT Press, Cambridge, MA, USA.

Vera, A. H. & H. A. Simon (1993). 'Situated action: a symbolic interpretation.' *Cognitive Science*, 17:7–48.

Viola, P. A. (1990). *Adaptive gaze control*. Ph.D. thesis, Massachusetts Institute of Technology, Cambridge, MA, USA.

von Bertalanffy, L. (1967). *General system theory: Foundations, development, applications*. Allen Lane, London, England. Originally publ. Braziller, NY, 1969.

von Helmholtz, H. (1873). *Popular lectures on scientific subjects*. Longmans, Green, London, England. Translated by E. Atkinson.

von Neumann, J. (1958). *The computer and the brain*. Mrs Hepsa Ely Silliman memorial lectures. Yale University Press, New Haven.

von Uexküll, J. (1921). *Umwelt und innenwelt der tiere*. Berlin, Germany.

von Uexküll, J. (1928). *Theoretische biologie*. Suhrkamp, Frankfurt, Germany.

von Uexküll, J. (1934). 'Streifzuge durch die umwelten von tieren und menschen.' Springer, Berlin. Translated in C H Schiller (editor) *Instinctive behaviour* Methuen: London.

Vygotsky, L. S. (1934/1961). *Thought and language*. MIT Press, Cambridge, MA, USA.

Vygotsky, L. S. (1978). *Mind in society: the development of higher psychological processes*. Harvard University Press, Cambridge, MA, USA.

Waddington, C. H. (1969). 'Paradigm for an evolutionary process.' In C. H. Waddington (editor), *Towards a theoretical biology*, volume 2: Sketches. Edinburgh University Press, Edinburgh, Scotland.

Wagner, G. P. & L. Altenberg (1996). 'Complex adaptations and the evolution of evolvability.' *Evolution*, 50:967–976.

Walcott, B. & C. Walcott (1982). 'A search for magnetic field receptors in animals.' In F. Papi & H. G. Wallraff (editors), *Avian Navigation*, pages 338–343. Springer-Verlag, Berlin.

Wallraff, H. G. (1978). 'Proposed principles of magnetic field perception in birds.' *Oikos*, 30:188–194.

Wallraff, H. G. (1984). 'Migration and navigation in birds: a present-state survey with some digressions to related fish behaviour.' In M. et al (editor), *Mechanisms of Migration in Fishes*. Plenum Press, New York, NY, USA.

Warren, J. M. (1965). 'Primate learning in comparative perspective.' In H. F. Schrier, H. Harlow, & F. Stollnitz (editors), *Behaviour of non-human primates: modern research trends*, volume 1, pages 249–281.

Wason, P. C. (1966). 'Reasoning.' In B. Foss (editor), *New horizons in psychology*, pages 135–151. Penguin, Harmondsworth, Middlesex.

Wason, P. C. (1977). 'Self-contradictions.' In P. N. Johnson-Laird & P. C. Wason (editors), *Thinking: readings in cognitive science*. Cambridge University Press, Cambridge, England.

Wason, P. C. & P. N. Johnson-Laird (1972). *Psychology of reasoning: structure and content*. Harvard University Press, Cambridge, MA, USA.

Watson, J. B. (1913). 'Psychology as the behaviourist views it.' *Psychological review*, 20:158–177.

Watson, J. B. (1925). *Behaviourism*. Kegan Paul, Trench, Trubner, London, England.

Watson, J. B. (1929). *Psychology from the standpoint of a behaviourist*. J B Lippincott, London, England, third edition. Revised.

Wehner, R. & F. Räber (1979). 'Visual and spatial memory in desert ants *cataglyphis bicolor*.' *Experientia*, 35.

Wehner, R. & M. V. Srinivasan (1981). 'Searing behaviour of desert ants, genus *cataglyphis* (FORMICIDAE, hymenoptera).' *Journal of Comparative Physiology*, 142:315–338.

Weiss, G. & E. von Puttkamer (1995). 'A map based on laserscans without geometric interpretation.' In *Intelligent autonomous systems-4 (IAS-4)*, pages 403–407, Karlsühe, Germany.

Weiss, S. M. & C. A. Kulikowski (1979). 'EXPERT: a system for developing consultation nodes.' In *Proceedings of the sixth international joint conference on artificial intelligence*, pages 942–947, Stanford, CA, USA.

Wheeler, M. (1995). *The philosophy of situated activity*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, England.

Whitehead, A. N. & B. Russell (1925). *Principia mathematica*. Cambridge University Press, Cambridge, England, second edition.

Whitfield, P. (1994). *The image of the world: 20 centuries of world maps*. Pomegranate Artbooks, San Francisco, CA, USA.

Wickelgren, W. A. (1969). 'Context sensitive coding, associative memory, and serial order in (speech) behaviour.' *Psychological Review*, 76:1–15.

Willatts, P. (1989). 'Development of problem-solving in infancy.' In A. Slater & J. G. Bremner (editors), *Infant development*. Lawrence Erlbaum Associates, London, England.

Williamson, A., D. D. Spencer, & G. M. Shepherd (1993). 'Comparisons between the membrane and synaptic properties of human and rodent dentate granule cells.' *Brain Research*, 622:194–202.

Wiltschko, W. & R. Wiltschko (1975). 'The interaction of stars and magnetic field in the orientation system of night migrating birds.' *Z. Tierpsychol.*, 37:337–355. and Vol. 39:265–282.

