Chapter 1

1. The Behaviour of Animals and Animats

Man has long sought to understand what constitutes life, and to understand the nature of living things. The new discipline of Artificial Life (Langton, 1989; Levy, 1992; Brooks and Maes, 1994) acts as a focus for research into a diverse set of topics relating to the modelling and understanding of life and the properties of living things. Artificial Life concerns itself with many aspects of those organisms we recognise as living entities. These aspects include evolution, morphology, swarming behaviours, behavioural models and learning, even the nature of life itself. The idea that “living” entities might yet be constructed artificially remains highly speculative and contentious, only in part due to the difficulties in agreeing a satisfactory definition of what does and what does not constitute the necessary properties of being alive. There is more general agreement that simulation can greatly add to our overall understanding of the nature of the structure and behaviour of living things. This work concerns itself with the behavioural properties of the individual. It will therefore touch upon the broader issues addressed by Artificial Life only in passing.

One question has engaged the minds of psychologists and those interested in a greater understanding of animal behaviour for decades. Is the behaviour of animals inherently driven by the current state of the world as perceived through the senses, or is it directed by goals, internally generated needs or requirements of the organism? Huge amounts of evidence supporting these two disparate viewpoints has been accumulated. It is an argument that is far from being resolved and one that has spilled over into the newer domains of Computer Science and Artificial Intelligence, where another generation of scientists is pondering the question and proposing new models of behaviour in an attempt to resolve the issue. The question was the subject of a meeting that invited this new generation of researchers to declare and defend their position - “models or behaviours” (Aylett, 1994). Paralleling this question is that of how learning is to be achieved in either of these possible situations. These problems have recently found renewed expression in an area of study broadly categorised as the “simulation of adaptive behaviour” (Meyer and Wilson, 1991; Meyer, Roitblat and Wilson, 1993; Cliff, Husbands, Meyer and Wilson, 1994; Maes, Mataric, Meyer, Pollack and Wilson, 1996). The debate is set to continue.

1.1. Three Components of Natural Intelligence

For the purposes of this thesis behaviour will be divided into three broad categories: (1) capabilities inherent to the individual from the moment it comes into being; (2) capabilities it may acquire as a result of interaction with its environment; and (3) capabilities acquired by processing or reformulating information or capabilities derived in any of the three categories. The first category will be referred to as “innate capabilities”, the second as “learned capabilities”, and the third will encompass a range of abilities broadly categorised as “problem solving”, and “inductive” and “deductive inference”. Some, possibly all, elements of the processes supporting categories (2) and (3) may also be an innate process inherent to the individual. Information from any category can potentially be utilised and exploited by any of the categories. Therefore the element of self and cross-reference of the categories is intentional. The “intelligence” of the individual will be based on some combination of these three basic activities (undoubtedly supported by many other activities of the individual and its structure). Intelligence will not be defined here by any specific ability, but rather by the degree or extent to which the individual can react and adapt to the circumstances that impinge upon it. One prevailing view holds that an individual can be considered intelligent solely on the basis of capabilities defined in the first category. Others argue that any useful degree of intelligence can only be displayed in individuals with significant capabilities in categories (2) and (3). This work will concentrate on the nature of intelligence as it arises from categories (1) and (2). This chapter and chapter two will consider the approaches adopted by others. Perhaps interestingly, these capabilities may arise either as a result of an evolutionary or a creational process, with little impact on the observable performance of the individual under study.

The term *animat* (Wilson, 1985, 1991) will be used throughout this work to indicate an artificial or simulated model of an animal. The term will also occasionally be used to denote properties shared by these simulated and natural animals. Specifically the term animat is used in preference to *agent*, which is used by various authors to refer variously to either an individual, or to component parts of an individual. The term animat is not intended to represent any specific organism or species type. The term *ethogram* will be used to represent a description, in operational form, of the behavioural capabilities of the animat in each of the three categories at the moment it becomes a free standing individual. The term "ethogram", after ethology[1], is apparently due to Kirsh (1991, p. 167).

