Designing Agents to Understand Infants

Dean Petters

A thesis submitted to the
Faculty of Science
of the
University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

School of Computer Science,
The University of Birmingham,
Birmingham, B15 2TT.

October 2006

Abstract

This thesis attempts to understand infant behaviour by designing autonomous software agents to reproduce those behaviours in simulation. The infant behaviours that have been investigated are related to the phenomenon of infant attachment. Empirical studies of infant behaviour are abstracted and the function of the behaviours are assessed from an evolutionary perspective. The behaviours are then reformed as scenarios against which simulations can be evaluated. These studies include naturalistic observation of infants at home and exploring a park, and undergoing the Strange Situation Experiment.

A number of information processing architectures have been constructed that reproduce the infant behaviours described in scenarios. These vary in complexity from a reactive architecture with no capacity to learn, to reactive architectures that can learn by reinforcement, and deliberative architectures that can reason by forming simple plans. Computational experiments undertaken with interacting infant and carer agents show the presence of interesting dynamic properties, such as positive feedback loops. These feedback loops may provide an explanation for the empirical finding that patterns of infant attachment response cluster into three categories. This thesis demonstrates how the methodology that it is uses in investigating attachment behaviour in infancy may be extended to many other infant and developmental behaviours.

Acknowledgements

The opportunity to do a PhD in Cognitive Science has been a great privilege. I would like to thank the Computer Science Department for having accepted and funded me. I'd also like to thank Aaron Sloman for all aspects of his supervision. In particular for making numerous comments that have improved my written work, and providing me with the many tutorials that have helped shape my understanding of infant and adult minds.

Other people to thank include my other thesis group members - Ron Chrisley, Peter Hancox, Sheila Glasbey and Gill Harris - who both kept me on my toes and provided valuable advice; the computer science support team and other staff; and all the people who have read and made comments on my work at different stages.

I would particularly like to mention: Mark Roberts; John Woodward; Dave Gurnell; Gavin Brown; Dave Parker; Adrian Hartley; Aiden Harding; Simon Harding; Luc Beaudoin; Ian Wright; Russell Beale; John Bullinaria; Ian Styles; Jeremy Wyatt; Peter Coxhead; as well as my wife, Cath, and my mum and dad, sisters and brother. I very much appreciate the assistance, advice, encouragement and comments the aforementioned have offered regarding the thesis or previous papers.

Most of all I'd like to thank Cath, and Lauren and Beth, for putting up with me doing my PhD, and making my life in the last few years so much more than just studying and research.

Contents

List of Figures

Glossary and abbreviations

Analogical representation - a form of non-fregean representation where the *"system properties of and relations between parts of the representing configuration represent properties and relations of parts in a complex represented configuration, so that the structure of the representation gives information about the structure of what is represented"* (Sloman, 1971)

EEA - (Environment of Evolutionary Adaptedness), the environment within which the human species evolved, (and therefore the environment to which an infant's information processing architecture is adapted).

Fregean representation - Most logical and mathematical formulae and much of the syntax and semantics of natural language are Fregean. What these representations have in common is that they can be analysed in terms of the application of functions to arguments (Sloman, 1975). The units of 'sentences' in a Fregean representation are composed of discrete symbols that can often be combined productively according to a compositional semantics.

GS - (goal-switching), an architecture defined in section 2.3

GLA - (goal-learning-from-anxiety), an architecture defined in section 3.3

GLS - (goal-learning-from-socialisation), an architecture defined in section 3.3

HAR - (hybrid-action-reasoning), an architecture defined in section 4.3.3

RAL - (reactive-action-learning), an architecture defined in section 4.3.2

Temperament - This thesis will be using the term 'temperament' to describe traits that possess a high degree of heritability, such as activity level or irritability, and not its sense in more general application where it can refer to traits that possess non-innate aspects, such as general personality or character.

Chapter 1

Introduction

1.1 The nature of the problem

This thesis seeks to better understand the kinds of process which occur within a one year old human infant's mind/brain as it pursues everyday activities such as exploring a park, meeting a stranger or being reunited with its carer after a short absence. The aim is to explain these behaviours by designing and building software programs that reproduce the behaviours, at an abstract level, in a simulated virtual environment. Autonomous agents that exist in a two dimensional virtual world have been designed to reproduce the behaviours under study. The simulated agents are a simplified abstract representation of real infants, implemented at a high, goal-oriented level. They have been designed to capture only the key aspects of an infant's interaction with other agents and objects in its environment. Whether infants use particular information processing structures or mechanisms is an important question. Of additional interest is the organisation and interrelation of any structures and mechanisms within an overall information processing architecture. Hence this work can be summarised as an attempt to explain infant attachment behaviour in terms of possible structures, mechanisms and classes of information processing architecture.

The period of infancy under investigation is of interest because infants of 9 to 12 months of age show the beginnings of distinctly human forms of information processing (Tomasello, 1999a). This period in human infancy sees the *"appearance of unmistakably intentional, goal directed behaviour."*(Flavell *et al.*, 2002). Infants for the first time demonstrate the ability to separate means and ends (Willatts, 1999). In addition, infants engage in new forms of joint visual attention (Butterworth, 2004), such as joint engagement between carer and infant directed towards other objects (Tomasello, 1999b). Before this period human infants do not exhibit evidence of forms of information processing that would clearly distinguish them as more advanced than other species. During this period their ability to engage in means-end reasoning, for example whilst using tools, is comparable with more advanced examples of tool use observed in non-human animals[1]. Much after this period, as infants approach eighteen months of age, many behaviours, in particular the full use of language, are markedly different from anything produced by non-human animals. An inviting question is thus: In this transitional period, are attachment behaviours produced by a basic information processing architecture, that has more in common with other animals, or a more specialised architecture that has distinctly human features?

The thesis attempts to explain a number of related behaviours, the outlines of which have been organised into a number of scenarios which include: the PARK, WARY, COY and CONTACT (chapter two), SECURE-BASE (chapter three), and REUNION (chapter four) scenarios. The four scenarios in chapter two present relatively simple patterns of behaviour. Some sketches of behaviour from these four scenarios are:

> **Exploring a park (A sketch from the PARK scenario)** - A carer and one year old infant enter a park and move to a spot away from other people. The carer sits down and the infant makes exploratory forays into areas surrounding the carer. Periodically the infant stops exploring, turns and 'checks in' with the carer before returning to exploration. Sometimes the infant 'checks in' by making eye contact from a distance, at other times the infant returns and makes physical contact.
>
> **Response to an unfamiliar environment (A sketch from the WARY scenario)** - A carer and one year old infant enter a park that the infant has not visited before. The carer sits down and the infant makes exploratory forays. However the infant stays closer to the carer

---

[1] A detailed review of animal cognition is beyond the scope of this thesis, and whether there is true means-end reasoning in non-human species is very much an open research question. For some species the results are clear cut, for example dogs fail in a means-end reasoning task which is similar to a task in which human nine month old infants succeed (Osthaus *et al.*, 2005; Willatts, 1999). The results for chimpanzees are less clear. A problem in interpreting any behaviour as resulting from true means-end reasoning is that the same behaviour may result from trial and error learning, and association of particular tools with particular goals, without a full understanding of the causal relations between objects. Finding examples where animals use tools to use other tools provides stronger evidence of true means-end reasoning. In reviewing animals' ability to use two means (means-means-end) to solve a problem Santos *et al.* (2005) find that *"initial data suggest that non-human tool users have, at the very best, a rather limited ability to use two means in combination to achieve a goal and that means-means-end proficiency might be highly variable across individuals."*

> than it would in a familiar park, and also ‘checks in’ and makes more physical contact until it becomes familiar with the new environment.
>
> **Meeting a stranger (A sketch from the COY scenario)** - An infant is confronted by a stranger in its familiar home environment. The infant smiles at the stranger, but withdraws its gaze. It may move towards its own carer, perhaps even wrapping its arms around the carer’s legs, before moving back and gazing towards the stranger again. If the infant stays in contact with the stranger for a prolonged period, or meets the stranger frequently, the infant shows less wariness towards the stranger.
>
> **Gaining contact when fatigued (A sketch from the CONTACT scenario)** - An infant that is exploring becomes fatigued, and then moves towards its carer and signals to its carer to gain close physical contact.

In all the sketches above, the infants are using their carers as secure bases, from which to explore their environment and to return to when they are anxious, wary, or fatigued. For each of these types of behaviour, infants behave differently with an attachment figure, which is usually one of their main carers, than they do with less familiar adults or strangers. An aspect of infant behaviour that the sketches above hold in common is that the infants do not learn about the effectiveness of their carer in responding. The SECURE-BASE scenario has been created and extends the exploration behaviour shown in the PARK scenario. This scenario describes how an infant’s Security or Insecurity in their relationship with their main carer can be derived from their carer’s pattern of responsiveness to signals.

> **Learning whether to be Secure or Insecure (A sketch from the SECURE-BASE scenario)** - A carer is busy collecting food, and is periodically interrupted by an infant signalling for attention. Over time there are many occasions when the infant bids for attention and receives responses which may vary in their promptness. When the carer is generally prompt in response the infant explores the world with more confidence, knowing that the carer is likely to ‘be there’ if called upon. When the carer is not generally prompt the infant stays closer to the carer, not confident that the carer will respond adequately if the infant moves further away.

The PARK, WARY, COY, CONTACT and SECURE-BASE scenarios are derived from separate empirical studies and illustrate particular dimensions of infant behaviour independently. These various dimensions of behaviour can be observed together, integrated in a single study. Patterns of behaviour similar to all of the sketches of infant behaviour above can be observed with infants undertaking the Strange Situation Experiment at the end of their first year (Ainsworth *et al.*, 1978). The Strange Situation Experiment is not an experiment where subjects are randomly assigned to different conditions in the laboratory. Rather, it is a standardised laboratory procedure that presents all infants with the same controlled and replicable set of experiences. Different outcomes arise from infants possessing different predispositions for behaviours related to attachment. These predispositions have been primarily gained from the infant's particular experiences with their carer over the previous year. The laboratory procedure consists of eight episodes of three minute duration which are designed to have the effect of activating and intensifying the attachment behaviour of one-year-old infants in a moderate and controlled manner. Some episodes involve the infant meeting an unfamiliar 'stranger' in the laboratory, other episodes involve the mother being removed from the room. In all of these episodes the infant's behaviour is carefully recorded from behind a two-way mirror. In the final episode of the Strange Situation Experiment the mother returns to her one-year-old infant after the infant has been left alone for three minutes in the unfamiliar setting. A sketch of this episode illustrates the possible range of responses in the reunion episode:

> **Being reunited with a carer in an unfamiliar environment after a short absence (A sketch from the REUNION scenario)** - A carer and infant are in an unfamiliar room and after a few minutes the carer leaves the infant with a stranger or completely alone. After a further few minutes the carer returns. Some infants will show some initial distress but approach the carer in a positive manner, seeking and receiving physical contact, and then quickly return to play and exploration. Others may show little distress, will not seek contact and may avoid their carer's gaze and physical contact, instead resuming play and exploration but with less concentration. Still others may respond by being overly passive or showing anger towards the carer and not resuming play and exploration.

Ainsworth *et al* (1978) found that when the measures of infant response in the Strange Situation were evaluated the infants were found clustered into three major categories of attachment style, labelled: Avoidant (type A), Secure (type B) and Ambivalent (type C)[2]. Ainsworth *et al* (1978) found that an infant's response to reunion in the Strange Situation can act as a shorthand for the infant's home

---

[2] More recent studies have categorised a fourth type of disorganised attachment style that forms a very small proportion of infants in the general population. It usually arises from a dysfunctional environment, and is not currently represented in the simulation

relationship with their carer. This is because the pattern of responses made by infants in this particular episode of the Strange Situation correlate most strongly with patterns of maternal behaviour and infant responses intensively observed throughout the previous year.

We can therefore outline a further three sample sketches from the REUNION scenario that connect home observations and the laboratory results seen at one year:

> **Avoidant infants** respond to their carer on reunion in the Strange Situation by not seeking contact or avoiding their carer's gaze or avoiding physical contact with her. These children return quickly to play and exploration but do so with less concentration than secure children. Whilst playing they stay close to and keep an eye on their carer. It may seem that they are not distressed or anxious in the Strange Situation. However, studies which measured physiological stress correlates of infants in the Strange Situation demonstrated stress levels of avoidant infants were at least as high as the secure and ambivalent groups (Hertsgaard *et al* 1995). In comparison with average levels across all groups: A type carers were observed at home being consistently less sensitive and providing more physical contact of an unpleasant nature; at home these infants were more angry and they cried more. However, in the reunion episodes of the Strange Situation these infants showed the least anger and crying.
>
> **Secure infants** respond to their mothers on reunion in the Strange Situation by approaching them in a positive manner. They then return to play and exploration in the room quickly. They received care at home which can be summarised as being consistently sensitive. In comparison with average levels across all groups: B type carers were observed at home being more emotionally expressive and provided less contact of an unpleasant nature; at home these infants were less angry and they cried less.
>
> **Ambivalent infants** respond to their carers on reunion in the Strange Situation by: not being comforted and being overly passive or showing anger towards their carers. These children do not return quickly to exploration and play. They received care at home which can be summarised as being less sensitive and particularly inconsistent. In comparison with average levels across all groups: C type carers were observed at home being more emotionally expressive; they provided physical contact which was unpleasant at a level intermediate between A and B carers and left infants crying for longer durations; at home these infant's were more angry, and they cried more.

Since the development of the Strange Situation as a method to investigate individual differences in attachment style, various theories have been put forward to explain its correlational results that link early carer and infant behaviour with later infant behaviour (Goldberg 2000, chapters 4 and 5). The two main theoretical camps are: firstly, those which emphasise evidence that the correlations are due to temperamental traits possessed by the infant that are considered to be innate; and secondly, those that emphasise evidence that infants are adapting to the caregiving environments provided by their mothers. A recent review of several meta-analyses that covered hundreds of studies on thousands of participants found that maternal sensitivity is a more important factor in the formation of infant attachment than infant temperament (van Ijzendoorn and Bakermans-Kranenburg, 2004).

> *"In sum, the causal role of maternal sensitivity in the formation of the infant-mother attachment relationship is a strongly corroborated finding. Correlational, experimental, and cross-cultural studies have replicated the association between sensitivity and attachment numerous times, and through different measures and designs. In general, the maternal impact on the infant-mother attachment relationship has been shown to be much larger than the impact of child characteristics such as temperament. During the first few years after birth, parents are more powerful than their children in shaping the child-parent bond."* (van Ijzendoorn and Bakermans-Kranenburg (2004) page 252).

Therefore this work has created a set of agent-based simulations that support theories that propose an infant's attachment style is an adaptation to the caregiving style they have previously experienced. Theories that incorporate individual differences in developmentally fixed mechanisms, such as those linked to temperament, are not incorporated in this present work. However, the simulations are sufficiently flexible for straightforward inclusion of these mechanisms at an abstract level in future work.

Part of the research problem that this work is attempting to solve involves creating simulations that explain normative behaviours that all infants develop, such as the exploratory, wary, coy and contact patterns of behaviour illustrated in the first four scenario sketches. For these normative behaviours we are not concerned with how these patterns of behaviour might develop in different caregiving environments, we are simply concerned with designing an agent architecture that can reproduce those patterns at an abstract level. This is because all infants, of whatever attachment style, produce these patterns of behaviour at some time. Part of the research problem involves creating simulations that do explain individual differences in patterns of behaviour. The SECURE-BASE scenario describes the production of two types of attachment response: Secure and Insecure. The REUNION scenario takes a more detailed approach to individual differences and breaks Insecure infant responses into Insecure-avoidant and Insecure-ambivalent style of response. The question of how the individual difference categories are formed in the SECURE-BASE and REUNION scenarios are developmental research problems, and since we are excluding theories based upon temperament the research problem can be restated. The research problem is to design infant agent architectures that exhibit the required normative behaviours, and in addition, in interaction with three different kinds of caregiving environments, develop the three major styles of infant attachment behaviour. The problem is therefore designing architectures with adaptive mechanisms that are effective in reproducing the behaviours in the scenario, and plausible with regard to the mechanisms they use to do this. It is hoped that by constraining the problem in this manner the solutions may be putative interpretations that allow progress towards firmer explanations in future. Having provided an overview of the behaviours that constitute the problem for this project we will now turn to the nature of the solution.

1.2 The nature of the solution

This work is a novel approach to understanding infant attachment behaviour that straddles the field of Attachment Theory and the broader discipline of Cognitive Science. It differs from other types of investigation in these areas in various dimensions. The approach it takes might equally well be described as Synthetic Modelling, Computational Modelling, Mind Design, the Design Based Approach or Psychology by Reverse Engineering. This section will show that each of these labels has merit as each emphasises a different aspect of the approach that has been taken in this thesis.

1.2.1 The design based foundations of Attachment Theory

Attachment Theory has always possessed a strong design based dimension. It originated primarily from the work of John Bowlby, who formulated the core of the theory whilst working in a multi-disciplinary team of childcare professionals, psychoanalysts, ethologists and other researchers at the Tavistock Clinic, for over thirty years after the Second World War (Holmes, 1993). This position enabled him to form a theoretical framework by integrating ideas from these diverse fields. Holmes (1993) reports that the origination of the theory ‘came’ to Bowlby after a discussion about ethological theory with Robert Hinde, an eminent researcher in this field. Bowlby at this time must have been aware of much of the behavioural phenomena that he wanted to explain. Exposure to the design framework provided by Ethological Theory seems to have crystallised a design based explanation of these behavioural phenomena, which Bowlby then spent many years elaborating. Design based components that Bowlby (1969) incorporated include:

- Ethological concepts and mechanisms, such as Behaviour Systems, Reflex Actions and Fixed Action Patterns which can interact in complex ways by chaining and alternation;
- Concepts from the study of Control Systems such as feedback and goal directed mechanisms;
- Concepts from AI and Cognitive Science such as Internal Working Models (IWM's) and hierarchical organisation and control of behaviour using complex representational forms such as natural language.

Bowlby (1969) described the development of attachment in childhood as occurring in four phases. For our purposes the key phase change occurs at the middle to end of the first year. In this phase change infants progress from the second 'fixed action pattern' phase to the third 'goal corrected' (goal directed) phase of infant attachment. This third phase is maintained until the age of three, and therefore in this phase goal directed mechanisms become augmented with higher level mechanisms such as IWM's and linguistic representations.

Bowlby's original design based ideas have been developed by other researchers in a variety of directions. Some developments have been focused at the same high level of description as Bowlby's original account (Main *et al.*, 1985; Main, 1991; Bretherton and Munholland, 1999). Attachment Theory has also progressed to include low level physiological descriptions and explanations of attachment phenomena (Fox and Card, 1999; Kraemer, 1992; Hofer, 1995; Polan and Hofer, 1999; Suomi, 1999).