Wiltschko, W. & R. Wiltschko (1976). 'Interrelation of magnetic compass and star orientation in night migrating birds.' *Journal of Comparative Physiology*, 109(A):91–99.

Wimsatt, W. (1972). 'Complexity and organisation.' In K. Schaffner & R. Cohen (editors), *Studies in the philosophy of sciences*, volume 20, pages 67–86. Reidel Publishing Company, Dordrecht, Holland.

Winograd, T. & F. Flores (1986). *Understanding computers and cognition*. Ablex, Norwood, NJ, USA.

Winston, P. H. (1972). 'The MIT robot.' In B. Meltzer & D. Michie (editors), *Machine Intelligence, 7*, pages 431–463. John Wiley and Sons, New York, NY, USA.

Winston, P. H. (1984). *Artificial intelligence*. Addison-Wesley, Reading, MA, USA, second edition.

Wittgenstein, L. (1953). *Philosophical investigations*. MacMillan, New York, NY, USA.

Wyatt, J. (1995). 'Issues in putting reinforcement learning onto robots.' DAI Research Paper 784, Department of Artificial Intelligence, University of Edinburgh.

Yamauchi, B. & P. Langley (1997). 'Place recognition in dynamic environments.' *Journal of Robotics Systems*, 14(2):107–120. Special issue on mobile robots.

Yamauchi, B., A. Schultz, & W. Adams (1998). 'Mobile robot exploration and map-building with continuous localization.' In *Proceedings of the 1998 IEEE international conference on robotics and automation*, pages 3715–3720, Leuven, Belgium.

Yamauchi, B. M. & R. D. Beer (1993). 'Sequential behaviour and learning in evolved dynamical neural networks.' Technical report CES–93–25, Case Western University, Cleveland, Ohio, USA.

Yerkes, R. M. (1916). 'The mental lives of monkeys and apes.' *Behaviour Monographs*, 3(1).

Appendix A

Flow diagrams: acts

A.1 Sensing

A.1.1 compass-discrepancy(desired) [theta]

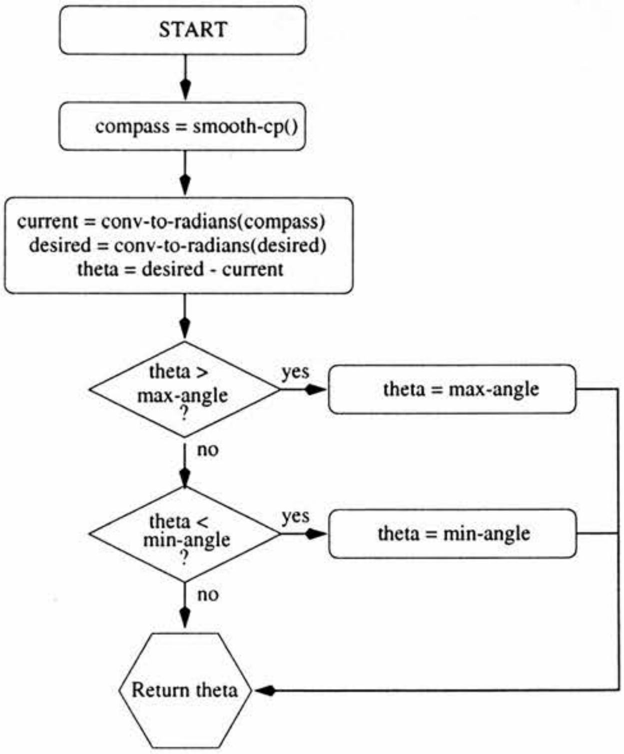

Figure A.1: compass-discrepancy(desired)[theta]

`compass-discrepancy(desired) [theta]` This function (see Figure A.1) calculates and returns the discrepancy (`theta`) between the current bearing and a `desired` bearing.

A.1.2 get-angle(x,y)[θ]

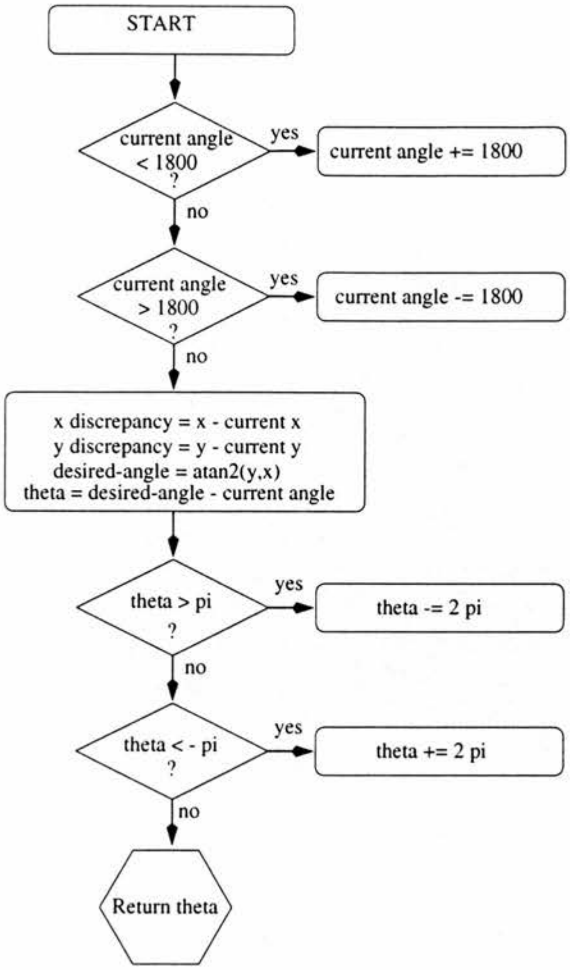

Figure A.2: get-angle(x,y)[theta]

get-angle(x,y)[θ] This routine (see Figure A.2) calculates and returns a new steering angle (θ), based on the current position, and the position of the target (x,y).