1.2. Reactive Models of Intelligence

This section considers some of the issues relating to the first category of intelligent behaviour, variously named *behaviour based* (Maes, 1993), *reactive*, or *situated agent* models of behaviour (Agre, 1995). Brooks' (1991a) view of *intelligence without reason* and his (Brooks, 1991b) *intelligence without representation* arguments follow in a long tradition of stimulus-response (S-R) *behaviourism*. All argue that the majority of observed and apparently intelligent behaviour may be ascribed to innate, pre-programmed, processes available to the individual. This viewpoint is not without its critics, Kirsh (1991) for instance. Category (1), innate, capabilities of the individual derived from an evolutionary process are shared by all members of the same species (allowing for some variation between individuals). Individuals derived by a creational process acquire innate intelligence from their constructor. Similarly, we may be impressed by the advice from an expert system and yet be aware that the intelligence displayed is still derived from the knowledge of a human expert. In both cases the intelligence seems diluted. To a certain extent capabilities derived in this first category may be regarded as "intelligence without intelligence".

Innate intelligence is not, however, defined by degree. The behavioural repertoire of an insect may be completely mapped, and its ability or inability to react to any situation comprehensively modelled. At a distant end of this scale Pinker (1994)

---

[1] (OED): ethology n. Science of character formation; science of animal behaviour

argues that human language ability, for all its complexity, is primarily innate. He cites much evidence that all undamaged humans develop language abilities to a largely uniform level of complexity by simply interacting with others, essentially regardless of (and possibly in spite of) any form of education or teaching. Specifics of vocabulary and grammar are environmentally determined, but vocabulary and grammar develop in all undamaged individuals as a matter of course during their infancy. Notwithstanding differences in their vocal tracts it is clear that, while non-human primates may be taught a limited vocabulary of symbols, attempts to teach or activate any significant tendency to structured grammar remain largely unsuccessful (Premack, 1976). Where significant progress has been reported this has lead to suggestions of observer bias.

The innate behavioural repertoire of many species has been extensively studied. Where this is done primarily by observation of the animal in its natural surroundings, the term *ethology* is often used. An alternative approach, adopted by behavioural scientists, places the subject animal in controlled experimental conditions to investigate the subject's reactions. Innate behaviour patterns are reasonably investigated by the former procedure, but aspects of learning and problem solving are often better researched by the latter method. This appears in part due to the wide range of innate activities a subject may perform, masking or hiding specific learning phenomenon under investigation.

1.3. Action Selection Mechanisms

Action Selection Mechanisms (ASM) attempt to provide a model to understand how behaviour is generated in response to the current requirements of the animal. These are specific implementations of category (1) notion of intelligence, that of unlearned or innate behaviour. They do so in a manner intended to illuminate the properties observed of living creatures. The systems discussed here tend toward the modelling of natural systems, but are not drawn exclusively from those that do so. For largely historical reasons these models concentrate on a variety of non-primate vertebrate species, including small mammals, birds and fish, whose behaviour may be closely observed and recorded. Tyrrell (1993, Ch. 8) provides a useful summary of a variety of action selection mechanisms drawn from both natural and artificial examples. Despite the huge body of observational evidence from the discipline of ethology and the subsequent introduction of computers allowing detailed simulation and testing of the various theories, there is still much controversy as to which of the many possible architectures represents the most appropriate description.

Tinbergen (1951, Ch. 5) devised a model for the organisation of behaviour based on observations by himself and others of a variety of species, including the digger wasp, the three-spined stickleback and the turkey. Tinbergen's model is a hierarchic control model of action selection. The creature is embodied with several central "instincts". Figure 1-1 models that of the reproductive instinct of the *three-spined stickleback*. Each central instinctive behaviour is inherently part of the creature, but it is not always manifest. Reproductive behaviour in the stickleback is a complex set of activities spread over a period of many weeks during the breeding season. Once initiated, say by the onset of warmer weather or lengthening hours of daylight in the spring, second level behaviours become active. In this model such behaviours are normally inhibited by a blocking mechanism. When circumstances appropriate to the conduct of some aspect of the innate behaviour are sensed an *innate releasing mechanism* (IRM) removes the block, so enabling behaviours at a lower level in the hierarchy. These sub-ordinate behaviours may then also be released by their IRMs, shown in figure 1-1 as grey coloured areas, when the conditions appropriate for their use are encountered. Lorenz had earlier proposed a simple hydromechanical analogy to illustrate the operation of the IRM (Lorenz, 1950).