1.2.2 Uphill Analysis and Downhill Synthesis

The design basis of attachment behaviour has been elaborated with a variety of methods. Until this current work, these methods have not included actually synthesising artifacts which simulate processes of attachment. Haugeland (1997) introduced the term Mind Design to describe a synthetic approach to model formation in psychology and Cognitive Science:

> *"MIND DESIGN is the endeavor to understand mind (thinking, intellect) in terms of its design (how it is built, how it works). It*

> *amounts, therefore, to a kind of cognitive psychology. But it is oriented more toward structure and mechanism than toward correlation or law, more toward the "how" than the "what", than is traditional empirical psychology. An "experiment" in mind design is more often an effort to build something and make it work, than to observe or analyze what already exists. [...] Of course, natural intelligence, especially human intelligence, remains the final object of investigation, the phenomenon to be understood. What is distinctive is not the goal but rather the means to it."*(page 1, Haugeland (1997)).

We can contrast the synthetic approach taken by this work with the analytic approach taken by previous work in Attachment Theory. For example, Ainsworth *et al.* (1978) characterised the three styles of attachment that are the target of our work by analysis that involved repeated observations, numerous coding schemes and intensive application of statistics. So why not explicate the internal structures that give rise to attachment behaviours with further analysis? Advantages of taking a synthetic approach are captured by Braitenberg's 'law' of 'Uphill Analysis and Downhill Synthesis' (Braitenberg, 1984; Dawson, 2004):

> *"It is much more difficult to start from the outside and try to guess internal structure just from the observation of the data"* (page 20 Braitenberg (1984) [3]).

Braitenberg claims that not only can synthesis be easier, but the theories that result from synthesis can be simpler than those arrived at from analysis. This can be because psychologists over-estimate the complexity of the phenomena that they are analysing, or it may be that complex patterns of behaviour emerge from simple building blocks in unforeseen ways. We can see the operation of Braitenberg's law in chapters three and four of this thesis, when we provide a novel explanation of the statistical pattern of the clustering of the results from Strange Situation studies.

1.2.3 Complete Agent Architectures

In his 2001 review article: *"The Agent-Based Approach: A New Direction for Computational Models of Development"* Schlesinger (2001, page 129) suggests that *"aspects of social interaction open to investigation include the evolution and development of attachment"*. This work is a fulfilment of this suggestion. It is the first implementation of a simulation that explains the three principal types of infant attachment as adaptations to caregiving style. Agent-based simulations

---

[3] also cited in Dawson (2004) page 99

are an emerging trend in computational developmental psychology because they are models of whole systems which include perception, action and internal processes embedded in a dynamic environment. This means that they can simulate continuously changing environments where actions have important short term consequences for subsequent behaviour.

This work differs from many developmentally oriented agent-based simulations and other relevant work, such as social interaction in infant-like robots, because of its concentration on central processing and impoverished perceptual capacity (Schlesinger, 2001; Breazeal and Scassellati, 2000; Likhachev and Arkin, 2000).

1.2.4 Reverse Engineering

Labelling the work in this thesis as examples of Mind Design or the Design Based Approach (Petters, 2004; Beaudoin, 1994; Wright, 1997) is apt because these terms emphasise the important engineering aspects of this project. We can make an important distinction between different types of design based work. This is to consider whether it is involved in forward engineering or reverse engineering. Forward engineering is a familiar concept. Much of what the software industry does is forward engineering. When forward engineering is also attempting to perform the same information processing operations that occur in humans it is engaging in Mind Design. Reverse engineering has important differences. Dennett states:

> *"Reverse engineering is just what the term implies: the interpretation of an already existing artifact by an analysis of the design considerations that must have governed its creation."*(page 684, Dennett (1994))

The difference can be illustrated if we consider how forward and reverse engineers would tackle a challenge of reproducing some human capability, say visual tracking. The forward engineer would just build an artifact (robot or software program) that accomplished this capability, *"however he wanted to"* (page 683 Dennett, 1994). The reverse engineer would show, *"through building, that he had have figured out how the human mechanism works"* (page 683, Dennett, 1994). Another way of saying this is that by using a reverse engineering approach to study a living system means explaining the kinds of structures and mechanisms which might be required in a systems design, to enable that system to be reproduced in simulation. This is precisely the aim of this work. The infants that have been observed in studies of attachment are treated as designed objects and their behaviours are treated as though they have arisen from designed information processing architectures. The designer is taken to be the process of evolution. Dennett (1994) notes that when this is done the engineer makes the assumption that although *"the historical design process of evolution doesn't proceed by an exact analogue of the top-down engineering process. ... Reverse engineering is just as applicable a methodology to systems designed by Nature, as to systems designed by engineers."* A presupposition has been made *"that even though the forward processes have been different, the products are of the same sort, so that the reverse process of functional analysis should work as well on both sorts of product."*(page 684, Dennett (1994)).

There are several issues that arise when studying phenomena using an approach based upon reverse engineering. Firstly, a reverse engineer uses data and theory to constrain the design that they produce. In the case of simulating attachment behaviour of one-year-old infants, the designer should be invoking mechanisms that are plausible for infants of this age to possess. The artifact being reverse engineered has to have some kind of evolutionary function, or be the bi-product of some evolutionary function. The issue of adaptive function and reverse engineering is dealt with in more detail in chapter two.

1.3 Overview of the design process

In general terms, using an engineering methodology to produce a scientific explanation of some phenomena involves three stages[4]. First the phenomenon is framed as a problem to be solved. This means the problem is presented as a set of requirements to be met, or specification to be fulfilled. Once source material has been selected which describes the required observable behaviour and the constraints that must be met have been decided upon, this information must be repackaged in a form that will facilitate later evaluation. This is done by abstracting the source material and reforming it into concrete mini-scenarios which will act as clear 'goal-posts' for the designs to meet. Part of the requirement is to reproduce observed patterns of behaviour. Another part is that these behaviours are produced in pursuit of plausible functions or goals. Infants switch between attachment behaviour and other activities, such as exploring their environment. Therefore a simulation of attachment needs to include mechanisms for how infants interact with and learn about objects they perceive.

The next stage involves producing one or more designs that attempt to solve the problem by meeting the requirements. The process of designing solutions

---

[4] The nature and description of this process is evolving as the project progresses. The process has previously been described as a five part process ((Beaudoin, 1994) page 5 and (Wright, 1997) page 22), and a seven part process (Petters, 2004).

to the behaviour, as set out in the scenarios, is accomplished initially in linguistic terms and then implemented in autonomous agent simulations using the Sim-agent toolkit.

Finally the solution designs are then evaluated in various ways as to how well they meet the requirements. A set of constraints needs to be decided upon that allows acceptable, plausible solutions to be discriminated from unacceptable, merely superficial solutions. The simulation must be consistent with what is known about infant abilities between nine months and one year of age in other psychological domains. For instance our representation should be rich enough to include mechanisms that support the range of behaviours that can be observed in nine month to one years old infants in studies of cognitive development (Diamond, 2000; Willatts, 1989). The implemented designs are evaluated in two broad ways. First the simulations are evaluated for internal consistency: do they reproduce the behaviour set out in the specification of requirements? Secondly, the simulations are assessed for goodness as theories of the target phenomena. Any architecture that can reproduce the behaviour required by the scenario has passed a form of sufficiency test and is a 'proof of concept' for that theory (Cooper, 2002). External validity is related to how well the specification of requirements, in the form of a scenario, represents the behaviour we are trying to explain. Architectures that fully reproduce the scenario can be assessed against each other, and may differ in how they fulfil the scenario, whether by principled or ad-hoc means. Additional constraints can be derived from a wealth of linked empirical data and theory from cross-species, evolutionary, neurophysiological and cross-cultural branches of Attachment Theory. Constraints can also be obtained from links to other proposed mechanisms explaining different phenomena in "broad" architectures.

The design stages can be followed in series in any order, in parallel or some mixture of the two. These processes may loop many times.

1.4 The need for an analysis of function

The essence of this project is reproducing some observed behaviours in a simulation. One way that we might proceed is to reproduce the behaviours in a highly detailed and yet superficial manner. The resulting simulation might be something like the Kismet robot's interaction with humans (Breazeal and Scassellati, 2000), with the additional constraint that it closely follows the behaviours set out in the source material for this thesis. It would be superficial in the sense that it would be the observable behaviours that are important. In a deep representation of behaviours the observable behaviour may be less convincing copies of real behaviours, but the underlying mechanisms that support the simulated and real behaviours would be more similar. Superficially producing behaviours would be a forward engineering project. However what we are interested in is reverse engineering these behaviours. This means we are not interested in simple but psychologically implausible solutions such as matching correlated behaviours with look up tables or 'If...then' rules. We do not require detail in the representation of the behaviours, neither do we want to simply capture the stimulus-response contingencies present in the target data. Rather, we want to use the simulation to capture the actual causal relations that drive patterns of behaviour in real infants.

Dennett compares reverse engineering in Cognitive Science with reverse engineering that an industrial company might carry out when it wants to understand the products of its rivals. In this context:

> *"if the reverse engineers can't assume that there is a good rationale for the features they observe, they can't even begin their analysis."*(page 683 Dennett (1994))

This notion applies as much to reverse engineers working with artifacts designed by the process of evolution as with human designed artifacts. To reverse engineer an infant behaviour or pattern of behaviours we should be able to assume these behaviours either: serve a function; are a necessary consequence of some other function; or are dysfunctional consequences of mechanisms that are normally functional. For example one might try to reverse engineer some pathological behaviour in order to decide which part of the brain might have been damaged.

1.5 Attachment and modern evolutionary thinking

1.5.1 Attachment Theory in the hierarchy of evolutionary theories

Since its publication, Bowlby's (1969) original formulation of Attachment Theory has become closely integrated with modern evolutionary thinking, being drawn into a broader evolutionary framework as a middle level theory (Simpson, 1999). In addition, the conceptual apparatus developed by Bowlby to explain attachment phenomena, such as the concept of the 'Environment of Evolutionary Adaptedness' (EEA), has been borrowed from Attachment Theory and used widely in other modern evolutionary theories (Laland and Brown, 2002).

Attachment Theory and the high level theoretical framework into which it now fits were both developed in the 1960's. The theory that initiated modern

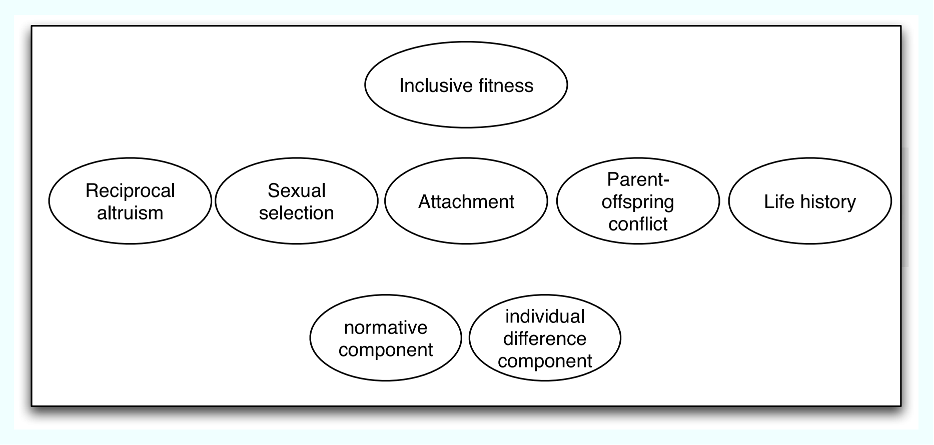

Figure 1.1: The hierarchy of evolutionary theories, adapted from Simpson (1999)

evolutionary thinking was Hamilton's (1964) 'Theory of Inclusive Fitness'. This theory advanced Darwin's ideas by emphasising the importance of taking a 'Genes Eye View' of evolution[5]. The concept of Inclusive Fitness includes Darwin's concept of fitness, that is: fitness due to one's own reproduction, but in addition it incorporates Hamilton's concept of fitness, which is fitness due to one's biological relatives (Simpson, 1999). This new concept allows us to understand phenomena which puzzled Darwin, such as why some organisms do not attempt to reproduce, instead assisting the reproductive efforts of individuals with which they share a large proportion of genes.

The 'genes eye view' catalysed the production of a stream of middle level evolutionary theories. Simpson writes that *"Inclusive fitness theory ... is the superordinate, general theory of evolution from which nearly all middle level theories are derived"*. This relationship is represented in figure 1.1, which includes Attachment Theory as a middle level theory alongside four others theories most relevant to understanding the behavioural phenomena of interest in this thesis. Reciprocal altruism is relevant to understanding attachment because it explains the benefits which can be reaped by possessing well developed social skills. The theory of sexual selection explains why an infant's mother rather than its father is likely to have been the main carer in the human 'Environment of Evolutionary Adaptedness' (EEA). The theory of Parent-offspring conflict proposes that it is sometimes evolutionarily advantageous for carers to act in decidedly non-

---

[5] Darwin can't be blamed for not taking a gene's eye view, his work was written before Mendel's discoveries of heritable traits in plants, and also therefore before a full understanding of how sexual recombination and genetic mutations provide the variation from which better adaptations are selected (Simpson, 1999).

	The behaviour IS adaptive	**The behaviour IS NOT adaptive**
The behaviour IS an adaptation	**Current adaptation**	**Past adaptation**
The behaviour IS NOT an adaptation	**Exaptation**	**Dysfunctional by product**

Figure 1.2: The difference between adaptive behaviour and adaptations, (adapted from Laland and Brown (2002)).

benign ways towards their own offspring. Life history theory assesses the possible functions of the different styles of infant attachment at all periods through an individual's lifespan. Below the middle level theories are the major principles or components of those theories. In figure 1.1 these have only been added for Attachment Theory. The normative component is described in more detail in section 2.2. It deals with relatively universal, stable patterns of behaviour. The individual difference component is described in section 3.2.2 and 4.2.3 and assesses the evidence that each of the three main patterns of infant attachment behaviour is an adaptation.

Although the first part of Bowlby's trilogy was published in 1969 it didn't cite Hamilton's 1964 paper or the concept of inclusive fitness. Attachment Theory was framed by Bowlby in terms of evolution acting upon differential survival rather than differential reproduction. Attachment Theory is a lifespan evolutionary theory, and so ultimately should be concerned with reproduction. However, we are applying this theory to the behaviour of one-year-old infants. Whether we are only concerned with survival, or whether we view the infant as only surviving through the perils of infancy as one of the barriers to inclusive fitness, does not affect the assessment of evolutionary function for the behaviours we want to simulate (Simpson, 1999).

1.5.2 Adaptationism versus adaptivism

Laland and Brown (2002) characterise two contrasting evolutionary approaches to understanding the function of behaviour. If a behaviour possesses an adaptive function then it benefits the individual performing the behaviour, in the sense of increasing that individual's reproductive fitness. If a behaviour was selected for its effectiveness in a particular role in the evolutionary past of the individual's species, then it is termed an adaptation. Figure 1.2 presents these differences in tabular form and shows the four alternatives.

Studies that invoke evolutionary analyses of function can take an adaptivist approach. This form of study is often characterised by making an assumption that humans are cognitively flexible enough that their behaviour is optimally adaptive (Laland and Brown, 2002). The foraging study of Smith (1985) is one example of an adaptivist approach of this type. The investigators made detailed observations of foraging behaviour in the Inuit population of Arctic Canada and then compared this behaviour with theoretically derived models of optimal foraging behaviour that gave predictions of optimal net capture rate. Deviations from the predicted optimal foraging behaviour allowed Smith (1985) to conclude that social obligations, in the form of taking others along on hunts, were what decreased the actual net capture rate from the predicted optimal value. In this study the individuals being studied were adults, the hunter-gatherer environment they lived in can be assumed to be close to that in which humans evolved, and the time-scale of the decisions they make can be assumed to allow optimally rational processing. When individuals are too young, in a modern environment that doesn't match with the human EEA or we are interested in behaviours that are produced over short time spans, adaptivist assumptions of optimally rational behaviour are more difficult to support.

Since this project is considering infant behaviours observed in modern environments that have been measured over short timespans it has therefore taken an adaptationist approach. If our aim is to reverse engineer behaviours we should avoid exaptations, which do enhance fitness but were not built by natural selection for their current role (Laland and Brown, 2002). The two cases of interest are current adaptations and past adaptations (see figure 1.2) . Both of these types of behaviour will have been adaptive in the evolutionary past but in the case of past adaptations a change in the selective environment means these behaviours are no longer adaptive.

When researchers focus on behaviours that are considered to be adaptations what they are really interested in is the mechanisms that give rise to those behaviours. Several authors have argued that natural selection doesn't operate directly on behaviours, rather it operates on the behavioural mechanisms that underpin behaviour (Symons, 1990; Tooby and Cosmides, 1990). Assessing whether behaviours are species specific adaptations for humans means considering the human Environment of Evolutionary Adaptedness (EEA), which is the subject of the next section.

1.5.3 The nature of the evolutionary niche in which our ancestors evolved

Evolutionary analysis of the functions of currently observed infant behaviours was undertaken by Bowlby (1969), who coined the phrase 'Environment of Evolutionary Adaptedness' (EEA) to describe the conditions which acted as selective pressures in human evolutionary past. An EEA is therefore a type of biological niche, that may differ from the current environment for human infants.

The concept of niche has aspects in common with what an engineer would call a 'requirements specification', however the concept of a niche (such as the human EEA) *"is more subtle than a standard set of engineering requirements since the nature of the niche, i.e. the collection of requirements for an organism to function well, is not some fixed externally determined specification, but depends in part on the organism's existing features and capabilities."*.(page 1 Sloman (2000)). So when we are theorising about the nature of the human EEA, we need to be aware that the nature of an EEA changes for an individual as that individual develops. The EEA for a newborn, a crawling infant and a walking and talking child will all be different, as infants or children at these ages will be required to interact differently, will have the potential for getting into different types of danger and will be driven to carry out different forms of exploration and learning.

In addition, the EEA is more abstract than simply being a location or set of locations and doesn't apply to a particular time period, such as the Pleistocene or Paleolithic. As Tooby and Cosmides (1990) note:

> *"The EEA concept does not refer to a single place or habitat, or even a time period. Rather it is a statistical composite of the adaptation relevant properties of the ancestral environments encountered by members of ancestral populations, weighted by their frequency and their fitness consequences"* (Tooby and Cosmides (1990) pages 386-7 cited in Laland and Brown (2002)).