The external state variables **max-theta** and **min-theta** are used in the case where θ falls outside the range **min-theta** $< \theta <$ **max-theta**.

A.1.3 get-swerve()[θ]

get-swerve()[θ] Returns a steering angle based on proximity to lateral objects (see Figure A.3).

A.1.4 ir-object-left()[0,1]

ir-object-left()[0,1] Determines whether an object is detected by the left-side IR

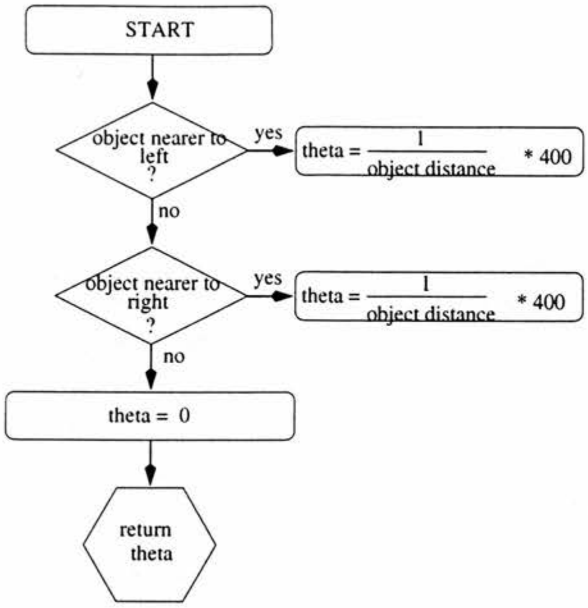

Figure A.3: get-swerve()[theta]

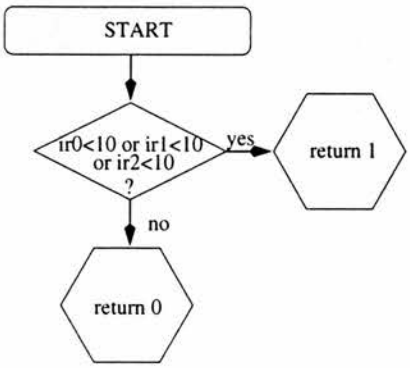

Figure A.4: ir-object-left()[0,1]

sensors. Returns 1 if an object is present, otherwise returns 0 (see Figure A.4).

A.1.5 `ir-object-right()[0,1]`

`ir-object-right()[0,1]` Determines whether an object is detected by the right-side IR sensors. Returns 1 if an object is present, otherwise returns 0 (see Figure A.5).

A.1.6 `object-left(threshold)[0,1]`

`object-left(threshold)[0,1]` Determines whether an object is detected by the left-side sonar sensors (see Figure A.6). Returns 1 if an object is below threshold distance, otherwise returns 0.

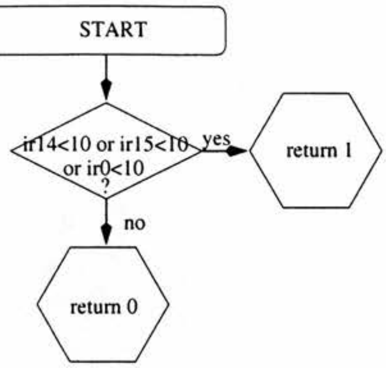

Figure A.5: ir-object-right()[0,1]

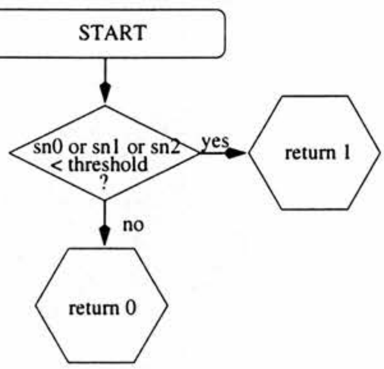

Figure A.6: object-left(threshold)[0,1]

A.1.7 `object-right(threshold)[0,1]`

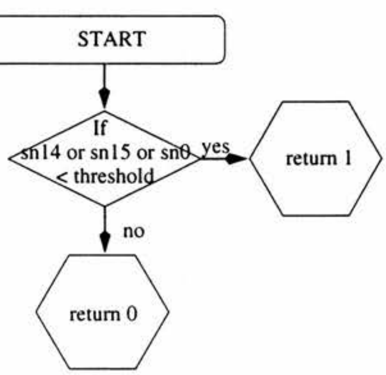

Figure A.7: object-right(threshold)[0,1]

`object-right(threshold)[0,1]` Determines whether an object is detected by the right-side sonar sensors (see Figure A.7). Returns 1 if an object is below threshold distance, otherwise returns 0.