Tinbergen distinguishes between *appetitive actions*, those which establish the conditions needed to continue or complete a sequence of behaviours and *consummatory actions*, which appear to "satisfy" the motivation for the action sequence and so complete it. Level 3 subordinate behaviours represent these appetitive and consummatory behaviours, and are observed and recorded by the ethologist. These behaviour units are considered to be *fixed action patterns* (FAP), groups of low level actions that may be initiated to complete some aspect of the overall instinct. Level 3 behaviour units may themselves be further sub-divided into the co-ordination of, for example, fin (level 4), and fin ray (level 5) movements, muscle activations (level 6) and so on.

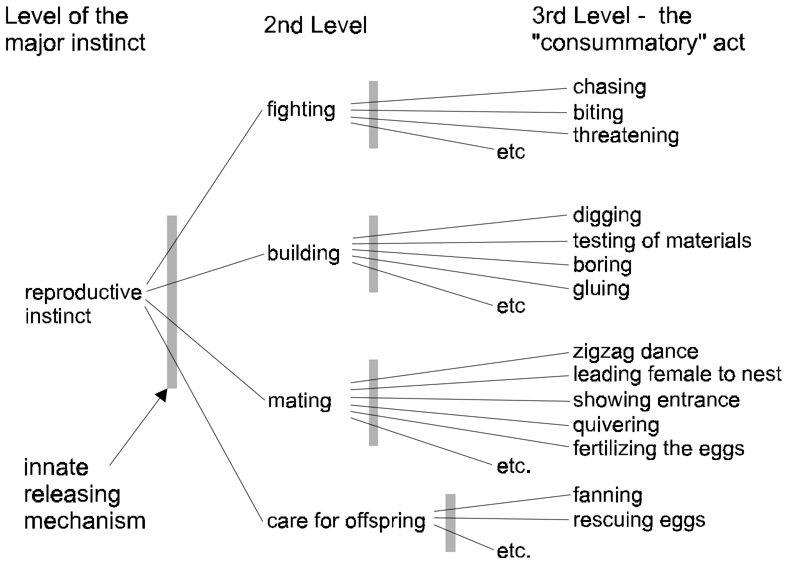

Figure 1-1: Tinbergen’s Principle of Hierarchical Organisation

adapted from Tinbergen (1951), p. 104 & p. 124

Baerends (1976) presents a hierarchical model to account for the incubation behaviour of the herring gull. This model adds inhibition between superimposed control centres (level 2 behaviours), in which active centres suppress the effects of others. Friedman (1967) prepared a computer model and simulation of the concepts of innate behaviour. He retained the notion of an innate releaser mechanism, but argued that viewing level 3 behaviours as fixed action patterns was too simplistic. To counter this apparent oversimplification Friedman introduced *behavior units*, behaviour patterns controlled and maintained by feedback loops at level 3. His system was tested with a simulated artificial animal, *ADROIT*. Travers (1989) presents a computer simulation of the stickleback’s innate reproductive behaviour; Hallam, Hallam and Halperin (1994) a simulation of aspects of behaviour in the *Siamese fighting fish*.

Rodney Brooks has described the *subsumption architecture* (Brooks, 1986). While not strictly an ethologically inspired model of behaviour it has proved influential in the design of subsequent reactive and behavioural models. Figure 1-2 illustrates some of the main features of the subsumption architecture. In a conventional model of robot task behaviour, Brooks argues, behaviour is decomposed into functional modules such as “perception”, “modelling”, “planning”, and so on. Each module will be involved in the completion of many different task types. In a subsumption architecture the robot control system is decomposed into individual task-achieving modules, a “level of competence”. Lower levels being responsible for simpler or more primitive activities. Each level is nevertheless responsible for a complete behaviour, having access to the sensory information it requires and the ability to send instructions to actuators. Examples of such behaviours include “obstacle avoidance” (level 0), “wandering behaviour” (level 1), “explorational and map building behaviour” (level 2), up to, say, the ability to reason about objects in the world and create plans.