1.5.4 We only need to focus on the infant EEA

Such an abstract conceptualisation can make it difficult to draw concrete conclusions about the nature of the EEA. It will have had no single pattern of weather and no set source of food. However, since we are focused on the nature of the EEA as it applies to infants of the age range nine months to one year old considerations are more simple, and we can make some generalisations. The main factor that is likely to have been in common in all the different environments that make up the composite EEA is the nature of the mother-infant relationship, and the place of this relationship in the broader social group. In a new born human's EEA the infant was almost completely reliant upon its mother for all its needs. The mother and infant will have lived in a social group that had from about a dozen members to perhaps two hundred (Bowlby, 1969). Therefore much of an infant's interaction will have been mediated through its relationship with its mother, and its mother's position in a broader group. Section 2.2.2 discusses the components of the human EEA that give rise to the normative aspects of infant behaviours, particularly those that apply to exploration-security trade-offs and fear of unfamiliar people and environments. Sections 3.2.2 and 4.2.3 deal with the components of the human EEA that gives rise to individual differences in infant behaviours, particularly those recorded in Strange Situation studies. If the capability to develop different styles of attachment is an adaptation then it may have been selected to deal with carers who are ineffective or constitute a threat.

1.6 Scenario formation

The central concern of this thesis is to simulate infant behaviours that may vary in numerous dimensions. The important characteristics of the behaviours that we want to capture may involve numerical quantities, such as the frequencies that particular behaviours are carried out. However there may be aspects that are more difficult to quantify, such as capturing rule based patterns of behaviour. An effective method of assessment in this case is to use structural descriptions of what has and has not been achieved. This thesis therefore uses a scenario based method of evaluation because it provides precise metrics for elements of the behaviours we want to capture that are not easily represented in simple quantitative ways[6].

Scenarios should capture abstract patterns of behaviour and reform them at a level of concreteness and detail appropriate for the requirements specification of the design and implementation phases. Scenarios are central to the validation of any designs that are implemented because they form a detailed and graduated requirement against which competing implementations can be assessed. The structure of the scenarios that have been produced in this thesis have been adapted from a scenario template produced for the Cogntive Systems (COSY) project[7]. Each of the scenarios in this thesis possesses sections describing a number of key sections, namely:

---

[6] this section is based upon the motivation and description found in: http:www.cs.bham.ac.uk/research/cogaff/gc/targets.html

[7] The full template can be found at: http://www.cs.bham.ac.uk/research/projects/cosy/scenarios/scenario-template.txt and a report which describes the production of a partially ordered network of scenarios, based on a competences grid, can be found at: http://www.cs.bham.ac.uk/research/projects/cosy/papers/sloman-aaai06-member.pdf

- Scenario summary
- Task analysis
- Scenario ontology
- Scenario scripts
- Negative Scenarios
- Kinds of integration
- Form of evaluation

Each scenario summary gives an overview of the behaviours that are to be reproduced. Since we are reverse engineering the infants they should do something in the world; they should have some purpose in being. Therefore there is a task analysis of the behaviours that assesses what function, if any, the behaviours possess. Either implicit or explicit goals can then be ascribed. In addition to any goals gained from assessing the particular behaviours of the infants there are also more general goals that arise because the infants are being represented as autonomous agents. For example, each infant agent should to be able to maintain multiple goals and cope appropriately and in a timely manner with changes in its dynamic environment. This criterion includes the ability to interrupt actions if more urgent actions become necessary and the ability to act opportunistically, such as the ability to act upon available resources when they occur (Clark, 1997).

The scenario ontology provides a description of the kinds of objects, properties, processes and interrelationships between them that will be found in the virtual world. Scenario scripts are descriptions of sample behaviours that demonstrate competences that need to be explained. Amongst other things, scripts indicate variability of behaviours and the types of learning that occur. Negative scenarios describe patterns of behaviour that should not occur in the simulation. They are important because we are targeting a particular stage of development so we want to rule out behaviours that the mechanisms we are engineering should not be able to produce. The kinds of integration may be between: different levels of processing, kinds of task, forms of representation or between kinds of learning and development. The form of evaluation may be an analysis of the running program or trace printing from previous runs. It may involve specification of different types of computational experiment that should be run.

1.7 Overview of results: scenarios and architectures

The designs that are produced to meet the requirement are produced at two levels, as linguistically described designs and as designs implemented in a software simulation[8]. The solutions are in the form of broad and shallow architectures (Bates *et al.*, 1991). The architectures can be described as broad agents because they possess the complete set of perceptual, action and other subsystems needed to act autonomously within their virtual environment. It is this breadth that allows incorporation of different AI and cognitive psychology techniques and theories in the simulation. They can be described as shallow agents because none of these subsystems is defined in great detail. This shallow aspect allows the simulation to be primarily directed at the level of constructs such as goals and action plans without being submerged in low level detail.

The designs described in this thesis have been inspired by existing theories from the attachment literature, derived from other sources or developed from scratch during the design process. Of these influences, the most fundamental influence in shaping the project is Bowlby's original formulation of Attachment Theory (Bowlby, 1969). The commitments of the theoretical framework followed by this work, and which derive from Bowlby's theory of Behaviour Systems, are:

- Infant behaviour is organised by structure termed behaviour systems
- Each behaviour system carries out a species specific function
- The behaviour systems most closely related to attachment are inherently motivated
- Behaviour systems possess a flexible repertoire of behaviours
- Behaviour systems involve a hierarchy of forms of information processing
- Behaviour systems are described at a high level of implementation

1.7.1 Information processing explanations in Attachment Theory

The key construct that Attachment Theory uses to explain how and when actions are selected is the concept of a control architecture formed by interacting behaviour systems. Behaviour systems have been theorised by Bowlby (1969) and

---

[8] The architectures have been implemented using the simagent toolkit, which is described in appendix A

other researchers of infant attachment phenomena to possess a number of properties that the simulations have attempted to incorporate in their designs (Cassidy, 1999).

1.7.1.1 Each behaviour system carries out a species specific function

Attachment Theory states that behaviour systems control a broad set of behaviours. In animal ethology, behaviour systems are theorised as controlling behaviours such as mating, fighting and feeding. Each behaviour system carries out a species specific function, and has been selected for this function in the evolutionary past. Hinde recounts that:

> *"The concept of a behavioural system is, in fact related to one meaning of the term instinct. [...] It has been used in a rather special sense by ethologists to refer to systems postulated as controlling a group of behaviour patterns that together serve to achieve a given biological end"((Hinde, 1983) page 57).*

The behaviour systems that Bowlby linked to attachment behaviour in human infants are the attachment, fear, sociability and exploration systems (Bowlby, 1969). In the view of Attachment Theory, behaviours resulting from the attachment behaviour system and the fear system have the predictable outcome of increasing the proximity of the attachment figure to the infant, the principal means by which infants gain security. The exploratory behaviour system activates behaviours that result in learning and the sociable system results in social interaction. The result of this theoretical commitment is that the agent based simulations reported in this thesis possess goal or motive based architectures. When a behaviour system in Bowlby's sense becomes activated it has been interpreted as requiring the activation of a goal or motive, though these may be implicitly represented.

1.7.1.2 The behaviour systems most closely related to attachment are inherently motivated.

Infants will work to experience exploration, socialisation and security because these outcomes can be considered primary drives. They are not activated as the by-product of any more fundamental process (Cassidy, 1999). This has come to be the mainstream view in child development, but this was not always the case. For example Smart and Smart (Smart and Smart, 1967) held the view that infants had to have returning behaviours reinforced by their carers or else *"there would be no more babies, since babies creep and toddle right into danger"*.

This inherent motivation to explore was recognised by Piaget, who described how a child's interest in exploration drives cognitive development (Marvin and Britner, 1999). What this means is that infants don't learn, by some process of association or re-inforcement, to use their carers as secure bases from which to explore the world (Bowlby, 1969). Running away, freezing and using a carer as a secure base are all behaviours that humans and other animals may instinctively 'know' to do when faced with danger of particular types. This has important implications for the design of the simulations. As designers we are not required to build architectures where desired behaviours somehow emerge from more basic types of behaviour.

The consequence of this theoretical commitment is that the goal activation mechanisms for the goals of attachment, fear, exploration and socialisation have all been hand-coded into the infants from the start of their simulated existence, rather than being required to somehow emerge.

1.7.1.3 Behaviour systems possess a flexible repertoire of behaviours.

What defines a behaviour system is not a set repertoire of behaviours but the outcomes that predictably follow from the carrying out of those behaviours. Similar behaviours may be produced by different behaviour systems. For example, behaviours such as locomotion can serve more than one system, such as the attachment and exploratory system. (Marvin and Britner, 1999). Also any given behaviour system may produce a wide range of differing behaviours. In the attachment system, if the infant possesses the goal of increasing its proximity to a carer the infant may cry to bring the carer closer or crawl towards the carer themselves. This is an example of behaviours within systems being interchangeable with other functionally equivalent behaviours.

According to Attachment Theory, the attachment behaviours used to gain proximity range from: subtle signals, such as gazing towards a carer, to overt signals, such as calling a carer, and active behaviours, such as locomotion. Exploratory behaviours range from locomotion to object manipulation and have the predictable outcome of improving the infant's ability to manipulate the external world. The consequence of this theoretical commitment is that is has allowed the simulations to be set at a high, goal/motive oriented level of description.

1.7.1.4 Behaviour systems in humans involve a hierarchy of forms of information processing

In his description of Behaviour Systems Bowlby invoked a variety of mechanisms. These ranged from simple mechanisms, such as reflexes and fixed action patterns, through goal corrected mechanisms, to complex mechanisms and forms of representation, such as Internal Working Models and natural language.

Reflexes are behaviours with a highly stereotyped form. Once activated by a stimulus at a specific threshold they are ballistically carried to completion (Marvin and Britner, 1999). Fixed action patterns are similar to reflexes because they are stereotyped but differ from reflexes because they are open to learning. The thresholds for activation and termination adapt according to the state of the organism and past experiences. Fixed action patterns are also less ballistic, with for example, proprioceptive feedback during execution. Examples of this type of behaviour include grasping, crying and smiling (Marvin and Britner, 1999). Reflexes predominate in the first few months after birth and fixed action patterns predominate from three months until the middle and end of the first year. The reflexes and fixed action patterns that infants perform may seem a very simple form of control but are highly effective in eliciting adult actions that benefit the infant. Different reflexes and fixed action patterns are coordinated together, the sum of these behaviours working together is greater than their parts. As Marvin and Britner (1999) note:

> *"The immediate effect of any behaviour is to bring about a change in the environment, which itself serves as an activating condition for another behavfiour, often forming a lengthy sequence with an eventual outcome that is necessary for the individual's survival."* (Marvin and Britner, 1999, page 48).

These complex patterns produced by fixed action patterns can be mistaken for behaviours directed by goal corrected mechanisms because of the sensitive matching of response to stimuli. The goal corrected phase of attachment commences from the middle to the end of the first year. When a mechanism is goal-corrected it is updated or retaken according to feedback on how well the goal has been satisfied.

The concept of Working Models came to Attachment Theory from the work of Craik (1943). In a broader sense Working Models are not confined to attachment but apply to all representative models of the world. In his work Bowlby restricted the term Internal Working Models (IWM's) to models of self and other in attachment relationships. IWM's capture the relation-structure of phenomena, not every aspect of reality but enough to make possible the evaluation of alternative actions.

These include spatio-temporal causal relations among the events, actions, objects, goals and concepts represented. IWM's of attachment are what hold an infant's expectations of the levels of predicted availability and responsiveness for a given carer. These expectations are derived from the carer's past performance. IWM's of self and attachment figure develop in a complementary manner. For example if the carer is responsive the self is valued.

Bowlby's invocation of such wide range of forms of representation and types of mechanism has meant that the design considerations in these areas have been minimally constrained by commitments from Attachment Theory.

1.7.1.5 Behaviour systems are described at a high level of implementation.

Behaviour systems are described at the level of software or virtual machine operation, as opposed to a neural level of description. Hinde (1983, page 57) explains:

> *"Although the behavioural systems were postulated to explain the observed behaviour, there was no necessary implication of isomorphous mechanisms in the brain. The explanation is a software one, comparable to a computer program which performs a particular job irrespective of the details of the hardware of the computer into which it is fed."*

This theoretical commitment reinforce the case for positioning the initial efforts at simulation at a high, goal orientated level. Since Bowlby's work a number of studies, including cross-species investigations, have uncovered many low level details of attachment mechanisms, and future work may incorporate these findings (Kraemer, 1992; Hofer, 1995; Polan and Hofer, 1999; Suomi, 1999; Fox and Card, 1999).

1.7.2 Description of the six scenarios

This thesis has produced six scenarios and produced five implemented architectures (see figure 1.3). The scenarios present challenges of increasing difficulty and the architectures produced in response therefore increase in the complexity of the structures, mechanisms and representations that they possess. The PARK, WARY, COY AND CONTACT scenarios represent short periods of time.

In the case of the park exploration in the PARK and WARY scenarios, the time period may be up to 20 minutes. In the case of the COY scenario the maximum time would be less than this. The CONTACT scenario may be longer. What the PARK, WARY, COY and CONTACT scenarios have in common is that in each

Scenarios	**short name**	**Architectural solution**
Illustrating how infants balance exploration and security in a park	PARK scenario	**goal-switching (GS)**
Illustrating how unfamiliarity modulates patterns of exploration	WARY scenario	**goal-switching (GS)**
Illustrating how infants use coy behaviour when meeting a stranger to balance socialisation and security	COY scenario	**goal-switching (GS)**
Illustrating how infants use their carer as a secure-base when fatigued	CONTACT scenario	**goal-switching (GS)**
Illustrating how style of caring gives rise to the distinction between Secure and Insecure infants	SECURE-BASE scenario	**goal-learning-from-anxiety (GLA)** **goal-learning-from-socialisation (GLS)**
Illustrating how style of caring gives rise to the distinction between Avoidant and Ambivalent infants	REUNION scenario	**reactive-action-learning (RAL)** **hybrid-action-reasoning (HAR)**

Figure 1.3: Relating the four scenarios to the six architectural designs.

scenario the infant possesses two competing goals and in each scenario the infant spends some time attempting to achieve both goals. In the PARK, WARY and CONTACT scenarios the infant is attempting to explore inanimate objects whilst remaining safe. In the COY scenario the infant is attempting to interact with a stranger whilst remaining safe. Each of these scenarios have been simulated with a single architecture, termed the **goal-switching** (**GS**) architecture. This is the simplest architecture implemented in the thesis. It is reactive and doesn't explain any individual difference data or incorporate any long term learning mechanisms. It explains how infants switch between goals in a short time period.

The SECURE-BASE scenario is an extension of the PARK scenario. It represents the whole first year of the infant's life. The scenario incorporates long term learning. The SECURE-BASE scenario describes how infants adapt to caregiving styles that vary along a dimension from very sensitive and responsive to insensitive and unresponsive. The required behaviour of the infant is therefore conditional. The infants in each conditional strand of the scenario must be able to adapt to the pattern of caregiving that they experience. The solution to the requirements are infant architectures, and any candidate architecture has to adapt to be in a Secure state if the carer is sensitive and responsive and adapt to be Insecure if the carer is insensitive and unresponsive.

Two implemented infant architectures have been created that both fulfil the broad outlines of the SECURE-BASE scenario, but that differ in the details of how they fulfil the scenario. Each implementation is from a class of architectures that are termed **goal-learning** (**GL**) architectures. These use the **GS** architecture as a foundation and so might strictly be termed 'goal-switching and goal learning' architectures. The two variants of **GL** architecture are termed the **goal-learning-from-anxiety** (**GLA**) and the **goal-learning-from-socialisation** (**GLS**) architectures. These architecture differ in the information they use in assessing the sensitivity and effectiveness of the carers. The **GLA** architecture only assesses these characteristics from events where the infant determined that it was in danger. The **GLS** assessed these caregiving characteristics from events where the goal for socialising was activated, and therefore possessing greater activation than the goal to be secure. Computational experiments involving these two architectures have attempted to investigate the statistical distribution of Secure and Insecure attachment styles in real infants. When populations of real infants are rated for security of attachment it is found that the results cluster. What this means is that values for security are not spread evenly in the population or distributed according to an approximation of a normal distribution, where the most popular values would be the mean value. There is a more frequent incidence of infants that are either clearly secure or insecure and a less frequent incidence of intermediate cases compared to a normal distribution of security level. Computational experiments show that when the style of caregiving is set to be intermediate between secure and insecure the **GLA** architecture results in infants forming two clusters with high and low values for security whereas the **GLS** architecture results in a spread of varying levels of security, including more intermediate cases and less clear clustering. These results make an important contribution to the understanding of how clustering of individual difference categories in infant attachment behaviour may be formed.

The REUNION scenario describes how the two Insecure styles of attachment come to be distinguished. The principal difference between the experience of Avoidant and Ambivalent infants in the scenario is the amount of unpleasant close physical contact infants experience in the training stage. In empirical studies there are a range of measures related to close physical contact that show significant differences between the performance of Avoidant and Ambivalent caregivers. These include: Avoidant caregivers being more averse to close contact ($p<0.0001$), and providing more close physical contact of an unpleasant nature ($p<0.02$). Ainsworth *et al.* (1978) note that *"aversion is only rarely expressed openly, and could well have been missed had the home visits not been long and frequent enough to encourage the mothers to behave naturally. ... Our hypothesis is that the underlying aversion and implied rebuff nevertheless communicates itself to the baby"* (page 151, Ainsworth *et al.*, 1978).

The brief test stage of the REUNION scenario captures the pattern of Avoidant and Ambivalent behaviour in the second reunion stage of the Strange Situation. The Avoidant infants initially move towards their carers, but then stop. Ambivalent infants move and signal loudly even when in close proximity. Two designs have been implemented that reproduce these requirements. The first implemented architecture to respond to the fourth scenario is a reactive architecture that incorporates the **GLA** structures and mechanisms, and in addition allows the infant to choose which action to take when the goal of safety is activated. This architecture is termed the **reactive-action-learning** (**RAL**) architecture. The principal design element that sets the **RAL** architecture apart from the **GLA** architecture is the manner in which the generated goals interact, before final selection.

A hybrid architecture has also been created that explains the same behaviour as the **RAL** architecture. This architecture is termed the **hybrid-action-reasoning** (**HAR**) architecture, and it possesses the same goal generators and the same repertoire of possible actions as the **RAL** architecture. The **HAR** architecture differs from the **RAL** architecture because instead of reactive, preprepared arbitration of the competing goals, this architecture allows on-line simple reasoning about what actions to take.