A.1.8 `obstacle(threshold)[0,1]`

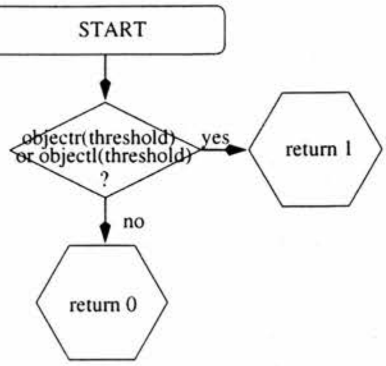

Figure A.8: obstacle(threshold)[0,1]

`obstacle(threshold)[0,1]` Returns 1 if an obstacle is detected within (optional **`threshold`**) distance or the default value of 20 inches, 0 otherwise (see Figure A.8).

A.2 Orientation and alignment

A.2.1 `align-left(distance)[1]`

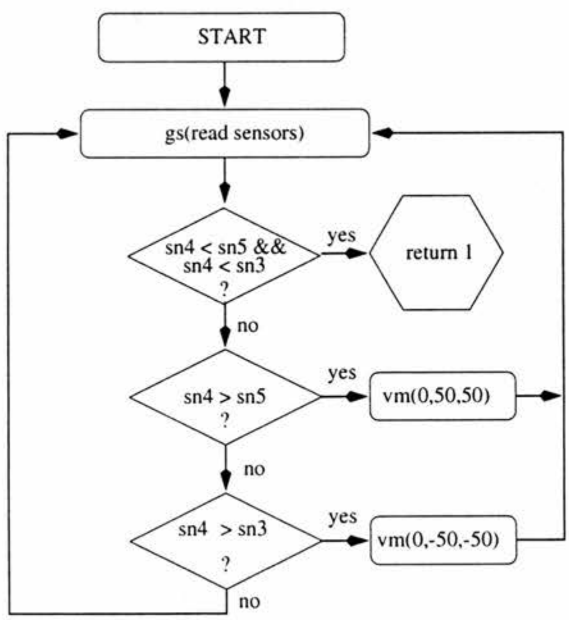

Figure A.9: align-left(distance)[1]

`align-left(distance)[1]` Aligns the Nomad to a rear surface (see Figure A.9), returning 1 on completion.

A.2.2 align-rear(distance)[1]

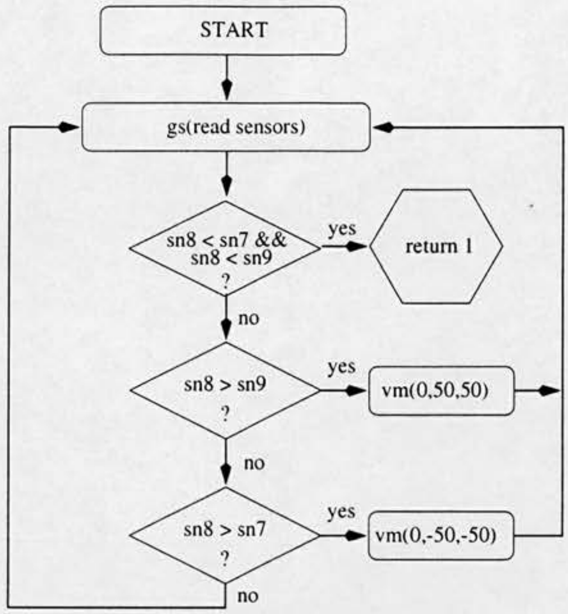

Figure A.10: align-rear(distance)[1]

`align-rear(distance)[1]` Aligns the Nomad to a rear surface (see Figure A.10), returning 1 on completion.

A.2.3 align-right(distance)[1]

`align-right(distance)[1]` Aligns the Nomad to a right-hand surface (see Figure A.11), returning 1 on completion.

A.2.4 conv-to-degrees(rads)[theta]

`conv-to-degrees(rads)[theta]` This function (see Figure A.12) is passed a value in radians (`rads`) and returns the equivalent value in degrees (`theta`).

A.2.5 conv-to-radians(theta)[rads]

`conv-to-radians(theta)[rads]` This function (see Figure A.13) is passed a value in degrees (`theta`) and returns the equivalent value in radians (`rads`).

A.2.6 fine-align(bearing)[1]

`fine-align(bearing)[1]` This function (see Figure A.14) is called after `turn-to-bearing(bearing]`

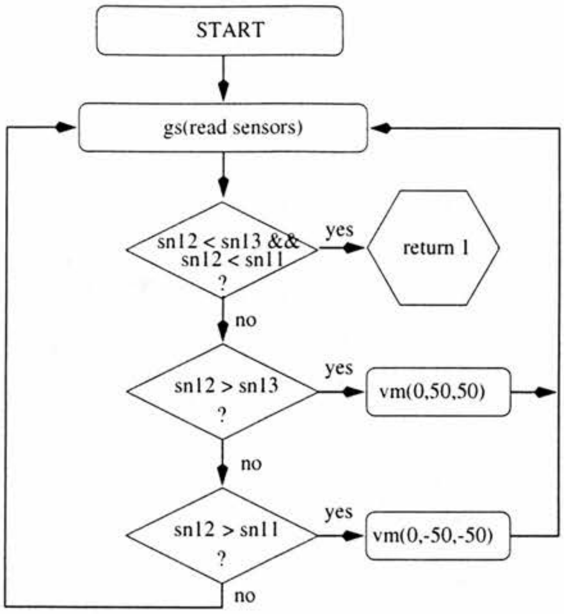

Figure A.11: align-right(distance)[1]

and rotates base and turret very slowly (1 $\frac{inch}{second}$) towards the desired **`bearing`**, returning 1 on completion.