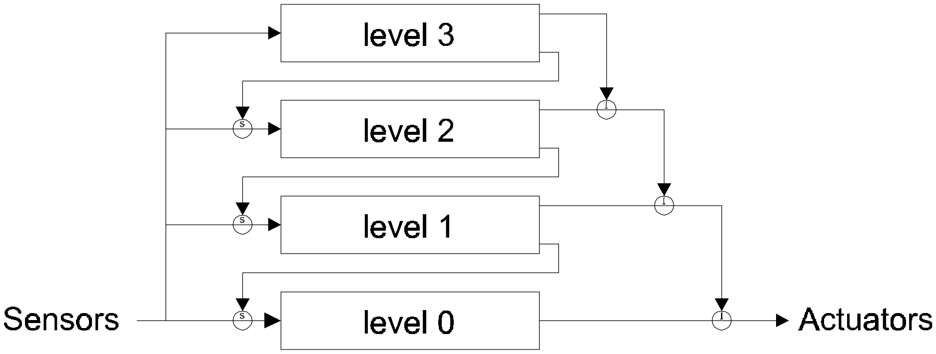

Figure 1-2: Brooks’ Subsumption Architecture

adapted from Brooks (1986), p. 17 & p.18

In Brooks’ model each level is created as a finite state machine. Every higher layer may subsume the behaviour of a lower layer, by modifying its input information (shown as a circled “S” on the input side of each layer in figure 1-2) and therefore adapt the lower level behaviour to its requirements. Alternatively the higher level may inhibit the output of lower layers to take control of the output behaviour (shown as a circled “I” on the output side of each layer in the figure). Brooks (1990) describes the *behavior language*, which allows behaviours defined in terms of the subsumption architecture to be complied into the native code for a variety of processor types including the Motorola 68000 and 68HC11, Hitachi 6301 and to Common Lisp.

Tyrrell (1993) argues that actions are not best selected on an all or nothing basis. Rather each module should contribute “evidence” for one or more of the possible actions available to the animat, with a “winner-take-all” strategy in place to select the final outcome to be sent to the actuators. His model is based on one devised by Rosenblatt and Payton to automatically control and navigate a mobile vehicle (Rosenblatt and Payton, 1989; Payton, Rosenblatt and Keirsey, 1990). Rosenblatt and Payton’s model overcame the potential loss of data in the subsumption architecture by allowing each behaviour module to feed (positive or negative) activations via weighted links to summation points for each action type.

Brooks’ subsumption architecture proposal is reminiscent of Paul Maclean’s *triune brain hypothesis* (Albus, 1981, p. 184). Each of three layers represents a stage in the evolution of the modern mammalian brain. All the layers have access to sensory mechanisms and motor outputs and are organised as a control hierarchy. The inner layer, layer one, is the primitive reptilian brain, equipped with reflexive and instinctive behaviours. Built over this primitive layer is the “old mammalian” brain, providing additional attributes, elements of planning, predictive abilities and some elements of memory. In turn the third layer, or “new mammalian” brain provides another set of capabilities including the sophisticated manipulation of arbitrary symbols and concepts, language and a distinct model of self. As in the subsumption architecture, each layer has access to information available to a lower layer but may also intercept and override the output of a lower layer.

Maes describes a bottom-up mechanism for action selection (Maes, 1989, 1991, 1993), which, while being primarily a computer based animat controller, addresses the problems of action selection from a broadly ethological viewpoint. Figure 1-3 illustrates the main points of her action selection model. The animat has a number of innate motivations (or, synonymously, goals), which are in turn connected to consummatory activities. Consummatory activities will, if performed, lead to a reduction or satisfaction of the attached motivation; eating assuages hunger, drinking slakes thirst and so on. Consummatory activities may in turn be linked to appetitive activities, ones that prepare the animat to complete the behaviour. Some appetitive activities lead directly to a consummatory activity; others are linked into chains of activities that lead the animat closer to the motivating goal. Thus eating food is preferable to moving towards food that can be seen, which in turn is preferable to moving to a location where food is remembered to be located, to having to explore for food.