Each of the implemented architectures is evaluated for how well it satisfies the requirements set by the REUNION scenario. Both architectures fulfil the internal requirements, that is: they pass a sufficiency test in that some set of parameters can be set to achieve the correct behaviour. However, both show limitations for different reasons in terms of their external evaluation. The **RAL** architecture reproduces the behaviour because of a ‘hand-coded’ local mechanism that is not general enough. The **HAR** architecture does possess general mechanisms, but these seem to be postulated too early in development. As a result of the evaluation of the **RAL** and **HAR** architectures additional structures and mechanisms are proposed in section 5.3. These detail how the **HAR** architecture might be constrained to match what is known about the neural basis of action selection; to be implemented at a lower level of granularity; and to incorporate different forms of representation.

1.8 Overview of forms of evaluation

Various types of evaluation include assessing:

1 - if the source material and constraints fairly reflect the broader literature;

2 - if the mini-scenarios produced are a fair abstraction of the selected source material, and are appropriately concrete and evaluable;

3 - how well the designs meet the scenarios;

4 - whether the software implemented designs are true representations of the linguistically described designs (and vice-versa);

5 - whether the software implemented designs are approximations of true representations of real infants;

6 - how the fitness of the designs changes following changes in design space;

7 - how the fitness of the designs changes when the scenarios are varied, perhaps beyond observed parameters

The internal evaluation is concerned with the internal consistency of each competing design. It iteratively assesses whether each part meets the requirements in relation to other parts of the overall process. The scenario must be a faithful abstraction of the source material, the design must meet the specification provided by the scenario and the implementation must possess a design that matches that already described. After each internal evaluation the implementations and designs are both, if possible, fitted to satisfy their requirements.

Evaluation of surrounding design space is undertaken to fully understand a design. We need to understand the space of possible designs (Beaudoin, 1994). Of interest are issues such as: how the design might have been different, and trade-offs implicated in the design. We want to see the effect of slight changes in the requirements and design. Approaches such as Monte Carlo simulations, computational experiments and comparative modelling may be used to evaluate the design space of simulations (Cooper, 2002):

- Monte Carlo simulations involve summarising behaviour from an analysis of multiple runs of the simulation.
- Computational experiments involve manipulating independent variables as they would be manipulated in an actual psychological experiment.
- Comparative modelling means comparing two alternative explanations

External validation is concerned with assessing the competing designs and implementations with regard to how good they are as theories of the target phenomena. Part of the external evaluation is an assessment of the source material. For example, might the source material be improved upon? Has the process of designing a simulation suggested particular observatins that would distinguish candidate theories. If possible, we would like to form our theories so that they can be rejected. According to Lakatos (1970), old (or rival) theories are only rejected when new theories are proposed that explain all the successes of the old, predict additional novel facts and some of these novel facts are empirically corroborated. Following this high standard for rejection means that complex architectural theories may need a prolonged period of development before we can reject them. This is because broad architectures possess many degrees of freedom, and ad-hoc and arbitrary corrections can be made at the implementational level. So how do we rationally validate theories if there is no reasonable likelihood of rejecting them in the near future?

Lakatos (1970) provides a solution to this problem by his reconception of static theories into dynamic Research Programmes, which change over time and are evaluated as undergoing progressive or degenerative problem shifts. In the context of this project, progressive problem shifts to simulations may involve adding detail at lower explanatory levels, such as at neurophysiological levels, that are within acceptable principled constraints. Degenerative problem shifts might involve matching the simulation to its required behaviour by ad-hoc, unprincipled hand-coding of details whose only rationale is to reproduce the target behaviour. Attempting to validate architectural theories of information processing in these terms is not new. For example Newell presents SOAR as a Lakatosian Research Programme (Newell, 1990). Lakatos separated the theoretical contents of Research Programmes into two types of information: core assumptions (which identify the theory and are not expected to change) and peripheral assumptions (which form a *"protective belt"* around the core assumptions and change in response to new empirical data).

1.9 Contributions to knowledge

The contributions to knowledge that this work has made has been structured according to five types of modelling outcome[9]:

- Using the process of computational modelling to clarify problems and solutions
- Making interdisciplinary contributions
- Creating computational models to act as sufficiency tests or proofs of concept

---

[9] Aside from the interdisciplinary contributions, the other four types of contributions made by this thesis are described as modelling strategies in Cooper (2002).

- Using computational models to analyse the a priori properties of architectural designs
- Using computational models to produce empirical predictions as part of a hypothetico-deductive scientific method.

1.9.1 Clarifying problems and solutions

This work has contributed to knowledge by proposing plausible ‘prima facie’ solutions to problems. The solutions have been in the form of autonomous agent architectures. The problems have been finding putative explanations of observed infants attachment behaviours that may require further development and testing. A major challenge has been that the project started with not only an undefined solution, but, in addition an undefined problem. At the start of this project verbally specified theories with enough detail and precision to be straightforwardly adapted into computational simulations did not exist. A major contribution of this thesis is the interpretation of Attachment Theory and observations of attachment phenomena so that they could be reconstructed within the ontology of a two-dimensional virtual world. The challenge has not been that the field of Attachment Theory is lacking in candidate theories. There are numerous theories that differ in small ways from each other and this work has made a contribution by the selection of a set of consistent ideas that have been formed into scenarios.

Part of the contribution to knowledge has been in selecting a coherent set of infant behaviours that form the six different scenarios. The scenarios come in three ‘flavours’. In chapter two scenarios capture patterns of behaviour that might be caused by the interplay of only two goals, whose interaction is treated independently of goals that are inactive in those scenarios. In chapter three the behaviours are very abstract, having been distilled from meta-studies that consider dozens of different studies. In chapter four the behaviour patterns are more complex, and try to capture a year in the life of infants that results in the Strange Situation experiment. In addition this thesis has identified additional consonant behaviours for validation and future integration. These behaviours include search behaviours. As well as constraining the architectures that have been produced with additional target behaviours this work has also identified a broad selection of constraining data, arising from physiological, anthropological and cross-species studies.

Perhaps the greatest ‘hidden’ contribution is the abstraction of behaviours that has been undertaken. A quotation attributed to Einstein states: *“everything should be as simple as possible, but not simpler.”* This is the approach that has been taken in this thesis. The very complex behaviours reported in many empirical studies have been condensed down to a very simple ontology that nonetheless captures the key aspects of carer-infant interaction. This contribution is ‘hidden’ in the sense that it arose from a process of reflection rather than a systematic and recorded process, and its simplicity belies the difficulty of its creation.

A final part of the contribution made in clarifying the research problem was the analysis of the functions of the infant behaviours described in the scenarios. This work has provided a novel synthesis of evolutionary theorising as it applies to infant attachment. Bowlby provided a comprehensive treatment of this issue in the first volume of the attachment trilogy (Bowlby, 1969). However, the review presented in this thesis updates Bowlby’s treatment and brings together ideas from various contemporary sources. In addition, the assessment of evolutionary function is dealt with in an atypical way because it considers the issue of adaptive function behaviour by behaviour in the six different scenarios.

Chapter two contributes to the debate on evolutionary function of infant attachment behaviour by highlighting a key evolutionary trade-off. This is the issue of how to stay safe and learn about the environment in preparation for the future. Bowlby suggested that the Behaviour systems for Attachment, Fear, Exploration and Socialisation were key to understanding infant attachment phenomena. Chapter three rationalises the operation of these four Behaviour Systems in terms of the previously described trade-off between safety and learning.

Each of chapters two, three and four analyse the cause of infant mortality and emphasise different risks for an infant. Chapter two emphasises external threats that increase the further the infant is away from its carer. Chapter three emphasises that infants are at greater risk when their carer is unreliable, and shows that in some circumstances increasing signalling can moderate the increased risk an infant faces. Chapter four emphasises that there are also circumstances when increasing signalling may have an adverse effect on an infant’s security by triggering a negative response from its carer.

1.9.2 Interdisciplinary contributions

This work is a quintessential example of interdisciplinary Cognitive Science research. This is because it combines data and theoretical constructs from Attachment Theory with an approach to modelling and structures and mechanisms derived from Cognitive Psychology and sub-disciplines of Artificial Intelligence, such as the study of multi-agent and autonomous agent systems (Wooldridge, 2002), reinforcement learning (Sutton and Barto, 1998) and planning (Nau *et al.*, 2004). Working with a mixture of disparate disciplines such as these can exert a type of load not found when researching safely within the confines of a single discipline:

> *"One fundamental cognitive reason why disciplines exist seems apparent - it comes from resource limitations to human cognition. There are only so many hours in the day; they can be filled with only so much talk and only so many articles. These limits assure that no scientist will be familiar with more than a small fraction of science; thus science cannot be the seemless web we all desire"* (page 105, Newell, 1983)

One of the ways that intra-disciplinary science marries the twin requirements of minimising the load of interaction across disciplines whilst maintaining channels of interdisciplinary communication is through transferring publications. As Newell notes,

> *"There is no need to examine the original work in the parent field or to follow in detail its continued advancement there. In general, only the gross aspects of an idea - to wit, the idea itself - is useful in the recipient field. Actually, a significant transform is usually necessary to make the idea applicable to the recipient field, which is accomplished by the transferring publication. Thus, there is no reason to go beyond it"* (page 108, Newell, 1983)

This work can make a strong claim to possess the potential to act as a two-way transferring publication. Infant attachment behaviour has been reviewed, selected and abstracted so that agent based modelers who are interested in psychological modellers have pre-formed scenarios to start from. Section 3.3.4 demonstrates a novel formal problem specification that makes a contribution to the field of reinforcement learning. Chapters four and five make a contribution to the study of hybrid agent architectures by demonstrating intermediate cases between reactive and fully deliberative architectures. This work also acts as a demonstration model for what modelling can do for the Attachment Theory community. Section 1.9.4 discusses how the work in this thesis that presents infant-carer interactions as dynamic feedback loops may act to change the conception within Attachment Theory of how attachment styles are formed.

In Newell's characterisation, researchers use a transferring publication so that they don't have to leave the comfortable confines of their home discipline. Transferring publications fulfil a bridging function, but developments arising from a transferring publication occur back within a receiving discipline. If research gives rise to longer term sustained efforts to collaborate between two disciplines then this work may be more accurately characterised as forming an Interfield Theory. Darden and Maull (1977) suggest that Interfield Theories

> *"bridge two areas of science ... and may provide answers to questions which arise in one field but cannot be answered within it alone, may focus attention on domain items not previously considered important, and may predict new domain items for one or both fields."* (page 43, Darden and Maull, 1977)

It may be that Attachment Theorists and agent based modellers come together to form new scenarios, developing new problems and solutions that derive from consideration of infant-carer pairs as dynamic complex systems. However, a qualification to the possible contributions that this work may make as a transferring publication or in founding an Interfield Theory is that simply by possessing the potential to act in these ways does not mean that a piece of work will fulfil its potential. This is dependent on how the broader community views and acts upon this work.

1.9.3 Sufficiency testing and proofs of concept

When problems are well specified simulations can act as sufficiency tests. Sufficiency testing *"involves implementing a pre-existing (verbally specified) theory, and determining if the theory behaves as claimed"* (page 26 (Cooper, 2002)). A simulation that implements a theory sufficiently is a proof of concept for that theory. This doesn't mean the theory is proved, merely that it is internally consistent and is a possibility. Parameter fitting may be needed to enable the simulation to sufficiently reproduce the required behaviour.

In addition to producing an implemented simulation that acts as existence proofs of particular designs, this work makes a contribution to knowledge because of its novel interpretation/translation of empirical description into scenarios. During the process of designing and implementing it may also find new mechanisms to act as part of the solution.

Attachment Theory has lacked, until this current work, an implementation of these design concepts in a working simulation. There are implementations that invoke concepts from Attachment Theory. For example Likhachev and Arkin (2000) developed a robot that possessed a comfort zone, though this zone was not defined in social terms and thus only minimally relevant to our purposes here. Other work with robots is social and simulates types of interactions that are more relevant to attachment phenomena (Breazeal and Scassellati, 2000). However, the work presented here is novel because it is attempting a much closer mapping of the behaviours observed in psychological studies of attachment. It is also the first attempt to simulate the formation of the three main attachment styles.

1.9.4 A priori analysis of theoretical properties

The simulations that have been created have been used to carry out computational experiments that have explored the design space of different architectures. The experiments have involved creating caregiving environments that it would be unethical to create deliberately in real life. For example, in chapter three Insecure infants are formed by placing them in environments where they are ignored for long periods of time.

The discovery that the simulations possessed dynamic properties that included positive feedback loops is a compelling result that arose during a priori analysis of the simulated system. Chapter three and four describe experiments where carer behaviour is varied across a range of values for sensitivity and the response of infants can be assessed in these differing environments. However, section 5.2.1 describes limitations in exploring the neighbourhood in design space of a simulation for ‘its own sake’.

1.9.5 Empirical predictions

The ‘gold standard’ for scientific contributions made by a simulation is to generate questions that either weren’t or couldn’t be asked before the construction of the simulation. There are several examples of these type of question being generated within thesis.

In chapter three this work has created two different models, (the **GLA** and **GLS** architectures) which make different predictions about the form of infant attachment learning. The results from simulations of these two architectures suggests more detailed empirical observation of real infants to ascertain if either or both of these types of mechanism exist in real infants. In addition, the results from the **GLA** architectures suggest studies which might carry out more detailed observations of infant-carer interaction to see if positive feedback loops exist which accentuate the formation of clusters of different infant attachment types. Studies of this type may provide valuable results for clinical practice.

In chapter four this work has made predictions about the effect of inconsistent caregiving versus unpleasantness in close physical contact. An adapted Strange Situation Experiment has been suggested, which has been termed the ‘Sportshall Situation Experiment’, and which would attempt to separate effects due to aversion due to close physical contact and aversion due to a history of inconsistent caregiving. In addition, this thesis has raised a number of other questions, for example: what internal processes are occurring when real infants are publicly angry or seemingly not angry but possessing a physiological state related to intense arousal.

1.9.6 Viewing the contributions to knowledge from a broader perspective

There is a way of summarising the contributions which cuts across the categories described above. In this view there are three broad areas within which this work makes a contribution to knowledge. Each of the five modelling outcomes described above demonstrates these three broader outcomes to a different degree.

The first kind of contribution is that this work extends and creates theories in the domain of infant psychology. This work makes a small contribution to some broad areas of psychological inquiry. The theories that are produced in the form of simulations are putative 'first drafts' theories that are proposed with the expectation that they will be improved. This is not to belittle their contributions because they are openings into unexplored areas of developmental psychology. The theories make a particular impact in Attachment Theory but are also applicable to the broader study of representation and information processing architectures in infancy.

The second kind of contribution to knowledge derives from the design methodology that this work introduces. It makes openings into unexplored areas of developmental psychology because of this novel methodology. What is novel about the methodology is the way that existing approaches have been integrated. A way of characterising why the methodology is original is to say that it has merged new autonomous agent techniques with some older cybernetic, ethological and representational theories combined, but never implemented before, in simulations within Attachment Theory.

Finally, as part of the design process, this work speculates on the phyologenetic origins of behaviours. This work makes a modest contribution to Evolutionary Psychology by reworking ideas from Evolutionary Psychology so that they can be used to motivate agent based simulations.

1.10 Conclusion and overview of the remainder of the thesis

This chapter started by providing some sketches of infant behaviour. The explanation of how and why these behaviours are produced constitutes the problem that this thesis is attempting to solve. The solutions to this problem are in the form of designs and implementations of autonomous agent simulations. This thesis is following a research methodology that involves three stages: setting a problem, solving a problem, and then evaluating the solution. The problem involves explaining what kind of information processing architectures may give rise to a set of infant behaviours linked to the phenomenon of attachment. The solutions are conjectural and are set within the existing framework provided by Attachment Theory. The next three chapters all follow a similar format. Each chapter frames behaviours in terms of an evolutionary function, presents architectural designs and then evaluates those designs. Where the next three chapters differ is that: chapter two is concerned with a set of separate normative behaviours; chapter three is concerned with individual differences in behaviour whose description is derived from many different studies, and are therefore very abstractly described; and chapter four is concerned with simulating individual differences in behaviour that have been derived from one study and so make more specific inferences about the causal effects of particular behaviours. Each of these three results chapters provides a putative attempt to answer or solve the problem. The solutions in each chapter are evaluated and the final chapter concludes with an attempt to evaluate the thesis as a whole.

Chapter 2

Switching goals

> *"All of us, from the cradle to the grave, are happiest when life is organised as a series of excursions, long or short, from the secure base provided by our attachment figures"* (page 62, Bowlby, 1988)

2.1 Introduction

This chapter describes the **goal-switching (GS)** architecture, which supports behaviour patterns in which the active goal that appears to be held by the infant repeatedly switches between alternatives. These cyclic patterns of behaviour are produced by infants in every day situations, such as exploring a park or meeting an unfamiliar adult. The structure of this chapter follows the methodology described in the previous chapter:

Stage 1 - **The problem is established**: behaviour in a study of infant exploration in a park is abstracted and formed into the PARK scenario.

Stage 2 - **The function of behaviour is analysed from an evolutionary perspective**: the patterns of behaviour described in this chapter are assessed as resulting from a trade-off between the requirements of learning and safety.

Stage 3 - **The scenario is extended**: The PARK scenario is extended to the WARY, COY and CONTACT scenarios by the additional requirements of the infant being able to support wariness of unfamiliar surroundings, wariness of unfamiliar agents, motivation to socialise with agents and motivation to experience close physical contact

Stage 4 - **An initial solution is proposed**: the **GS** architecture can activate goals, select the goal with the highest activation value and implement actions that have the predictable outcome of achieving that goal.

Stage 5 - **The design space of the GS architecture is investigated**: experiments are carried out that vary key parameters or change the way particular mechanisms are constructed.

Stage 6 - **The GS architecture is evaluated**: how the GS architecture supports the behaviours required by the scenarios is discussed. Behaviours that the scenarios missed out are noted, and consideration is given to how the scenario and architectural design might be deepened and broadened.

The linear nature of the stages in the list above is presented here as an organiser for the structure of this chapter. It under-represents the switching back and forth that occurred between the formation of problems, solutions and evaluation during the course of research, and is therefore partly a 'rational reconstruction' of the design process.

2.2 Framing the problem

The empirical studies selected in this chapter each describe a narrow aspect of infant behaviour. The studies have been chosen because of the restriction in the range of behaviours they illustrate. First, a study that motivates the PARK scenario, where the interaction of exploration and security in isolation from the other types of behaviour is illustrated. Next are studies which separately describe wary and coy behaviour and close physical contact. Choosing to represent infant behaviour in separate 'slices' is somewhat artificial. For example, the infants who were observed to be wary of objects were just as capable of engaging in coy behaviour. The observed behaviours are presented in isolation, but this chapter shows that the same generic architecture is capable of producing all individual patterns of behaviours. Furthermore, this chapter deals with these behaviours separately as a preliminary to chapters three and four, which integrate these behaviours in more complex scenarios.