A.2.7 `orient-left(distance)[1]`

`orient-left(distance)[1]` Moves the nomad to a position **`distance`** inches away from a left-side surface (see Figure A.15), returning 1 on completion.

A.2.8 `orient-rear(distance)[1]`

`orient-rear(distance)[1]` Moves the nomad to a position **`distance`** inches away from a rear surface (see Figure A.16), returning 1 on completion.

A.2.9 `orient-right(distance)[1]`

`orient-right(distance)[1]` Moves the nomad to a position **`distance`** inches away from a right-side surface (see Figure A.17), returning 1 on completion.

A.2.10 `turn-to-bearing(bearing)[disc]`

`turn-to-bearing(bearing)[disc]` This function (see Figure A.18) rotates turret and base to ± 1 degree of the supplied `bearing`, returning the final discrepancy (`disc`) on completion.

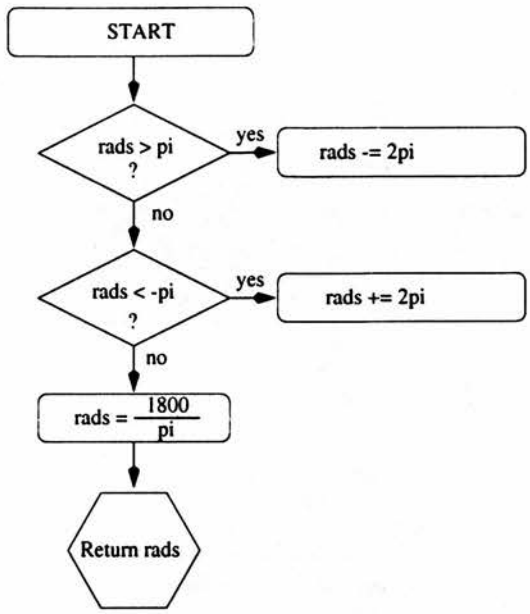

Figure A.12: conv-to-degrees(rads)[theta]

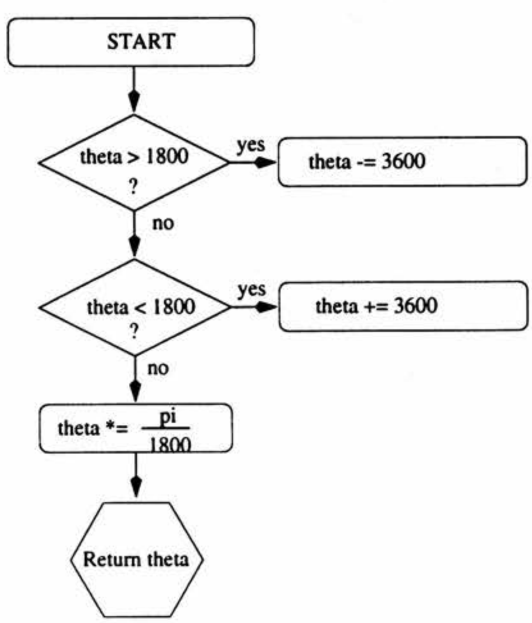

Figure A.13: conv-to-radians(theta)[rads]

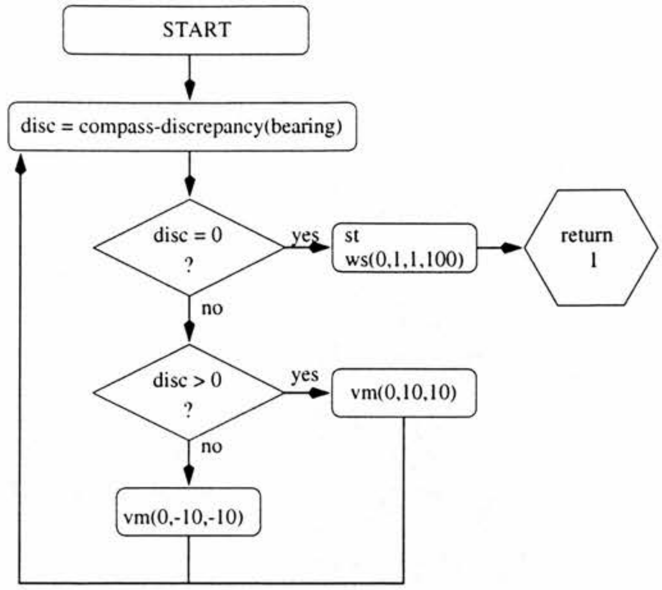

Figure A.14: fine-align(bearing)[1]