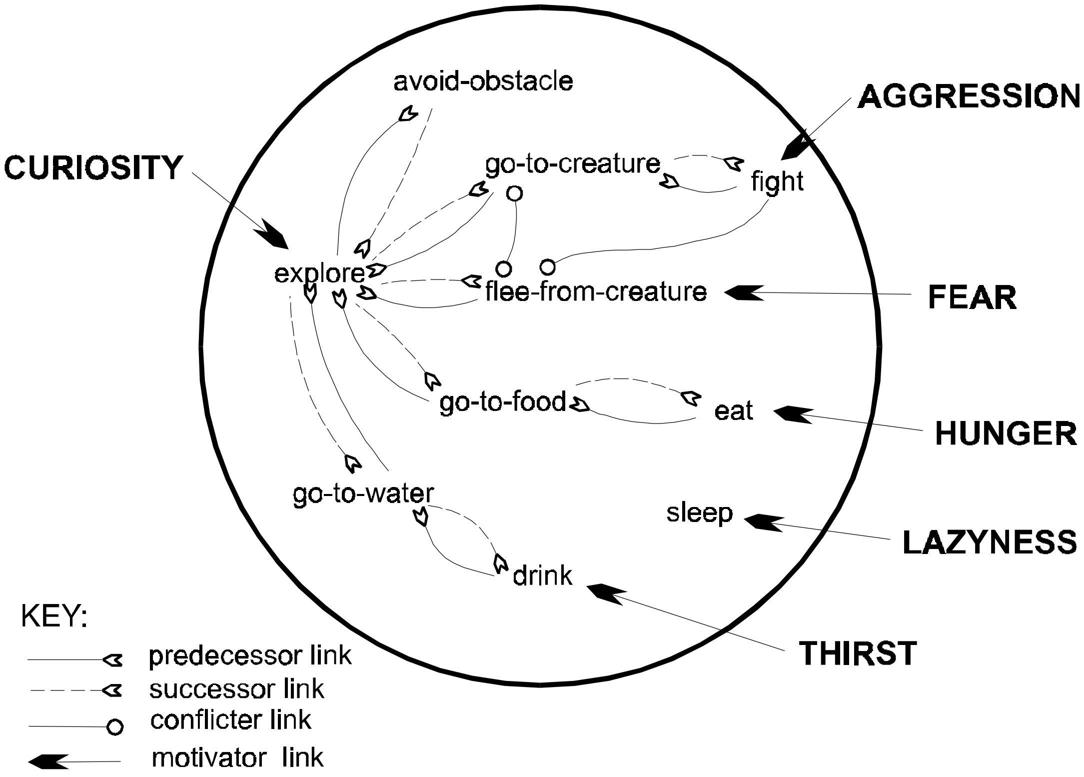

Figure 1-3: Maes' Action Selection Architecture

adapted from Maes (1991), p. 240 & p. 242

Activities are linked by a network of *predecessor links* ("——<"), a list of pre-conditions necessary to initiate an activity and by *successor links* ("- - -<"), add-list conditions arising as a consequence of performing the activity. Activities may also inhibit other activities with a *conflictor link* ("——o"). At any time each of the motivations will be characterised by a level of activation, a degree of "hunger", "thirst", "fear", etc. Motivation activations spread throughout the network of activities through the predecessor links, the activation level being relative to the strength of the motivation and to the number and type of links between motivation and activity. At the same time appetitive and consummatory activities attain a level of activation based on the degree to which their preconditions are met, either by activation via their predecessor links, or directly from sensory conditions associated with the activity. *Activation* spreads in two directions, along both predecessor and activator links, inhibition via conflictor links. At any time, then, the animat may select an action based on both its current needs and the prevailing environmental circumstances in which it finds itself. Tyrrell (1994) implemented and tested Maes' action selection mechanism with a wide range of parameters and concluded that there were some significant drawbacks to the mechanism she had described.