2.2.1 Selection and abstraction of source material

When infants develop the ability to crawl, and later to walk, they can explore the world more effectively and the rate at which they acquire knowledge about the broader world accelerates. However, this new found ability to explore brings with it the potential for access to many more hazards. Anderson (1972) undertook a study of toddlers between the ages of 1 and 2 which investigated how infants balance the opportunity for exploration and the security provided by their carers. In London parks, infant and carer pairs were observed without their awareness, in observation periods averaging 15 minutes' duration. The study attempted to focus only on the exploration/security balance, therefore only cases where the pair were far enough away from interruption by other people were included in the study. Also episodes where the infant ceased exploration due to the carer attracting the infant with toys or food were excluded. This study found that most infants kept within a line of sight of their carers, and used their carers as secure bases from which to explore. It was usually at the infant's rather than carer's instigation that the infant returned to 'base'.

We want to select behaviours from this study and use them as requirements to be produced by an infant agent and a carer agent in a virtual world. In the PARK scenario this world will be also populated by a selection of static objects with varying visual attributes. The carer will remain stationary while the infant moves towards objects, motivated to explore any novel object. When the infant touches an object, and therefore learns about it, the object will be less novel and the infant will then move towards another more novel object.

The Anderson studies give more detail on the observed infant behaviour than we want to incorporate in a simulation. For example, Anderson records that infants use several different types of arm gesture during natural play. Since we are reproducing this behaviour in broad and shallow software agents these details will not be incorporated in our more abstract representation of the problem.

2.2.2 Analysis of function

The analysis of function that will be carried out for the behaviours being simulated is concerned with the function that these behaviours may have had in the human Environment of Evolutionary Adaptedness. This section will make some suppositions about the nature of the human EEA and the selective pressures that it imposed on infant behaviour. The trade-off between the requirements of learning and safety in the EEA will be examined. A suggestion will be proposed that the patterns of behaviour selected and abstracted above result from an adaptation that provides a means for infants to balance these requirements.

Researchers have drawn upon reports from contemporary hunter-gatherer societies, from comparative studies of primate behaviour, and archeological evidence, to make inferences about the nature of the human EEA (Bowlby, 1969). For our purposes there are two important aspects of the EEA that are relevant to the behaviours being simulated in this chapter:

- sources of infant mortality included predation by other species and infanticide by human newcomers into the group;

- infancy was (and still is) a period of play, exploration, and experimentation, with guidance received from the infant's mother and others who furnished the individual with the skills and knowledge needed for adult life.

2.2.2.1 Causes of infant mortality in the EEA

The EEA will have included a range of different predators, but it is likely that for infants there was a single strategy for safety that generalised across all variants of external environment. This strategy was to stay close to their main carer. In his review of research on contemporary hunter-gatherer societies, Bowlby (1969) noted that across different groups the common response to meeting a predator was for the adult male members of the group to confront and drive off the predator whilst the younger and female members of the group retired behind the males for safety. Thus the coexistence of predators and this adult behavioural strategy will have acted as a selective pressure for the adoption of mother-seeking as an adaptation. In an EEA of this type, cues to danger included: unfamiliarity, sudden change of stimulation, rapid or looming approach, heights, darkness and being alone (Cassidy, 1999). However, proximity-seeking even in the absence of any cues to danger may have evolved as an adaptation. The infant would be gaining safety as a precautionary measure, in case a danger might arise. It is this type of proximity-seeking that is apparent when animals such as ducks or geese imprint to their mothers, or sometimes imprint to humans studying them (Flavell *et al.*, 2002). Proximity-seeking by loud signalling can be seemingly dysfunctional. Infants cry when left alone, even if their carers are out of ear-shot and the crying may attract predators. Infants may possess predispositions such as this because in our evolutionary past selection has acted upon mechanisms that give rise to policies such as the general action 'cry when carer is far'. This policy would have evolved if it were more beneficial than the alternatives open to the infant (Weinfield *et al.*, 1999). In our modern environment, infants may be unaware of threats that are obvious to an adult, and conversely sometimes act as if threatened by people or objects that pose no threat. Infants can meet a number of strangers who are all totally benign in a modern environment, and yet they will still continue to be wary of strangers they meet for the first time. This is because attachment behaviours such as stranger anxiety are not learned from individual experience but are inherently motivated adaptations to the EEA.

Hrdy (1999) reviews various studies on the causes of mortality in some primates species, such as Langur Monkeys. She highlights one of the major causes as infanticide by unrelated males. These males are newcomers to groups and kill existing infants to free their mothers for immediate reproductive availability. This form of behaviour in humans may seem extreme when considered against modern social norms. However, Daly and Wilson (1988) have found a less extreme effect in modern humans, where the chances of a step-father committing murder against a step-child are far greater than the chances of of a child being murdered by its biological father. It may be that in the EEA the vulnerability of infants to adult males to whom they weren't related was intermediate between the levels of infanticide observed in Langur Monkeys and in Daly and Wilson's study of infanticide in modern humans. It certainly seems reasonable that the universally cross-cultural development of stranger-anxiety at around eight months of age in humans is an adaptation selected for in the EEA.

2.2.2.2 Requirement for infant learning in the EEA

Rochat (2004) highlights the altricial nature of human development:

> *"In mammalian evolution, there is a general trend from large litters of altricial (fast growing and underdeveloped) young to small litters of precocial (slow growing and overdeveloped) young. In this general evolutionary trend, humans represent a noticeable exception: their litter is small and their young are most definitely altricial - helpless and underdeveloped at birth. Why do humans represent such a noticeable exception in mammalian evolution? In comparison with other primates, humans are born too soon. According to some estimates, for humans to have the growth level at birth of other great ape species, their gestation time should more than double (from nine months to approximately twenty-one months).*(Page 12, Rochat, 2004)

Rochat (2004) suggests three selective pressures that may have all contributed to the altricial nature of human development. One is anatomical, it may have been the narrowing of the pelvic bone due to bipedal locomotion that provided selective pressure for human infants to emerge sooner and smaller. Another is physiological, it may be that the greater brain growth of humans requires more food than can be supplied within the womb. The final contributing factor is of particular interest in our current task of interpreting the function of infant behaviour. This factor is that it was the ability to exploit the rich extra-uterine environment that became selected for in the EEA. Human forms of higher learning and unique psychological functioning may require a more supportive and stimulating infant environment from earlier in development than other species because of the unique form of cognition that develops. It may have been that the information required to cope adequately to achieve certain sorts of goals in certain sorts of environments could not be encoded in the genome for some reason, e.g. lack of opportunity in the evolutionary history, or lack of capacity in the genome, or some other reason. It therefore had to be acquired by individuals. This raises questions regarding learning mechanisms and how they relate to kinds of brain development. The suggestion is that staggering brain development during learning is far more effective than producing a fully formed brain and then making it start learning. It may have been the requirement for more advanced adult skills which selected for adaptations that brought about greater infant predisposition to learn from novelty.

What did infants need to learn about in the EEA? Infants needed to learn how to manipulate objects in their environment and learn skills directly from their main carers. In addition, a preparation for adulthood required an understanding of social relations with other group members and strangers who may have been trading partners or competitors for resources. The theory of reciprocal altruism shows us that we shouldn't expect interaction between different groups to involve just costs. Since adult humans in the EEA possessed advanced cognitive skills we should expect them to have engaged in complex forms of reciprocal altruism. For infants the most adaptive pattern of behaviour will likely have been to attempt not only to avoid unfamiliar adults because of the threat they present, but to develop social skills to exploit the opportunities in meeting strangers. Increasing levels of social skills will have required greater cognitive ability, and over many generations the opportunities and risks of interaction with strangers may have contributed to a 'cognitive arms race' which accelerated the cognitive development of the whole species (Pinker, 1998).

2.2.2.3 A trade-off between learning and safety

In the two previous sections we have suggested that infants possess two types of adaptations relevant to the behaviour we want to simulate. These are adaptations that motivate them to experience different forms of learning, and motivate them to maintain their safety. An infant primarily maintains safety by gaining proximity to its main carer. Since these two types of adaptations can work in opposite directions there is a trade-off between them (Chisholm, 1999). Increasing learning by moving to new objects can decrease proximity and hence decrease safety. Maximising safety by remaining close to the carer does not rule out learning, but it certainly constrains learning opportunities. There is an asymmetry between learning and safety that Chisholm makes clear:

> *"in evolution (as in all complex adaptive systems) the cost of 'stepping off a cliff' is always greater than the benefit of 'setting the stage' for some future fitness move. This is because the penalty for failing to avoid a cliff is immediate and severe - death or other failure*

> *to reproduce - whereas the penalty for failing to set the stage for some profitable future move is, by definition, both delayed (making it possible that another opportunity might arise) and not so severe, being an opportunity cost (a failure to increase rather than an actual decrease)"* (page 91, Chisholm, 1999)

Attachment Theory states that the functions of learning and safety are realised by information processing mechanisms that have been termed behaviour systems (Bowlby, 1969). Fear and attachment behaviour systems keep the infant safe while exploratory and sociable behaviour systems foster learning. The fear system is most likely to be activated in compound fear situations (Cassidy, 1999). Many contexts that will elicit fear will often also activate the attachment system. Bowlby defined external threats as fear provoking[1]. Attachment threats were more narrowly defined in Bowlby's work as those that gave rise to anxiety from the absence of the attachment figure.

The function of the sociable system is to learn about others, including the main carer. A quote from Bowlby helps us distinguish the attachment behavioural system from the sociable (affiliative) behavioural system:

> "*A child seeks his attachment figure when he is tired, hungry, ill or alarmed and also when he is uncertain of that figure's whereabouts, when the attachment figure is found he wants to remain in proximity to him or her and may want also to be held or cuddled. By contrast a child seeks a play-mate when he is in good spirits and confident of the whereabouts of his attachment figure; when the playmate is found, moreover, the child wants to engage in playful interaction with him or her. If this analysis is right, the roles of attachment figure and playmate are distinct*" (page 307 Bowlby, 1969).

The attachment, fear, sociable, and exploratory systems can be viewed as possessing a complementary relationship that would be predicted to provide an optimum balance of learning and safety in the EEA, but not necessarily in other environments (Cassidy, 1999). The overall structure of the system that supports these four types of behaviour will be treated as an adaptation by this thesis. This is not to suggest that the design of the architecture was achieved as a unitary step by optimising fitness. What it allows is that the adaptive value of the overal structure emerges from a collection of distinct features of the structure which were selected

---

[1] Bowlby did not discriminate between fear and wariness, and hence in his use, the fear behaviour system should be thought of as involving fear and wariness, as these terms are used in this thesis

for their adaptive value. This thesis is concerned with using a design process to better understand the nature of the overall adaptive architecture into which these four behaviour systems fit. What is of interest is not just the structural features of which the architecture is created, but also the way that these features combine to produce adaptive benefits. The next section spells out the requirements for this architecture and following sections describe its structure in more detail.

2.2.3 Scenario formation - the PARK scenario

Anderson lists five abstract aspects of behaviour as a summary of the infant's behaviour, and these statements provide the basis of the PARK scenario for our autonomous agents. These are that the infant will:

- frequently run to and from the carer;
- mostly keep within sight or sound of the carer;
- stop for longer by the carer's side than he does when he is away from her;
- glance at the carer from a distance
- engage in physical contact of a quick 'make-and-break' nature

An interesting aspect of the exploratory behaviour in the park is that after the commencement of each episode an infant's behaviour can continue to change whilst the environment stays the same. Internal changes in the state of the infant drive the variability in the infant's behaviour, rather than external changes in the infant's perception of the park or the stranger. This quality of behavioural change without environmental change needs to be incorporated into the scenario. In addition to viewing observable behaviours as the result of particular behaviour systems, this work also views the overall structure of the information processing architecture as an adaptation. This work makes the assumption that an infant's information processing architecture is designed to balance the behaviours that follow so that infants optimise safety and exploration. This means that the infant agent should not, for example, spend all its time keeping safe to the detriment of benefiting from opportunities to learn (about its environment).

Since this work is simulating infants as complete agents, capable of online processing in a virtual world, the design challenge can be reframed. This is now expressed as wanting to explain patterns of infant behaviour that arise when infants are attempting to form coherent responses when faced with simultaneous competing goals. The competing goals are learning about different kinds of novel situations and keeping safe.

The scenario requires that the world possesses two types of agent (an infant and carer) and a single type of ‘toy’ object. Infant agents should be able to sense the positions of the carer and objects in the world, and should be capable of acting by moving towards either the carer or any of the objects. The objects are inert, in that they possess visual features of interest to the infant, but can neither sense nor act upon their environment.

The details of perception, navigation and motor control are not part of the problem being tackled in this thesis. These are all difficult problems in their own right and this work is attempting to abstract away from them to focus on issues to do with central processing. Learned fear is also not part of the scenario. The principal integration is between different goals, showing how an agent manages to integrate its exploratory behaviours with its behaviour acting towards the goal of security.

The simulation is evaluated in a graphical mode where the parameters of the simulation can be interactively varied to see that it follows the scenarios. Screenshots of the simulation running in graphical mode can be seen in figures 2.5, 2.3, and 2.6 and appendix G.1. Trace print-outs of the running scenario also give numerical measures of the behaviour, such as the distance of the infant from the carer and the infant’s orientation.

The PARK scenario was created from observations that excluded social interaction with persons unknown to the infant. Obvious extensions to the PARK scenario are to give the infant agents a wariness of unfamiliar environments and unfamiliar agents, a desire to socialise and the ability to signal to get the carer or other agents to turn or move towards them in pursuit of security or socialisation.

2.2.4 Extending the requirements: the WARY scenario

The WARY scenario incorporates the phenomenon of infants experiencing a negative emotional response and aversion to objects or an environment that they have not experienced before. The wariness that infants experience in response to unfamiliar objects should be distinguished from fear of previously challenging situations. Bronson and Pankey (1977) undertook a study on the mechanisms that mediated infant avoidance of objects and environments. Forty infants were repeatedly observed throughout their second year in a variety of mildly challenging situations. Observations of wariness were gained by observing the infant’s response when taken with their carer to an unfamiliar room where there was an unfamiliar object that produced lights and sounds. Analysis of the infant’s behaviour lead to the positing of two distinct types of avoidant mechanisms. Avoidance of the unfamiliar was termed wariness and this was contrasted with fear, which resulted from previous experience of a mildly difficult situation.

The PARK scenario already incorporates objects, so what needs to be added to incorporate wariness is to give the infant the ability to be able to familiarise itself with novel objects, so that when it is in an unfamiliar environment it can discern the change.

In the WARY scenario the infant agent will return to its carer sooner when it is placed in a new and unfamiliar environment. However once the infant is placed in a new environment it will start to familiarise itself with that environment, until wariness drops and the infant becomes familiarised.

2.2.5 Extending the requirements: the COY scenario

When an infant meets unfamiliar adults coy behaviour often results. Marvin and Britner (1999) describes a pattern of coy behaviour that can occur when an infant meets an unfamiliar adult.

> *"If the wary system is highly activated, the infant tends to retreat to the parent as a haven of safety; if it is not, the infant may continue to stare at a non-intrusive stranger, or may initiate or respond sociably. In many cases, one can see a cycling of conflicting behaviour systems, with the infant moving back and forth from parent to stranger as the distance from each tends to activate one behaviour system and terminate the other."*(page 54, Marvin and Britner (1999)).

A scenario extension that includes coy behaviour provides a complimentary focus to Anderson's Park study, by minimising the exploration of the inanimate and concentrating on social learning. In the COY scenario the infant will exhibit the alternation between approach and avoidance that characterises the behaviours of the infants in Marvin and Britner's study. When an infant agent is with its carer and a stranger approaches the infant will:

- turn to and from the carer and stranger;
- keep close to the carer, but make movements towards the stranger and then retreat back to the carer;

This summary of the COY scenario shows a basic structural similarity with the PARK scenario. The differences are related to the substitution of sociability for object exploration. Extending the PARK scenario with COY behaviour can therefore be thought of as an investigation into the niche space of the PARK scenario. Instead of interacting with inert objects the infant agent interacts with a stranger agent, which the infant will be simultaneously motivated to socialise with and be wary of.

2.2.6 Extending the requirements: the CONTACT scenario

The complex and vital role of physical contact in the development of attachment has been highlighted by studies of attachment in rats (Hofer, 1995; Polan and Hofer, 1999). Close physical contact has been shown in rats to act as a hidden regulator of an infant rat's behaviour. When rat pups are separated from their mothers they showed slowed heart rate and decreased temperature and activity. However if the rats were kept warm their temperature still fell but they became much more active. After a series of systematic experiments it was found that different stimuli from the mother rat, such as tactile, nutritional and thermal stimuli, regulate different infant behavioural and physiological systems (Goldberg, 2000). Hofer (1995) described the mechanisms he had discovered as hidden regulators. It may be that similar hidden regulators form the physiological basis for the affective state of security that human infants experience when they gain close proximity with their carers.

The CONTACT scenario does not incorporate the physiological details of the studies of close contact and attachment in rats. The principal reason for this decision is that the low level of detail required to reproduce the details of close contact shown by these animal studies does not match with the high level of description of information processing constructs found in the other scenarios. In addition the physiological details of attachment regulation in rats may be somewhat similar to that in humans, but they are not identical. Therefore an initial simulation of hidden regulation based upon Hofer's findings might be best carried out with simulated rats rather than simulated humans. What the CONTACT scenario does include is the requirement that infants spend time in close contact with their carers. The physiological details remain for future work.

In the CONTACT scenario the infant is required to:

- periodically return to the carer for close physical contact.

This requirement may seem very simple, but it is a compromise between having no close physical contact and a more ambitious requirement that may have specified physiological details.

2.2.7 Integrating all the scenarios

The full set of requirements is that a single architecture can produce all four of the scenarios described above.