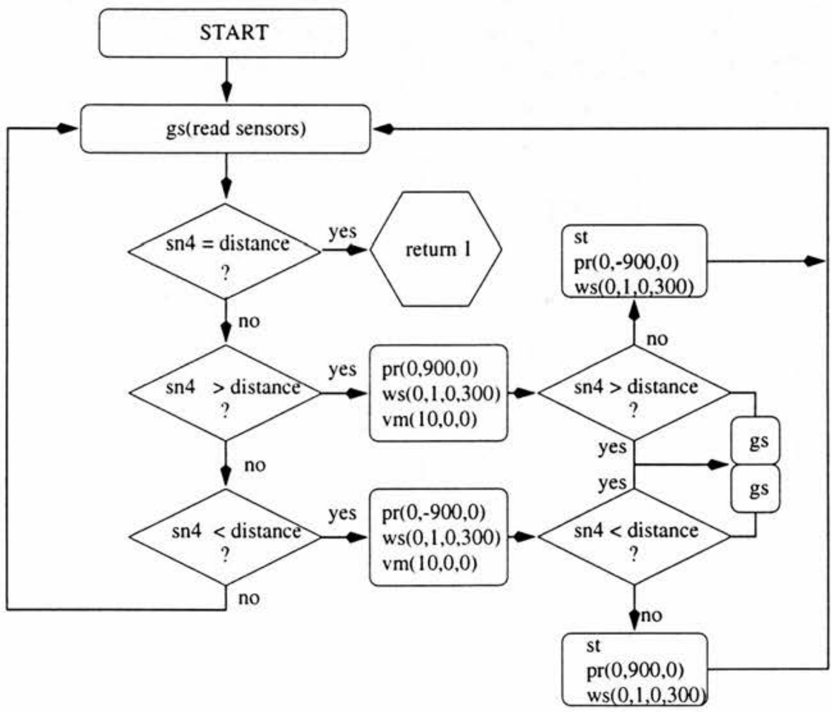

Figure A.15: orient-left(distance)[1]

A.3 Avoidance

A.3.1 `avoidthreshold[]`

`avoid(threshold)[]` `threshold` is an optional parameter, the default value is 20

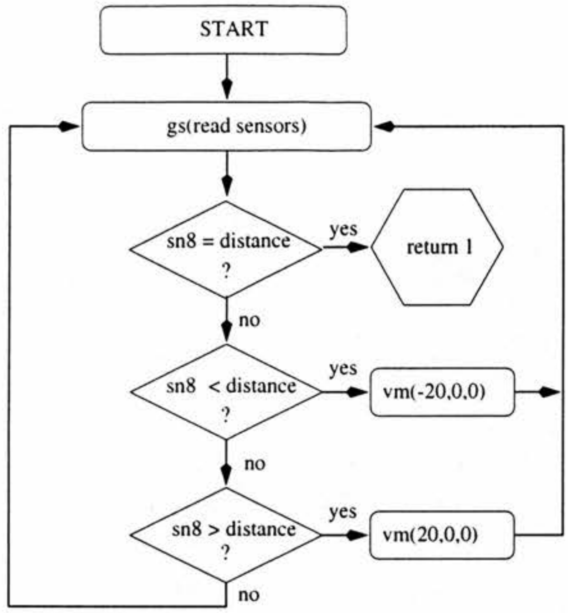

Figure A.16: orient-rear(distance)[1]

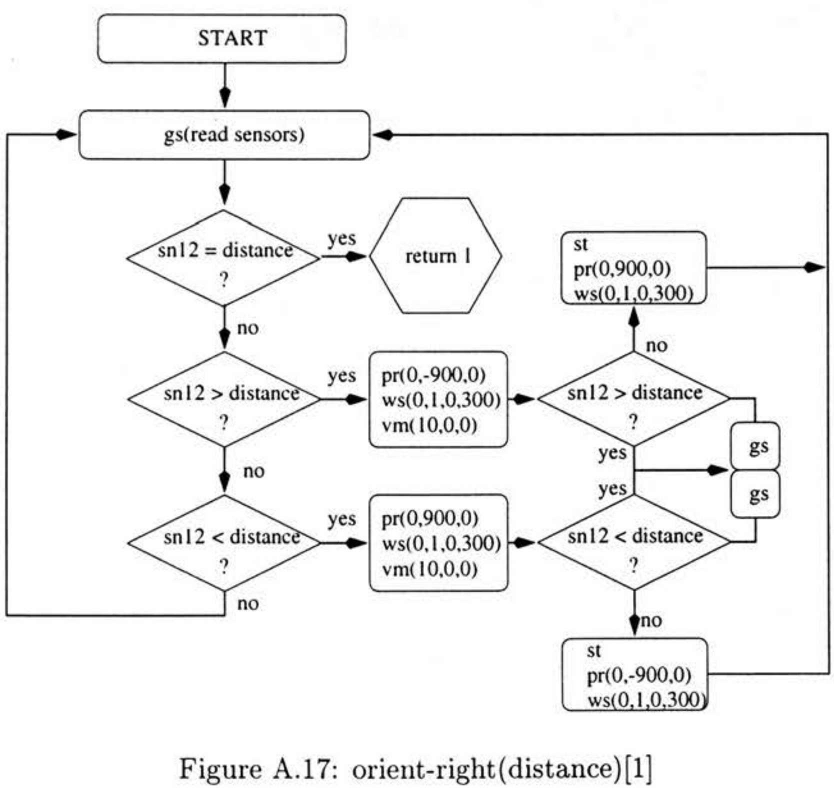

Figure A.17: orient-right(distance)[1]

inches. This function is designed to run as a *child* process and implements (see Figure A.19) an eternal avoid loop unless killed.