Action selection mechanisms only address the first category of intelligence as described previously. They are an important part of the process, but insufficient to account for the range of phenomena observed. The next sections concentrate on the second category, that of learning and learned behaviour.

1.4. Arriving at a Definition of Learning

It has not proved easy to generate an all embracing definition of exactly what does, and what does not, constitute the process of learning. Learning is by no means synonymous with change; it is clearly a form of change, but one that makes "*useful changes in the workings of our minds*" (Minsky, 1985, p. 120). This definition is imprecise and incomplete. Simon (1983) extends the definition to "*learning denotes changes in the system that are adaptive in the sense they enable the system to do the same task or tasks drawn from the same population more efficiently the next time.*" Razran (1971, p17) suggests that a "*commonsense view of learning*" would be "*profit through experience,*" but immediately qualifies this to "*more or less permanent central modifications of a reaction or reactions through reacting and interacting of reacting.*" He then further excludes transient changes such as fatigue and sensory or effector adaptation. Razran and Simon have both identified a clear property of learning systems - they improve what they do by doing what they do.

Bower and Hilgard's (1981, p. 11) definition of learning develops the theme:

"*Learning refers to the change in a subject's behavior or behavior potential to a given situation brought about by the subject' s repeated experiences in that situation, provided that the behavior change cannot be explained on the basis of the subject's native response tendencies, maturation, or temporary states (such as fatigue, drunkenness, drives, and so on).*"

This definition amplifies the notion of change precipitated by experience and made manifest in behaviour. Thus category (2) intelligence is distinguished from category (3) intelligence in that the change is mediated by the receipt of external information, rather than a reprocessing of internally held knowledge. The distinction becomes increasingly blurred as previously learned information is itself reformulated. This last definition also introduces an element of permanence, or at least semi-permanent change, which does not readily revert to the previous condition without further experience within the environment. It is clear from these definitions that while learning is a change in behaviour, not all changes in behaviour can be regarded as learning. Chapter two reviews possible behavioural mechanisms that can be described as learning, the next sections consider some forms of behavioural change that are excluded by the definitions.

1.4.1. What is Not Learning

Bower and Hilgard's definition also excludes a number of other sources of change that should not be classified as learning. These sources of temporary change, such as fatigue or the influence of drugs, are essentially reversible and the animal will revert to its original behaviour once the effects of the influence abate[2]. Similarly, the effects of *habituation* and *sensitisation* are normally excluded from definitions of learning. There are many situations in which an organism will come to react less frequently or with less vigour to a particular sensation apparently only due to the frequency of presentation of that stimulus. The organism "habituates" with respect to the stimulus. An organism may also react more vigorously to a stimulus that has been withheld for an abnormal period. The organism is "sensitised" with respect to the stimulus. Both conditions are transitory and reaction reverts to normal levels once the stimulus regime is stabilised.

Maturation, on the other hand, does represent a permanent change, but one that also falls outside the definition of learning. Maturation represents behavioural changes in the organism that take place essentially independently of the individual

---

[2] Which is not to say that the organism will not modify its behaviour as a consequence of these influences. A drinker might subsequently imbibe more due to the pleasing effects of inebriation, or less due to the consequences of a hang-over. In either case the intoxicating effects of the alcohol ingested may be considered essentially transient.

organism's experience in its environment. Such behavioural changes mirror physical changes due to growth, and may be linked to or co-ordinated with the development of physical attributes. As an example of the maturation process Altman and Sudarshan (1975) investigated the development of reactions to different environmental situations in new-born rat pups, showing the appearance of successively more complex behaviour patterns during the first weeks of life. These changes are apparently pre-programmed to occur in the organism, in much the same manner that innate tendencies appear as pre-programmed reactions to specific stimuli.