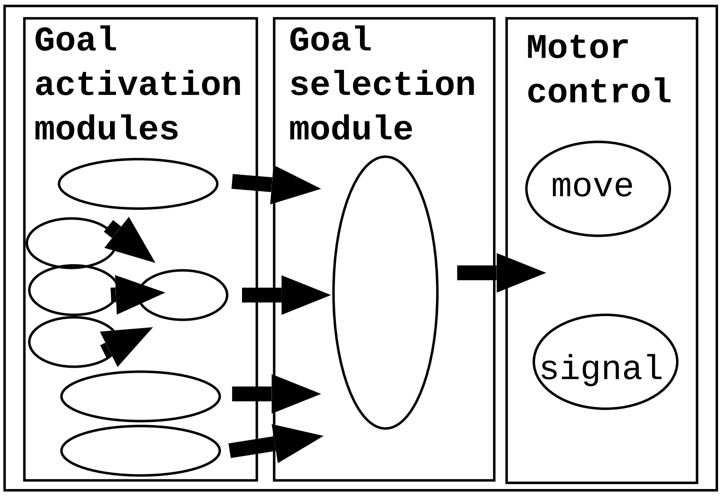

Figure 2.1: The **GS** architecture with activators for exploration, security, socialisation and physical need. The security goal activator includes components for anxiety, and object and agent wariness. Only the most highly activated goal is selected. When a goal is selected it directs the motor actions of the agent.

2.3 The GS architecture

2.3.1 Overview

The PARK scenario requires an architecture that supports regular and repeated interchanging of behaviours, alternating between attempts to carry out exploratory or social learning, and then swapping to seek the proximity of the carer. The **goal-switching** (**GS**) architecture, shown in figure 2.1, produces these required patterns of behaviour. The **GS** architecture has three major sub-divisions: a perceptual subsystem comprising goal activator modules, a central selection and arbitration subsystem, and an action subsystem. For each slice of time, each agent senses its environment. This sense data enters each goal activation system in a simulated parallel manner. Then each goal activation system outputs an activation level within the same time slice, and therefore simulates a fast, continuous process. In the current implementation the goals that are activated are explicit. This might not be the case in human infant systems.

There are four goal activator modules. Each can be viewed as possessing an implicit goal or function, and each goal provides proposals for action with a variable activation level. They include: the security module (maintaining safety from unfamiliar objects, unfamiliar people and remoteness of the carer); the exploration module (learning about objects); the socialisation module (learning about agents) and the close physical contact module.

In the simulated environment there are no actual threats that justify the infant activating a goal due to fear or wariness. This is often also the case for real infants in modern environments. The functions of goal activators in real infants can be viewed as adaptations that have been selected for in the EEA, they do not have to be adaptive in the current environment. This is the type of function that the simulation is trying to capture. The goal activators in the simulation are intended to act towards functions that would have been required in the EEA. Each goal activator module carries out three subtasks. The first is to gain an activation level as a result of sensing the state of the environment. To do this, each goal activator receives sensor readings that include the positions and visual features for all the agents and objects within the infant's sensor range. The second subtask is to modulate this activation according to internal factors. These include comparing the sensor readings to patterns of previously perceived objects and agents that are held in memory, and adding residual levels of activation from previous cycles. The third subtask is to make selections for targeting. For example, the exploration goal activator passes the bearing of the particular target toy that it has selected to explore, as well as a scalar activation value.

There are also two action subsystems, moving and signalling. These can be active simultaneously. In between the perceptual subsystem and the action subsystem is the selection and arbitration mechanism, which selects the action with the highest activation in a 'winner takes all' contest. Each of the goal activators may have a substantial level of activation, but only one goal activator, with the highest activation, defines the next actions to be carried out.

2.3.1.1 The exploration and socialisation goal activators

The exploration and socialisation goal activators in the simulation are intended to possess an evolutionary function similar to that which is held by part of the information processing architectures of real infants. This function is to prepare the infant for the future. The exploration and goal activators also possess a similar internal structure. Both goal activators sense the environment and pass increasing activation when a target of exploration or socialisation is close by. Activation for exploration and socialisation rises over time even if no objects or agents are within sensory range. This is because the exploration and socialisation goal activators increment their activation levels by a small amount each time slice in the simulation. One way to think about this property of the goal activators in the agents is to relate it to real infants possessing a kind of innate motivation for learning about objects and agents. The longer infant agents don't experience novel object or agent interactions the higher the activation of the goal for this kind of experience becomes. This state of affairs is hand-coded in the current simulation. However, in real infants the same behaviour may emerge from perceptual systems that amplify signals from novel sources, or habituate to familiar sources.

The exploration and socialisation goal activators focus on particular objects or agents. Thus if there are multiple objects or agents present only the attributes of one of each will be used in setting the amplitude of the respective goal. A real infant may choose to move towards a group of objects that are close together, and only choose to interact with a single object when it is closer. The exploration and socialisation goal activators cannot perceive groups of objects together.

2.3.1.2 The security goal activator

The security goal activator in the simulation has three components, that each provide part of its activation. The internal workings of each of these security components operate differently from the exploration and socialisation goal activators. The learning goal activators rise over time, whereas the security components don't. In real infants the activation of attachment systems is likely to change over time even when the carer is close. For example, infants are more likely to seek proximity when they are fatigued or hungry.

2.3.1.2.1 The anxiety component of the security goal activator The anxiety component of the security goal activator is intended to possess an equivalent function to the protective function needed by infants in the human EEA when they became isolated from their carers. Figure 2.2 illustrates how exploratory and attachment goals interact in the PARK scenario (figure 2.5 gives a more complete description of the graphical display of the virtual world).

The anxiety component of the security goal activator may pass no activation for extended periods of time. This occurs if the carer remains very close to the infant. Activation only starts to be passed when the agent senses that the carer is beyond a Safe-range distance.

At the moment when an infant moves beyond the Safe-range distance the activation level of the anxiety module is likely to be below the activation level of other modules, particularly if novel and unthreatening objects are present and driving the goal of exploration higher. The anxiety module passes increasing activation the longer the carer is beyond the Safe-range distance. If a carer does not go far and returns quickly the anxiety goal may not have risen high enough to win-out in the selection mechanism to cause behaviour to switch. If a carer remains beyond the safe range distance for a long period then the activation of the

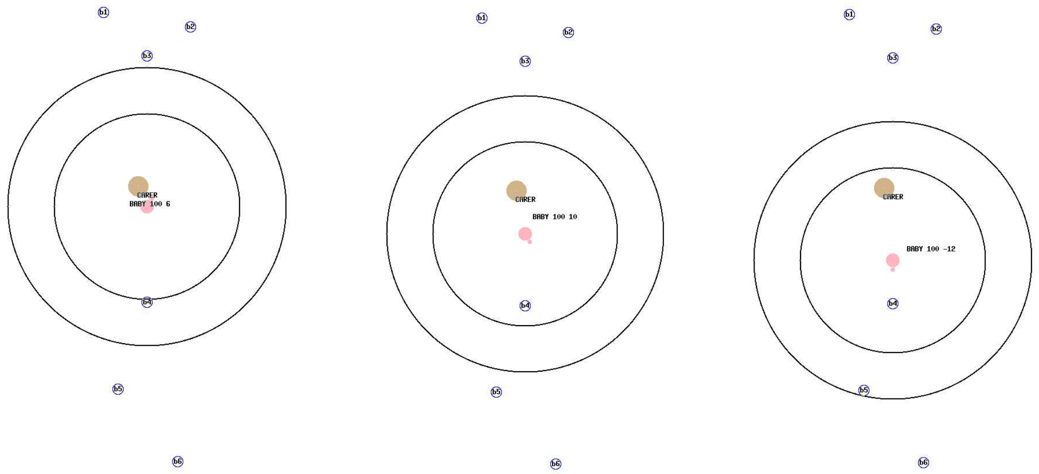

(a) Exploration is active, anxiety has zero activation

(b) Exploration is active, anxiety has zero activation

(c) Exploration is active, anxiety has zero activation

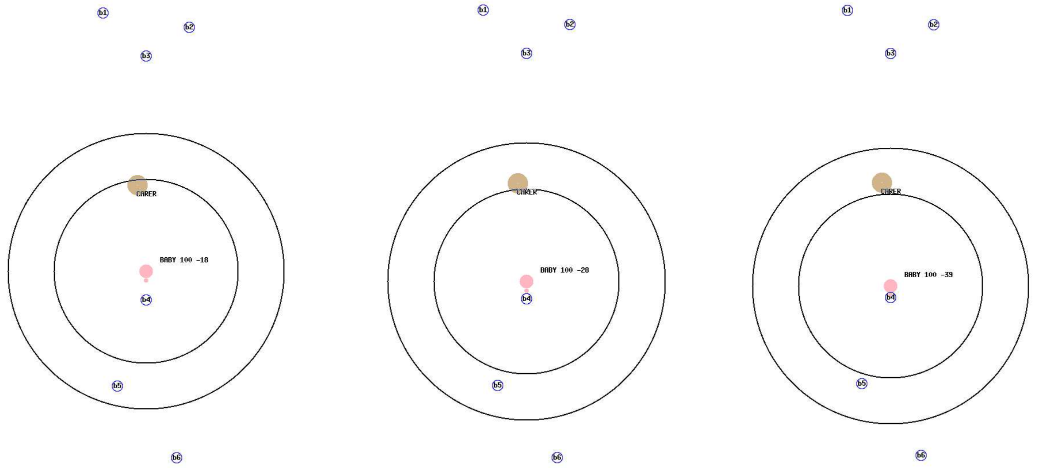

(d) Exploration is active, anxiety is increasing in activation

(e) Exploration is active, anxiety is increasing in activation

(f) Exploration is active, anxiety is increasing in activation

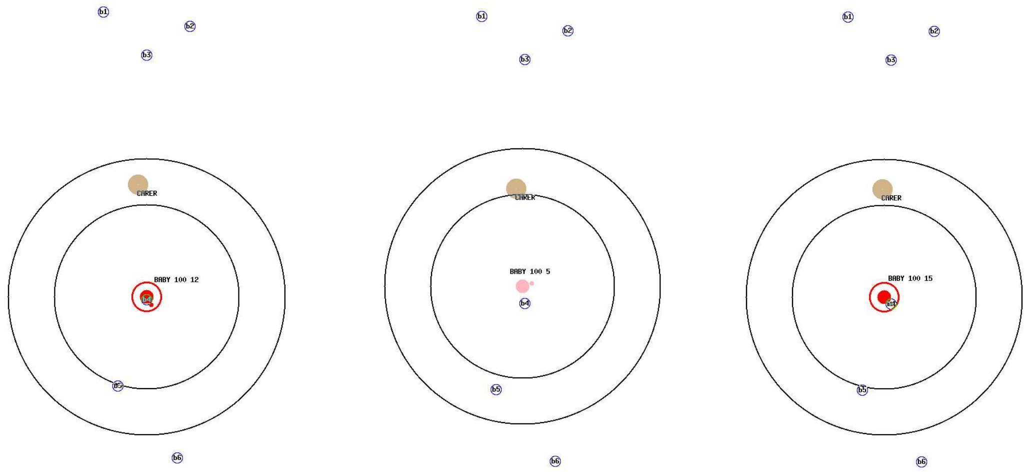

(g) Attachment is active, exploration is above zero activation

(h) Exploration is active, anxiety is increasing in activation

(i) Attachment is active, exploration is above zero activation

Figure 2.2: Nine snapshots of the an infant agent exploring during the PARK scenario, covering a duration of fifty six cycles of the simulation. The internal state of the infant at each of the nine moments in time is detailed in G.1

anxiety module will win-out over other goals and the infant will act to decrease the distance to the carer. This can be accomplished by moving towards the carer and by signalling so that the carer moves towards the infant.

After the security goal has been activated and the distance to the carer has become less than the Safe-range distance, then no new activation is passed as output. If the carer is closer than the Safe-range distance and the carer is signalling, the anxiety level drops to zero.

2.3.1.2.2 The object wariness component of the security goal activator The object wariness component of the security goal activator is intended to possess an equivalent to the protective function needed by infants in the human EEA when faced with possible hazards posed by unknown objects that may signify dangers such as predators. The overall measure of wariness has two components: wariness from objects and wariness from agents. Object wariness is not wariness of particular objects. The form of object wariness that is supported in the simulation is wariness of novel arrangements of objects. The object wariness component of the security goal activator stores an abstraction of the way that the objects are arranged. As it experiences new arrangements of objects the abstraction of these is also stored. When an infant agent experiences a new environment it compares the current pattern in its perceptual input against patterns stored from perceptions of previously experienced environments. It then outputs an index of the new pattern's familiarity. This is then summed with the other components of the security goal to make a total secure goal activation.

2.3.1.2.3 The agent wariness component of the security goal activator The agent wariness component of the security goal activator has the evolutionary function of protecting the infant against the possible hazards posed by unknown adults. This component receives sense data for all agents that are within visual range. For each agent there is data for the distance, the bearing and a description of the visual features of each agent. The carer's features are imprinted early in the running of the simulation and then held in memory. When the infant meets an agent it compares the new agent's features to the carer's features held in the infant's memory. Since the visual features are recorded as a string of numbers, the unfamiliarity of a stranger is a distance in multi-dimensional space. The raw difference is then scaled to give a figure which is the output from this component.

The outputs from the object wariness, agent wariness and anxiety components of the security goal activator are simply added together. Alternatively, they might be multiplied or combined according to some more complex function. The way in which this occurs in real infants is an empirical question.

(a) The infant's starting position was close to the carer, and it has moved between the carer and stranger a few times

(b) The same episode as figure 2.3(a) after a longer duration.

Figure 2.3: Illustrating the COY SCENARIO (implemented without agent wariness). To represent the repeated approach and retreat of the infant agent all the agents in the world are moved slightly to the right each timeslice, else the markers showing the movement of the infant would become overlayed.

Figure 2.3 shows a 'still' screen shot of the progress of an infant agent in the process of moving towards the stranger agent when the socialise goal is activated, and moving towards the carer agent when the security goal is activated. In the computational experiment which gave rise to the results shown in figure 2.3 the anxiety component was the only active component of the security goal activator. During this implementation the carer and stranger were set to move slowly in a right-wards direction so that the trail left by the infant did not become superimposed upon itself.

2.3.1.3 The close physical contact goal activator

This goal activator is more abstract than the other goal activators. It represents an amalgamation of all types of physical need, from feeding to warmth and close body contact itself. It is meant to represent the findings from studies such as Harlow (1958) and Polan and Hofer (1999), but it does so only in a very shallow manner. The chief benefit of this goal activator being present in the **GS** architecture is seen in chapter six, where its operation means that even when infants become averse to close physical contact, they are still motivated to achieve it.

2.3.1.4 The selection and arbitration subsystem

The selection and arbitration subsystem takes inputs from each goal activator and outputs the winning goal. The activation levels of the winning goal activator is the output of this system. The activation level of the winning goal is used to determine the broadcast level of the signalling action. The selection mechanism in the **GS** architecture is implemented very simply and rather implausibly. One way to improve the model of actions selection would be to base it more solidly on neurally inspired mechanisms. There are a number of different but related neural models of the Basal Ganglia that are theorised as performing a similar action selection function in humans and other animals (Redgrave *et al.*, 1999; Prescott *et al.*, 1999; Gurney *et al.*, 2001a,b; Montes-Gonzalez *et al.*, 2001; Girard *et al.*, 2002; Prescott *et al.*, 2002), (Frank, 2005), (Houk *et al.*, 2005). Basing the action selection that occurs in the infant agents upon the Basal Ganglia is also supported by the fact that this organ may be the the locality for part of the process of Contention Scheduling (Rumiati *et al.*, 2001).

2.3.1.5 The action subsystem

The infant agent possesses two possible actions by which it can manipulate the world: moving and signalling. The level of signalling is set by the level of activation of the goal which activates the signalling. In an environment with infant, carer and stranger agents and toy objects the different atomic actions that are possible are three movement actions and three signalling actions: to move to the carer (Move-carer); to signal to the carer (Signal-carer); to move to a target toy object (Move-toy); to not signal (Not-signal); to move to the stranger (Move-stranger); to signal to the stranger (Signal-stranger). When the Explore subsystem has the highest activation then the Move-toy and Not-signal actions are taken. When the Security subsystem has the highest activation then the Move-carer and Signal-carer actions are taken. When the Socialisation subsystem has the highest activation then either the Move-carer and Signal-carer or Move-stranger and Signal-stranger actions are taken, depending on which agent is the target of the socialisation goal. Figure 2.4 summarises this information, listing the perceptual subsystems alongside the actions that each subsystem may activate.

The **GS** architecture is a comparatively simple architecture, but it is dissimilar to many other simple behaviour based architectures because it activates its constituent goals rather than atomic behaviours (Wooldridge, 2002). Therefore the most similar architecture in the literature is probably the A similar but more

Security	**Move-carer** **Signal-carer**
Socialisation	**Move-carer** **Signal-carer** **Move-stranger** **Signal-stranger**
Exploration	**Move-toy** **Not-signal**
Close physical contact	**Move-carer** **Signal-carer**

Figure 2.4: The mapping from goal activator modules to the actions that are activated by those modules.

complex architecture is the Contention Scheduling mechanism of Cooper and Shallice (2000). Appendix B.1 shows in some detail that the **GS** architecture can be viewed as a simplified version of a Contention Scheduling mechanism Cooper and Shallice (2000). In summary, the structure of the GS architecture can be characterised as a flat, non-hierarchical version of the model of Contention Scheduling. Key differences are also apparent in the nature of the goal activators in the two designs, with the goal activators in the GS architecture possessing more heterogenous internal structure. The model of Contention Scheduling is more complex than the **GS** architecture, but human adults are more complicated than one year old infants. The kinds of intentional tasks performed by one year old infants do not possess the hierarchical nature of the tasks modelled by the Contention Scheduling subsystem. Therefore both the **GS** architecture and the model of Contention Scheduling possess a level of complexity appropriate to the phenomena that they are trying to represent.