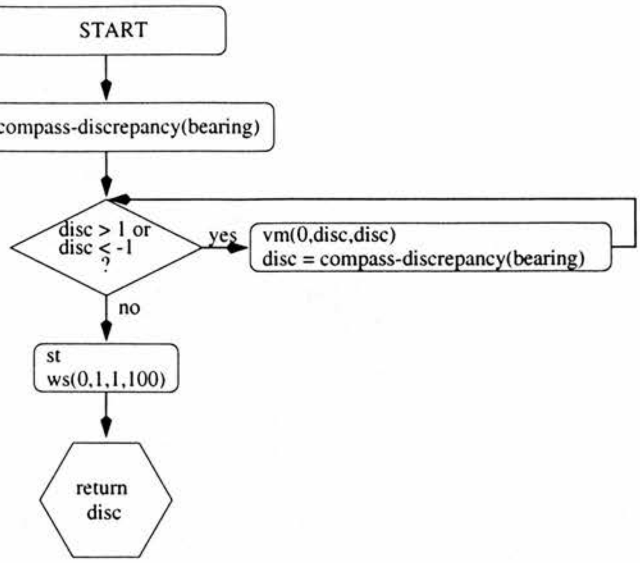

Figure A.18: turn-to-bearing(bearing)[disc]

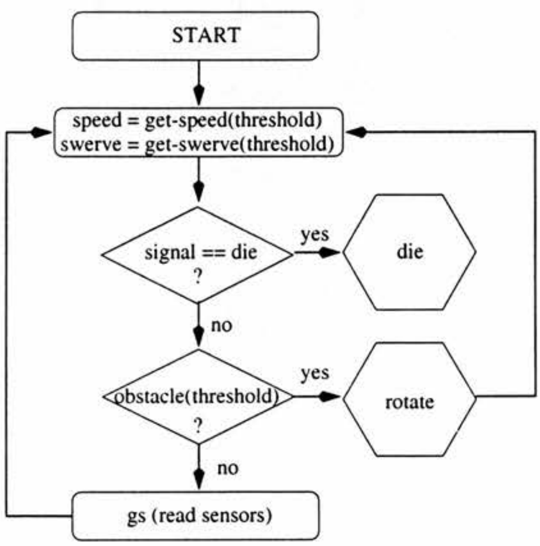

Figure A.19: avoid(threshold)[]

A.3.2 `rotate()[]`

`rotate()[]` Rotates to left or right whilst an obstacle is detected (see Figure A.20). The function exits when sensor readings indicate that no object is within a range of 15 inches from an arc of the front three sonar sensors sn_{15}, sn_0, sn_1.

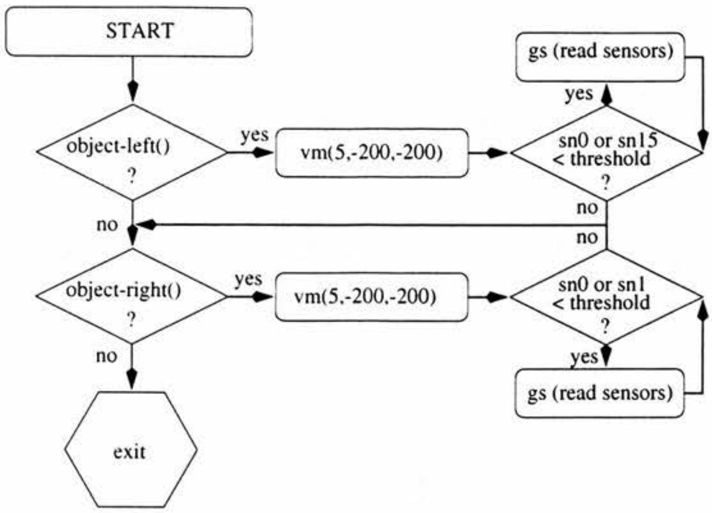

Figure A.20: rotate()[]

A.4 Wall following

A.4.1 `follow-left(distance)[]`

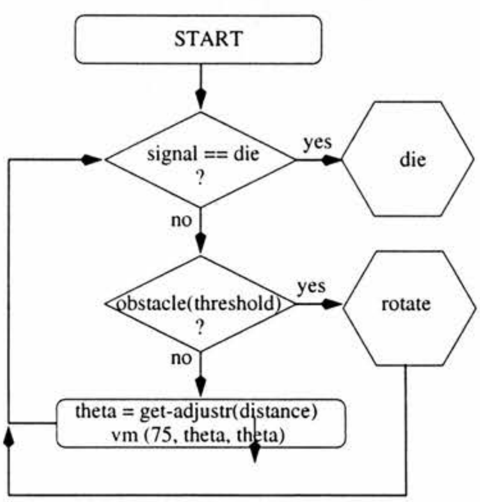

Figure A.21: follow-left(distance)[]

`follow-left(distance)[]` A wall-following loop for use as a *child* process (see Figure A.21). Allows the nomad to shadow obstacles at a range of **`distance`** (an optional parameter, the default is 22 inches).

A.4.2 `follow-right(distance)[]`

`follow-right(distance)[]` A wall-following loop for use as a *child* process (see Fig-

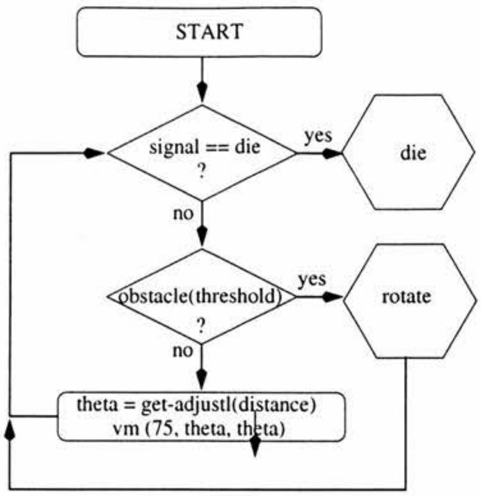

Figure A.22: follow-right(distance)[]

ure A.22). Allows the nomad to shadow obstacles at a range of **distance** (an optional parameter, the default is 22 inches).