Imprinting may be considered as a special case of maturation, in which the individual is pre-programmed to incorporate an external stimulus as releaser or trigger for some other pre-programmed behaviours. Only the stimulus adopted varies between individuals of the species, the mechanism to adopt some stimulus (often within recognisable limits), and the reactions it will subsequently elicit appear to be pre-programmed. Imprinting was first recognised by the ethologist *Konrad Lorenz* (1903-1989). He noticed that graylag goose chicks, which normally follow their mothers, would follow a human in preference to their mother if exposed to a human individual at a critical stage in their development. Imprinting is characterised by a typically rather narrow *sensitive period*, during which the effect develops easily. Ducklings (Hess, 1959) are most sensitive to the effect at between 13 to 16 hours after hatching. Attempts to imprint before 5 hours or after 21 hours from hatching invariably fail. The imprinting phenomenon has been widely researched and has been demonstrated in a variety of avian and mammalian species (Dewsbury, 1978, pp. 140-153).

1.5. A Caveat

This work strives to present a "biologically inspired" model of an animat controller; it is not intended as a specific model of any particular animal or species. Such models have been prepared, and often shed further light on the nature of the creature being emulated (Arbib and Cobas, 1991; Arbib and Lee, 1993; Hartley, 1993; Mura and Franceschini, 1994; and Webb, 1994, for instance). Cliff (1991) has promoted the term *computational neuroethology* for this type of study (Beer and Chiel, 1991), Sejnowski, Koch and Churchland (1988) the term *computational* *neuroscience*. Roitblat, Moore, Nachtigall and Penner (1991) propose *biomimetics* in relation to their neural network model of echolocation in dolphins.

In designing the animat controller some of the essentially engineering solutions that arise are resolved on the basis of *biological plausibility*. By adopting evidence drawn from many different species, under many different experimental regimes, general principles may be identified and integrated into a whole. However, it is unreasonable to assume that capabilities are evenly distributed across the animal kingdom. There is diversity at every point and at every level, so that a generalised model cannot be expected to account for detailed reactions in specific individuals.

1.6. Thesis Outline

This chapter has introduced the idea that animal intelligence is composed of three component parts, (1) innate behaviour, (2) learned behaviour and (3) behaviour directed towards inferring and deducing new knowledge from existing knowledge. As well as defining some terms, several models of innate behaviour were described and what does and does not constitute learning was also considered.

Chapter Two develops the theme of learning, concentrating on learning in reactive systems. A review of learning from a historical perspective introduces many important concepts and illustrates the spread of the problem being addressed. A review of recent and current research into computer models concentrates on work in reinforcement and *Q*-learning methods, classifier systems and artificial neural networks. The chapter also considers the evidence for a cognitive or goal driven view of learning and behaviour in animals. Existing models of intermediate level (sensory-motor) cognition are reviewed.

Chapter Three considers the role of hypothesis generation and verification by experiment at a behavioural level, consistent with reported observations of animal behaviour. A comprehensive set of postulates for a new *Dynamic Expectancy Model* is developed which combines the apparently disparate threads of reactive behaviour, perception and action, goal setting and pursuit, and learning.

Chapter Four develops a computer simulation algorithm (*SRS/E*) from the Dynamic Expectancy Model presented in chapter three. This chapter describes the data structures and processes required to implement the Dynamic Expectancy Model.

Chapter Five describes an experimental environment attached to the SRS/E program implementation and describes the facilities available to an investigator using the program.

Chapter Six reports a series of experiments with the SRS/E algorithm. These experiments are constructed to allow direct comparison with other published reinforcement learning algorithms, and to several well-established procedures from the behavioural sciences, which are adapted for use with the SRS/E program.

Chapter Seven describes some possible extensions to the Dynamic Expectancy Model to enhance the SRS/E algorithm.

Chapter Eight concludes by reviewing the relationship of reinforcement learning to cognitive structures and proposing *Expectation Based Learning* (*XBL*) as a fruitful line of research investigation for the future.

Appendix One gives a complete description of the execution cycle for the SRS/E algorithm, described in detail in chapter four.

A bibliography of references is attached, as is an index of topics and author citations by page number.