2.3.2 Implemented with the Simagent toolkit

This architecture has been implemented in autonomous agents using the simagent toolkit. Figure 2.5 shows a screen shot of the simulation in graphical mode. The activity in the main window shows the BABY agent moving away from the CARER agent towards an object. Of the three windows on the left hand of the screen the top window is a control panel, which allows different types of trace printing to be sent to the output window. The output window is below and partially hidden by the control panel. The window at the bottom of the left hand side of the screen represent elements of the internal state of the BABY agent. In

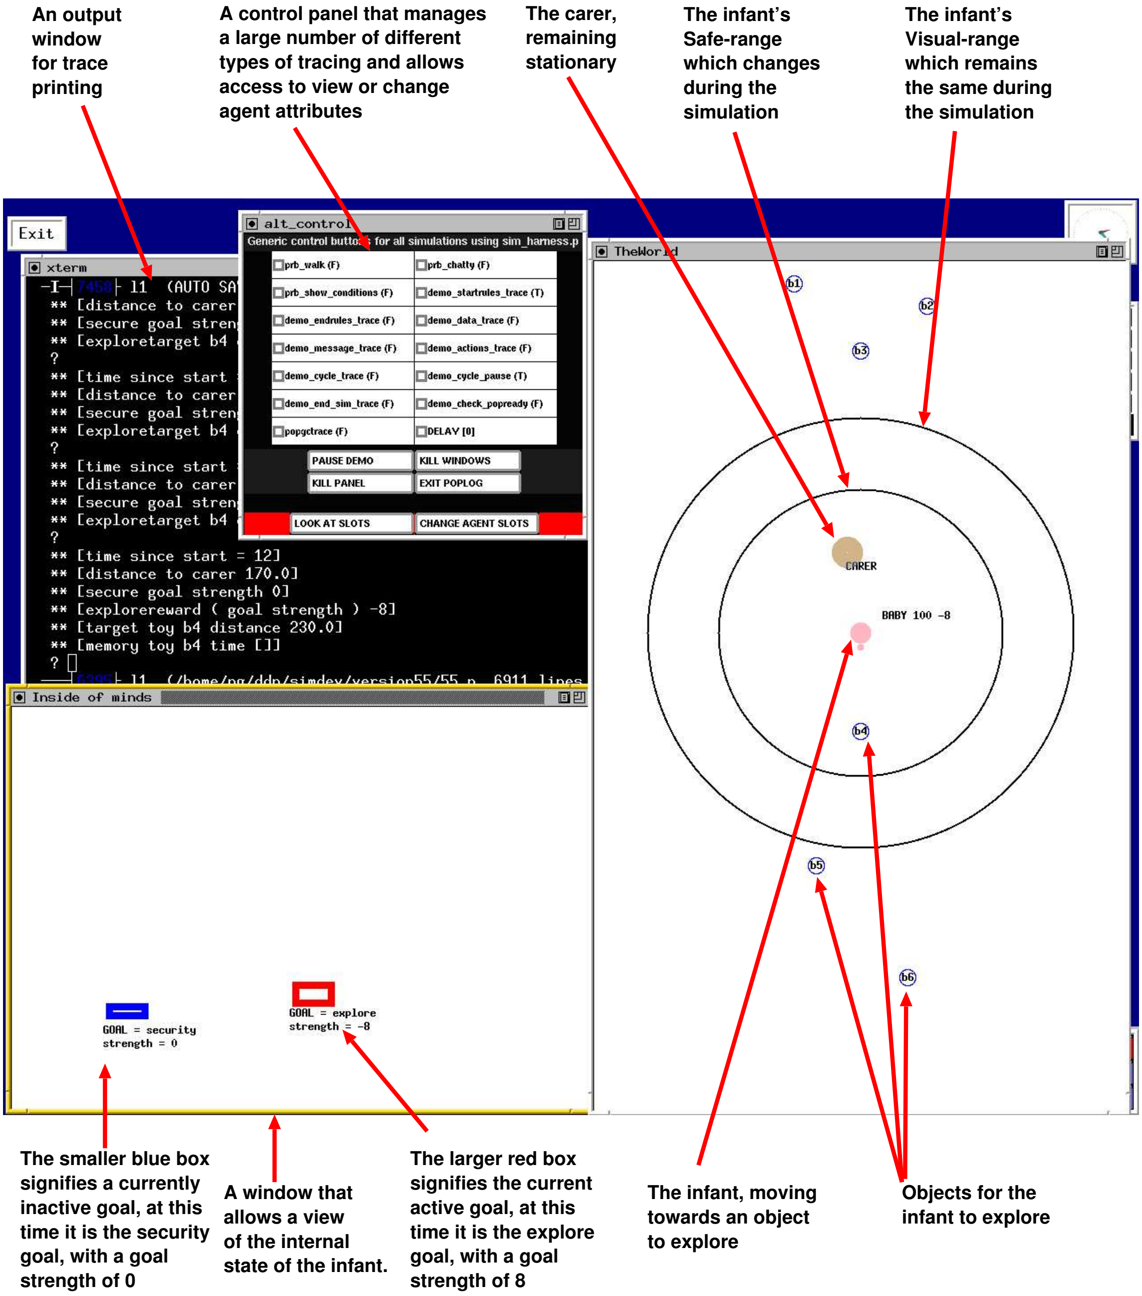

```
-I-| 7458 |- 11  (AUTO SA
  ** [distance to carer
  ** [secure goal stren
  ** [exploretarget b4
  ?
  ** [time since start
  ** [distance to carer
  ** [secure goal stren
  ** [exploretarget b4
  ?
  ** [time since start
  ** [distance to carer
  ** [secure goal stren
  ** [exploretarget b4
  ?
  ** [time since start = 12]
  ** [distance to carer 170.0]
  ** [secure goal strength 0]
  ** [explorereward ( goal strength ) -8]
  ** [target toy b4 distance 230.0]
  ** [memory toy b4 time []]
  ? 
----| 6395 |- 11  (/home/pg/ddp/simdev/version55/55.p  6911 lines
```

An output window for trace printing

A control panel that manages a large number of different types of tracing and allows access to view or change agent attributes

The carer, remaining stationary

The infant's Safe-range which changes during the simulation

The infant's Visual-range which remains the same during the simulation

The smaller blue box signifies a currently inactive goal, at this time it is the security goal, with a goal strength of 0

A window that allows a view of the internal state of the infant.

The larger red box signifies the current active goal, at this time it is the explore goal, with a goal strength of 8

The infant, moving towards an object to explore

Objects for the infant to explore

Figure 2.5: A screen shot of the simulation in graphical mode. The large window down the right hand side of the screen is the main window, and shows the positions of all the agents and objects present in the 2D virtual world in this experiment. These are a BABY, a CARER and six objects named b1 through to b6 which represent objects that the BABY agent can explore.

this particular experiment the only goals that the BABY agent can experience are for security and exploration, which conforms to the PARK scenario without its WARY, COY and CONTACT scenario extensions.

2.4 Investigating design space

2.4.1 Varying rates of anxiety change

Figure 2.6 shows screen-shots from four runs of the **GS** architecture in a simple simulated world. The larger brown object made up of two filled-in circles is the carer agent. The smaller brown circle shows this agents direction of gaze. In all four screen-shots it is facing the infant agent.

Figure 2.6(a) shows the virtual world set up with all toys nearer than its Safe-range distance. Therefore the infant can visit all the toys without ever having to pass the boundary of the Safe-range distance. This means that if the simulation was extended in time the infant would only return to the carer to satisfy it needs for close physical contact.

In figure 2.6(b) the toys have been moved further away. Now the infant has to go across the threshold of the Safe-range distance to reach the toys. The secure goal does not become activated immediately when the infant crosses the Safe-range distance. The secure goal receives small increases in activation each time-slice that it is outside the Safe-range. When the secure goal wins in the winner take all selection mechanism, the infant returns directly to the carer. When the infant has crossed back within the Safe-range distance the secure goal starts to drop in activation. When the activation level has dropped below the level of the exploration goal the infant turns and moves towards the next exploratory target. The parameters in this run of the simulation have been set so that the security goal activation drops relatively fast. This means that the infant never reaches the carer before the security goal ceases being active.

Figure 2.6(c) shows what happens when the rate of decrease of the secure goal is decreased. When this happens the secure goal stays higher for longer and the infant moves further back towards the carer agent before its active goal switches to exploration.

Figure 2.6(d) demonstrates that toys can be positioned so that the infant never reaches them. The rate of increase of the secure goal means that it will always become the highest goal before the infant has covered the distance from the edge of the Safe-range. Increasing all of the object's novelty would overcome an increase in anxiety, so that infants would move further to explore more interesting objects. Increasing the speed of the infant agent's physical movement would

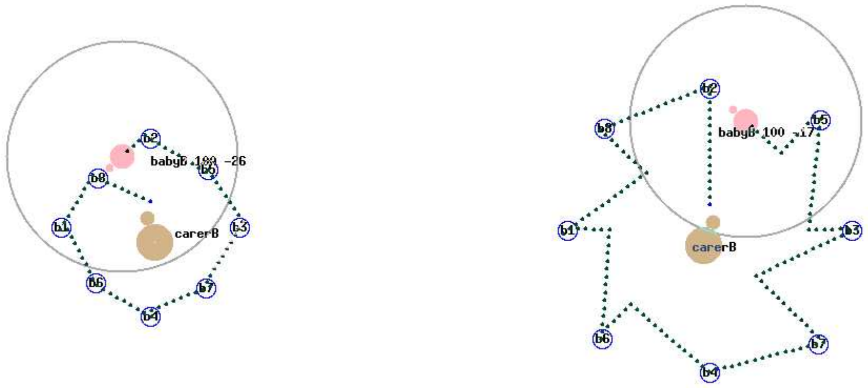

(a) The infant moves to explore each toy in turn, never moves outside of the Safe-range, so the secure goal gets no activation

(b) Infant goes out of its Safe-range, then returns to the carer until its secure goal activation has dropped below the level of the explore goal

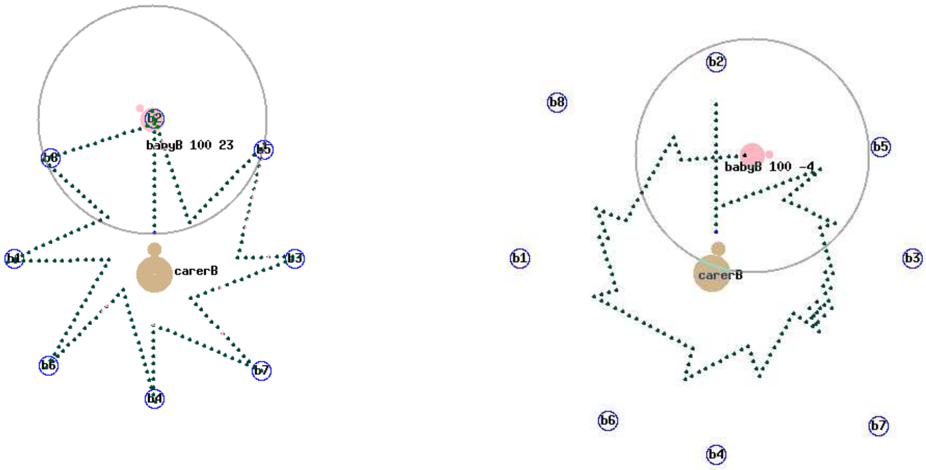

(c) The infant agents return closer to the carer because a parameter is set so that the security goal activation drops more slowly

(d) The objects have been moved further out whilst the Safe-range is maintained at the same level so that infants don't reach the objects before their security goal is activated. If the objects caused a greater activation of the explore goal or the secure goal rose less quickly then the infant would get close to the objects before breaking off exploration

Figure 2.6: Different initial configurations of infant, carer and novel objects in the PARK scenario.

produce a similar effect.

An initial response to the patterns traced by the infant agent's movement may be that they look too geometric, with straight lines and sharp turns, and not like the erratic wandering one would expect a real human infant to perform. The infant agents keeps to straight lines and makes sharp turns because of the impoverished nature of the virtual world, and the absolute concentration with which it follows its active goals within this world. Simon's 'Parable of the Ant' demonstrates an opposite phenomena (Simon, 1969). Simon sketches the path of an ant across a beach: *"It is a sequence of irregular, angular segments - not quite a random walk, for it has an underlying sense of direction, of aiming towards a goal."* (page 51 Simon, 1969). Simon notes that the ant's path is complex, but its complexity is really the complexity of the surface of the beach, not a complexity in the ant. In this simulation, if the infant's virtual world was filled with many more objects, such as different types of grass, bushes and trees in the virtual park, then there would be many more distractors to grab the infant's attention. The goal activators would need to be adapted to cope with the large number of objects that competed for attention. In this world of greater distraction the trace of infant movement may look much more erratic. In the current **GS** architecture the infant's attentional system does allow the goals of exploration, socialisation and security to interrupt each other. What this current implementation does not allow is for the perception of non-targeted toy-like objects to interrupt progress towards the toy that is the current target of exploration. If many more objects were placed in the virtual world and interrupts could occur between exploratory targets then the path traced by the infants would be less simple.

2.5 Evaluation

2.5.1 Deficiencies in the scenarios

As was mentioned at the start of this chapter, the behaviours that have been formed into scenarios in this chapter have been done so separately. This deficiency is remedied in the next two chapters where integrated scenarios are formed. There also exist other limitations of the scenarios. A major limitation of the scenarios where infants activate the goal of security is that in reality an infant considers more factors than just proximity when it is assessing its safety. Fatigue, illness or hunger give rise to an increased predisposition for real infants to carry out attachment behaviours. These types of mechanism are simply absent in the simulation. They might be added in two ways. The first way in that they might be added as new goal activators to operate alongside anxiety and wariness goal activators. The second might be that they modulate the activity of other goal activators.

Infants in reality are virtually helpless and require their carers to fulfil most of their needs. Therefore an apparent deficiency of the scenarios presented in this chapter is that the carers are not required to do much. However, the omission of different caregiving styles means that the infants have not needed to demonstrate any individual differences in their behavioural predispositions. Having to only produce normative behaviours have meant that the the **GS** architecture has not needed learning mechanisms in core areas of the design. The absence of differences in caregiving style is remedied in the next two chapters. The carer has an important role to play in the SECURE-BASE and REUNION scenarios that can be performed with varying levels of effectiveness. Consequently the infant agents in the SECURE-BASE and REUNION scenarios have to possess the means for differentiating different types of caregiving if they are to capture the adaptiveness of real human infants.

2.5.2 Deficiencies in the design

The overall structure of the **GS** architecture is intended to act as a theory that explains elements of the behaviour of human infants. However, not every aspect of the simulation is intended as a precise match of what actually occurs within the architectures of real infants. Cooper and Shallice (2000) note in the description of the model of Contention Scheduling:

> *"Some elaborations are clearly implementational, in that they are necessary for the complete specification of an executable computational model"* (page 304, Cooper and Shallice, 2000)

A major issue in evaluating the **GS** architecture is demarcating the aspects of the simulation that are of theoretical importance from those aspects that are present so that the simulation can be compiled and run.

Cooper and Shallice (2000) approach the problem of how to evaluate their model of Contention Scheduling by dividing the assumptions made in the simulation into three categories:

> *"we attempt to isolate all of the assumptions of the computational model and indicate their status with respect to the theory/implementation distinction. Theoretical assumptions are further divided into core and peripheral assumptions (corresponding respectively to Lakatos's (1970) hard-core and protective belt). We hold a strong theoretical commitment to core assumptions (CA), and believe them to be critical*

> *in producing the simulated behaviour. We hold less theoretical commitment to peripheral assumptions (PA), and regard them as being somewhat flexible and open to modification. In general, alternate theoretically acceptable mechanisms exist for the processes which these assumptions define, ... We have no theoretical commitment to implementational assumptions (IA), and are intended only for computational completeness, and hold that the behavioural characteristics of interest are largely independent of these assumptions"* (page 304 Cooper and Shallice, 2000)

The evaluation of the **GS** architecture will involve assessing deficiencies according to whether they have implications for implementational, peripheral or core assumptions.

2.5.2.1 Deficiencies of implementational assumptions

The deficiencies of implementational assumptions of the architecture include:

- the representation of perception and action. The infant agent's only means of sensing the environment is a form of vision. It is an unrealistic form of vision because although the infant can only focus on a narrow field of view directly ahead, it can see all around itself, including directly behind.
- representation of familiarity in objects and agents. The simulation uses an entirely artificial and implausible metric for measuring familiarity. Despite this drawbacks the metric does possess the benefit that it is easy for users of the simulation to understand and manipulate and is invariant to the direction that infants view objects.
- the learning that occurs in object exploration and socialisation. Learning occurs by the infant agent simply attending to unfamiliar objects or agents.
- the architecture is not represented at a consistent depth. Some of the components of the architecture have a 'black-box' quality, where the internal mechanisms and structures are intended to remain informationally encapsulated.
- the **GS** architecture is not intended to represent symbolic computation. For example, the processes of goal activation are likely to be parallel and distributed in real infants. If real infants possess four distinct types of goal that correspond to the security, exploration, socialisation and physical interaction goals represented in the final version of the **GS** architecture, these states may emerge as attractor states in a dynamic system.

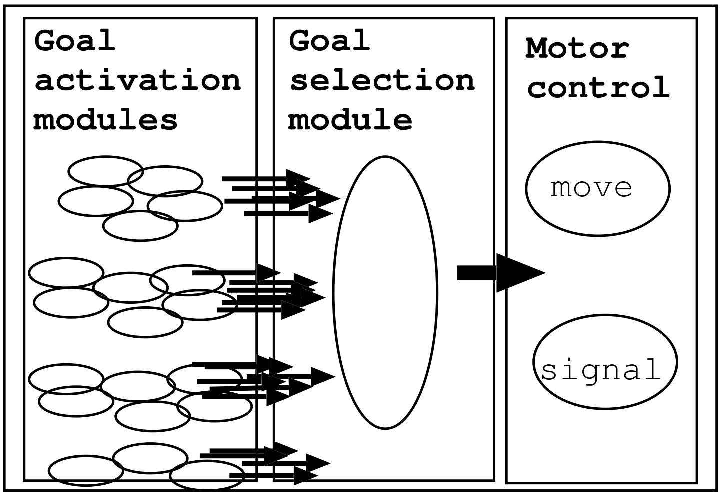

Figure 2.7: A speculative view of the **GS** architecture with multiple goal activators. This architecture has not been implemented but may be more similar to what exists in reality.

2.5.2.2 Deficiencies of peripheral assumptions

Deficiencies in the peripheral assumptions of the architecture includes the number of goal activators and how they interact. The number of goal activators is arbitrary. The **GS** architecture is described as having four goal activators, but we could equally have added the physical-need goal activator as another component of the secure goal activator or have broken the security goal subsystem down and labelled its three components as goal activators in their own right. To fulfil the PARK, WARY, COY and CONTACT scenarios requires four goal activators, but with one of the goal activators possessing heterogenous components. In reality there are likely to be multiple goals with multiple components that act together or compete with each other depending on dependencies such as what resources the goals call upon in the action subsystem. We might extend the simulation to include dozens of different types of fear, wariness or learning, as shown in figure 2.7. The important issue is how would these changes affect the theories being presented. Currently there is also no mutual inhibition between goal activators and the components of the security goal activator have been put together because they would invoke the same atomic behaviours.

A last deficiency of the design is not so much a limitation as a spur to deepening the design. The current high level representation doesn't use all possible constraints. However the design might be deepened to include lower level details. This might involve a completely different type of implementation, perhaps using low level neural processing.