A.4.3 `get-adjustl(distance)`[θ_{steer}]

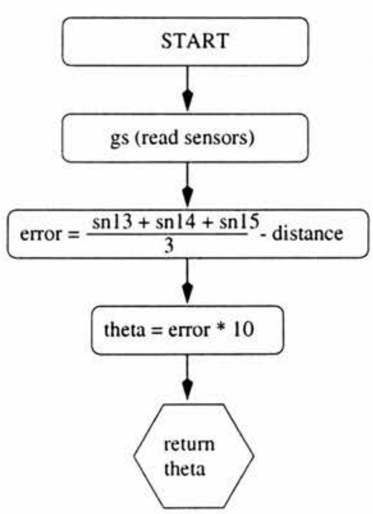

Figure A.23: get-adjustl(distance)[theta]

`get-adjustl(distance)`[θ_{steer}] Returns the steering angle for following a right-side object. `distance` is an optional parameter, the default is 22 inches (see Figure A.23). The external state variables `max-adjustl` and `min-adjustl` are used in the case of θ_{steer} falling outside the range `min-adjustl` $< \theta_{steer} <$ `max-adjustl`.

A.4.4 get-adjustr(distance)[θ_{steer}]

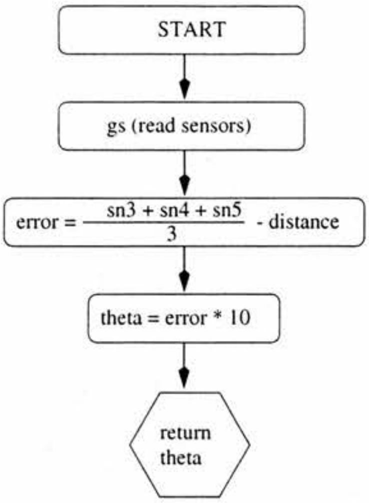

Figure A.24: get-adjustr(distance)[theta]

get-adjustr(distance)[θ_{steer}] Returns the steering angle for following a left-side object (see Figure A.24). **distance** is an optional parameter, the default is 22 inches. The external state variables **max-adjustr** and **min-adjustr** are used in the case of θ_{steer} falling outside the range **min-adjustr** $< \theta_{steer} <$ **max-adjustr**.

A.5 Navigation

A.5.1 go(x,y,range)[0,1]

go(x,y,range)[0,1] A navigation routine (see Figure A.25) designed for use independently or as a *child* process. The algorithm moves the nomad directly from its current position to within the specified **range** of (x,y).

A.5.2 go-monitor(x,y,timeout,range,frequency)[0,1]

go-monitor(x,y,timeout,range,frequency)[0,1] This function (see Figure A.26) serves to constrain **go** when invoked as a *parent* process. It takes five compulsory arguments: x and y are the goal coordinates; **timeout** is the function's timeout period in seconds; **range** within which recovery is possible; **frequency** is that of checking progress towards the goal position. It returns 1 if the child process exits successfully before **timeout** has expired, or if upon timeout the nomad is within recovery distance (**range**), 0 is returned in any other case.

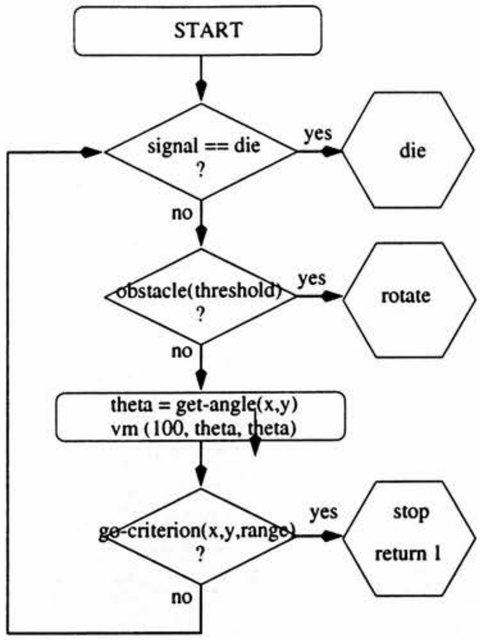

Figure A.25: go(x,y,range)[0,1]

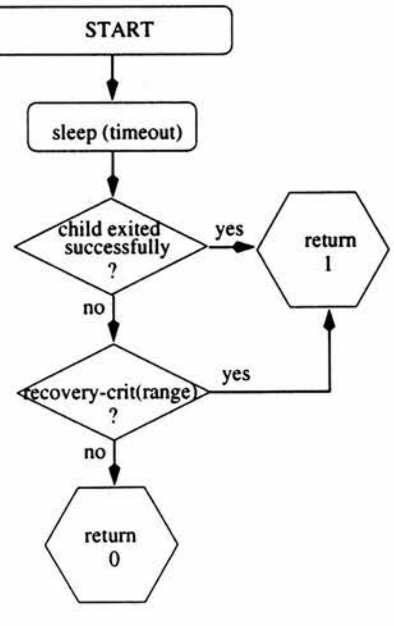

Figure A.26: go-monitor(x,y,timeout,range,frequency)[0,1]