2.5.2.3 Deficiencies of core assumptions

The core of the theory which the **GS** architecture instantiates is that infants of the target age range of nine to twelve months of age choose their current goal using reactive and encapsulated processing. However there exists evidence that this is not always the case with infants of this age. Tomasello (1999b) describes the emergence of joint attention between infants and their carers:

> *"Six month old infants interact dyadically with objects, grasping and manipulating them, and they interact dyadically with other people, expressing emotions back and forth in a turn-taking sequence. If people are around when they are manipulating objects they mostly ignore them. If objects are around when they are interacting with people, they mostly ignore them. But at around nine to twelve months of age a new set of behaviours begins to emerge that are not dyadic, like these early behaviours, but are triadic in the sense that they involve a coordination of their interactions with objects and people, resulting in a referential triangle of child, adult, and the object or event to which they share attention. Most often the term joint attention has been used to characterise this whole complex of social skills and interactions. Most prototypically, it is at this age that infants for the first time begin to flexibly and reliably look where adults are looking (gaze following), to engage with them in relatively extended bouts of social interaction mediated by an object (joint engagement), to use adults as social reference points (social referencing), and to act on objects in the way adults are acting on them (imitative learning). In short, it is at this age that infants "tune in " to the attention and behaviour of adults towards outside entities."* (page 62, Tomasello, 1999b).

The **GS** architecture could not accomplish behaviours such as those described above. With the additions of a number of specialised mechanisms it may be able to do so. It may be that infant social engagement starts with the infant being directed by an innate attraction to the human face and the direction of human gaze. However, it is clear that mechanisms that do not posit the infant having a sense of agency do not cover many of behaviours that appear after nine months of age. It is therefore likely that the **GS** architecture needs to be supplemented with some form of subsystem that can modulate the performance of the goal activators.

Any mechanisms that are added to the **GS** architecture to modulate how it performs need to distinguish between modulating behaviours to support the attainment of existing goals, and modulation that changes the goals that the infant will possess.

2.6 Conclusion

This chapter has selected studies of infant behaviour that deal with a set of behaviours with limited over-lap. Each study has been translated into a scenario where the infant behaviours are interpreted as arising from behaviour systems that activate a small number of goals. The common factor in all the scenarios is that infants use their carers as secure-bases. In the PARK scenario the goals that are activated are the goals of attachment and exploration, where attachment is treated narrowly as just comprising of anxiety. In the WARY scenario exploration is also activated, as is a broader representation of attachment which include anxiety and wariness of objects. The COY scenario incorporates social wariness into the goal of attachment, and this goal competes with social exploration. Finally, the CONTACT scenario adds a more abstract goal of physical contact that amalgamates needs which range from feeding to warmth. Each of the patterns of behaviour in all the scenarios has been represented in the same architecture. The **GS** architecture is a relatively simple architecture. However, this chapter has demonstrated that behaviours which appear complex may result from the performance of this architecture, particularly in complex environments. In addition, its simplicity has also allowed it to act as a foundation for two types of further study. Chapter four shows how the **GS** architecture has been adapted to represent behaviours observed in Strange Situation studies. The behaviours recorded in this studies are of a similar timescale and level of abstraction to the behaviours found in the studies from this chapter. Chapter three demonstrates the power of the **GS** architecture by additionally representing considerably more abstract behavioural measures that are described in two meta-studies. The architectures in the next chapter are based upon the **GS** architecture. What this foundation provides is an architectural framework within which different goals can compete with each other for control of the agent's behaviour over different levels of abstraction and different timescales.

Chapter 3

Learning and security

> *"A young child who had gained security in his relationship with his parents was emboldened thereby to strike out to explore the world, willing to risk the insecurity initially implicit in a learning situation because he could rely on his parents to be available, responsive, protective and reassuring. If his adventure evoked undue anxiety, the child could easily return to "home base," in the expectation that his parents would provide the reassurance he needed. If, on the other hand, his relationship with his parents was insecure, then he might not dare to leave them to explore, not trusting them to remain available to him if he left or to be responsive when he needed them. Lacking trust, he would stick close to his base, fearing to risk the anxiety implicit in exploration and learning."*(page xi Ainsworth *et al.*, 1978)

3.1 Introduction

This chapter is concerned with reproducing abstract patterns of behaviour similar to those sketched in the quoted passage above. The central research question being studied in this chapter is: What kind of information processing architecture might human infants possess, that enables them to assess carer sensitivity and then form Secure and Insecure patterns of behaviour? The patterns of behaviour in the simulation presented in this chapter have been anchored in empirical observations reported in the literature of Attachment Theory. The stages that have been gone through in adapting observed behaviours to solutions and then evaluating those solutions are similar to the previous and following chapters:

Stage 1 - **The problem is established**: There are hundreds of studies that have some relevance to the issue of how infants develop Secure and Insecure styles of behaviour. The simulations reported in this chapter draw upon two meta-studies that together consider the results from 205 studies of infant attachment behaviour. The key finding made in both these meta-analyses is that the principal factor in the formation of an infant's level of attachment security in its relationship with its main carer is carer sensitivity. It is this finding that is at the core of the SECURE-BASE scenario.

Stage 2 - **The function of behaviour is analysed from an evolutionary perspective**: Why have humans evolved information processing mechanisms, so that, early in infancy they can develop Secure or Insecure patterns of behaviour with their carers? Put another way: what adaptive benefit might the mechanisms that confer the ability for an infant to become Secure or Insecure have held for infants in the EEA? An analysis of the evolutionary function of the ability to develop Secure or Insecure predispositions is discussed. It is hypothesised that an ability to form these patterns of behaviour is an adaptation to the conditions of parent-offspring conflict that are likely to have existed in the human EEA.

Stage 3 - **Two solutions are proposed**: that are both based upon the **GS** architecture that is described in chapter two. These solutions are motivated by consideration of how the **GS** architecture might be adapted to provide information about carer responsiveness to different types of infant bid for attention.

- The **goal-learning-from-anxiety (GLA)** architecture tracks carer response times from bids arising from the activation of the goal of security, and does not consider responses to bids for socialisation.
- The **goal-learning-from-socialising (GLS)** architecture assesses carer responses according to how the carer reacts to bids for socialisation and does not assess carer reaction to bids when the infant has activated the goal of security.

Stage 4 - **Computational experiments have been carried out which investigate the design space of the GLA and GLS architectures.**: When caregiving responsiveness is set to be highly sensitive or very insensitive the two architectures give similar results. When carers are highly sensitive the Safe-range distance in both the architectures increases until it reaches the maximum bound of the infant's visual range. The reverse occurs in both architectures when carers are very insensitive. The Safe-range distance decreases so that eventually even when the carer agent moves a short distance away it crosses the Safe-range threshold. However the **GLA** and **GLS** architectures show different results when the carer agent is set to produce intermediate levels of sensitivity in the way they respond. The **GLS** architecture produces intermediate Safe-range levels when it experiences intermediate levels of sensitivity. In contrast the **GLA** architecture does not settle at intermediate levels. When a carer starts out with an intermediate level of caregiving towards an infant possessing the **GLA**, the success or failure of early responses changes the infant's subjective appraisal of the carer. If the infant experiences a string of sensitive responses then its appraisal is that the carer is providing security, and the same intermediate level of carer response will be perceived as sensitive. If the infant experiences a string of insensitive responses then its appraisal is that the carer is not providing security. These results occur because interaction of **GLA** with their carer environments involves positive feedback loops that amplify initial disturbances to the level of the Safe-range parameter.

Stage 5 - **Implications are discussed for the assessment of the validity of categorical versus continuous forms of classification of attachment**: These results are relevant to an evaluation of the categorical assessment used in Strange Situation studies. Classifying infants into categories relies upon the fact that the infants' behaviours are clustered. This means that when a random sample of infants is assessed for security, the values do not follow a normal distribution. Instead, high and low levels of security are proportionately more represented, and average levels of security less represented, than they would be if normally distributed. The behaviour of the **GLA** architecture shows how this clustering may occur in reality.

3.2 Framing the problem

3.2.1 Selecting and abstracting the source material

Infants learn about the effectiveness of their carers in responding to their signals for attention. If an infant comes to possess confidence in the carer's responsiveness then the infant is described as Secure. Alternatively, a lack of infant confidence in the carer's responsiveness means that the infant is described as Insecure. This distinction was brought to prominence by Mary Ainsworth and co-workers. In the 1960s, whilst investigating cross-cultural patterns of attachment, Mary Ainsworth was struck by a difference between the types of attachment behaviours observed in infants from rural villages in Uganda and infants in suburban communities in Baltimore, USA (Ainsworth, 1967; Ainsworth *et al.*, 1978). Under observation, the Ugandan infants exhibited significantly elevated levels of anxious behaviour, despite being in familiar surroundings. The presence of observers in Baltimore households did not generate similar levels of anxiety. Ainsworth *et al.* (1978) hypothesised that the intense behaviours seen in the Ugandan study might be evoked more incisively if the Baltimore infants were put in an unfamiliar environment. The Strange Situation was therefore created. In the first exploratory Strange Situation study three styles of attachment were distinguished, and in later studies a fourth was recognised. Of the four groups, Ainsworth *et al.* (1978) labelled Securely attached infants as type B. The infants that formed the other three groups, labelled the A, C and D types, are all held to be Insecurely attached. The distinction between type A Insecure infants, known as Avoidant infants, and the type C Insecure infants, known as Ambivalent infants, is the subject of investigation in chapter four of this thesis. The type D Insecure infants, known as Disorganised, form a very small proportion of infants in non-clinical populations, and are not considered further in this thesis.

One of the problems with the Strange Situation study is that the categorical nature of its classification system rules out the use of analytic techniques that rely upon continuously measured characteristics. To overcome this problem Waters and Deane (1985) developed the Attachment Q Sort methodology. This method of assessing an infant's style of attachment is based upon home observations of an infant and its carer's behaviour. It outputs values for infant attachment along a single dimension of security and these are linked to different aspects of care provided to the infant in the home.

There are a very large number of Strange Situation studies and Attachment Q Sort studies in the Attachment Literature. A major problem for any analysis of these studies is that many studies have minor variations that thwart a simple unified analysis. Wolff and van IJzendoorn (1997) and van IJzendoorn *et al.* (2004) have undertaken meta-analyses of respectively, sixty-six Strange Situation studies and one hundred and thirty-nine Q sort studies. The study of Wolff and van IJzendoorn (1997) is based upon results from 4,176 infants that were based upon the Strange Situation paradigm, and van IJzendoorn *et al.* (2004) describes a meta-analysis of 13,835 infants using the Attachment Q Sort methodology. These meta-analyses have both produced the same major finding, which is that it is carer sensitivity that is the principal cause of the level of attachment security in infants. This result does not need abstracting to be included in an autonomous agent simulation. The statistical analyses undertaken by Wolff and van IJzendoorn (1997) and van IJzendoorn *et al.* (2004) is already removed from the details of the individual studies and behaviour of individual infants.

Since the results from the Strange Situation and Q Sort studies are being combined the current chapter will be confined to investigating the distinction between Secure infants and a composite of the Avoidant and Ambivalent Insecure groups. This approach is justified because the Secure-Insecure distinction is the principal division within non-clinical infant populations. Insecure patterns of behaviour have more in common with each other than they do with Secure patterns of behaviour. For example, records of infant behaviour from the first three months after birth can be analysed with the knowledge of the classification those infants gain at one year. When this is done the early behaviours of Secure infants already show significant differences with the Insecure infants. However, no significant differences can be made between the Insecure groups of infants (Ainsworth *et al.*, 1978).

3.2.2 Analysis of function

As Cassidy and Berlin (1994) relate, within Attachment Theory infants are viewed as following particular strategies when interacting with their carers:

> *"Just as other organisms are genetically endowed to be flexibly responsive to the range of environments in which they may find themselves, so too ... does the infant possess the biologically based flexibility to adapt to a range of caregiving 'environments' or behaviour. Such tailoring of the infant's behaviour to conform to environmental circumstances is described as a 'strategy' which leads to behaviour that has the ultimate function of protection. Strategies are thought to be automatically employed and need not be in any way conscious for the individual"* (page 122, Cassidy and Berlin, 1994)

This section proposes that the strategies that infants possess are derived from the architectures that they are born with, and then adapted by caregiving experiences after birth. These architectures are an adaptation that allowed the infants to adjust their behaviour to match different forms of caregiving style in the human EEA. This hypothesis is not that patterns of contemporary caregiving behaviour observed in Strange Situation studies existed in the EEA, rather the hypothesis suggests that more extreme forms of caregiving behaviour in the EEA may have selected for mechanisms that are now triggered by aspects of moderate caregiving.

A key question that this section considers is: how likely is it that there were selective pressures in the EEA which were strong and directed enough to bring about the control mechanisms being discussed? Infants rely upon their carers to satisfy nearly all their needs, and in many circumstances carers attempt to fulfil those needs to the complete satisfaction of the infants. So for example, hungry infants cry for food and are then fed until they are satiated. There may be circumstances when carers simply cannot satisfy all their infant's requirements. A carer may have too few resources, or too many infants to completely fulfil all of their needs. Trivers' (1974) theory on parent-offspring conflict shows that even when carers can satisfy all of their infants' needs, natural selection would favour holding energy and other resources back to use for unborn and even unconceived future offspring. Parents are equally related to all their offspring; those that currently exist and those that may exist in the future. However, offspring in diploid species have a greater interest in their own future than in their siblings, so a conflict of interest arises. Trivers provided evidence of this conflict in bird species and primate species such as Lemurs, Baboons and Rhesus Monkeys, when offspring are about to go from being cared for and fed by their parents to being self-sufficient. Human conflict over the timing of weaning may be an example of this type of conflict. This is because infants may want to continue with breast-feeding whereas mothers want the infant to move to other food sources so that they can give more resources to existing infants or prepare for future infants.

Since Trivers' original paper, parent-offspring conflict has been more widely applied to explain human behavioural and physiological phenomena. Haig (1993) extended evidence of this conflict from weaning back into the womb, characterising the behaviour of the foetus, via its placenta, as similar to an internal parasite. Ridley (1996) reviews this research and recounts that:

> *"Cells from the foetus invade the artery supplying maternal blood to the placenta, embed themselves in the walls and destroy the muscle cells there, thus removing the mother's control over constriction of the artery. The high blood pressure and pre-eclampsia that often complicate pregnancy are caused largely by the foetus, using hormones to try to divert the mother's blood to itself by reducing the flow through her other tissues.* [...] *Haig's point is not to try to claim that all pregnancies are tugs of bitter war between enemies; mother and child are still basically cooperating in the business of rearing the child. The mother is still astonishingly selfless as an individual in the way she nurtures and protects her children. But, as well as shared genetic interest between them, there is also some divergent genetic ambitions. The mother's selflessness conceals the fact that her genes act as if motivated entirely by selfishness, whether being nice to the foetus or fighting it."* (page 23, Haig, 1993)

As with the evolutionary analysis in the previous chapter, the hypothesis put forward in this chapter implies a trade-off. Infants possess information processing architectures to maximise their learning opportunities whilst maintaining safety within acceptable bounds. If the carer acts in a manner that suggests it is looking after the infant's requirement for safety then the infant can concentrate on learning. If the carer does not act in a way that suggests it is ensuring the infant's safety then the infant must take its security more into its own hands. An obvious way that an infant will bring about more safety for itself is to get the carer to provide more safety. This the infant can do by signalling to the carer when the carer moves too far away. If a carer is juggling multiple goals, one of which is the infant's safety, and the infant increases the insistence of its signalling this may increase the priority of infant safety amongst the carer's goals. The theory that more intense infant signalling behaviours are linked to parent-offspring conflict directly parallels Trivers' notion that weaning in many species, and feeding behaviour in species such as birds, are examples of parent-offspring conflict. When a bird brings food to the nest to feed its chicks, all the chicks noisily demand food. Substituting safety for food, if a human infant assesses it is safe it does not demand more safety. If it assesses it is unsafe then it signals more vigorously to get more attention, and hence more safety. The next chapter deals with what happens if an infant also assesses its carers as a threat, in addition to being unreliable providers of security. In this case simply 'asking louder' for more safety may not be the best strategy.

So far we have considered why an evolutionary trade-off might affect mechanisms that are selected for in the EEA. This leaves the question: what kind of mechanism might evolve within the infant's information processing architecture that allows the infant to adapt to different styles of caregiving. This thesis takes the view that the mechanistic structures of attachment styles can be compared with other kinds of affective control states. Short term affective control states such as anger have functions such as communication and action readiness. The mechanisms that give rise to Secure and Insecure attachment styles can be viewed as longer term affective control states (Sloman, 1995) which are somewhat similar to personality, skills and attitudes (figure 3.1). Attachment styles can be formed from the experience of low level events and short term control states. Possession of a certain attachment style predisposes an architecture to particular short term patterns of states such as plans or emotions. The function of long term affective states is therefore to organise behaviour at a higher level. Attachment styles provide a filter through which all perceptions and actions are made. As an infant gets older their attachment style merges with other elements to form their adult personality (Marvin and Britner, 1999).

3.2.3 Forming the scenario

The SECURE-BASE scenario discussed in this chapter is a follow-on from the PARK scenario described in chapter two. Both the PARK scenario and the SECURE-

LONGER TERM	**INTERMEDIATE**	**SHORTER TERM**
Personality, **Temperament,** **Attitudes,** **Skills,** **Emotions such as love, grief,** ***Attachment style?***	**Moods,** **Beliefs,** **Preferences,** **Emotions such as joy, fear,** **Intentions,** **Plans,** **Desires**	**neural and physical events,**

Figure 3.1: Classes of semantic control state, which are compared with respect to the approximate duration that each class of control state may exist as a disposition within an architecture. Adapted from Sloman (1995).

BASE scenario are precursors to the REUNION scenario in chapter four. The motivation for the SECURE-BASE scenario is to respond to limitations of the PARK scenario by adding an ability for the infants to adapt to the style of caregiving they receive. So a prerequisite for designs to fulfil this scenario is that they also fulfil the previous PARK scenario. An important qualification on the objectives of the SECURE-BASE scenario is that it is only attempting to provide the infant with the ability to adapt to patterns of carer responsiveness. It is not attempting to include other types of learning that were missing from the PARK scenario, such as learning to fear objects or a more realistic form of object exploration.

In the SECURE-BASE scenario the carer agents are active within the virtual world. The carer agent has two tasks. The first task is to forage for food and the second task is to respond to the infant agent's signals. The carer agent follows a very simple routine when responding to the infant agent. The carer agent will only respond to the infant agent if the infant is signalling at a level above the level of the **response-threshold** parameter possessed by the carer agent. This means that when the infant agent is signalling at a level below the **response-threshold** the infant agent will be ignored. When the infant agent is signalling at a level above the **response-threshold** the carer agent will stop foraging, turn towards the infant agent, start signalling itself, and move towards the infant agent. If the infant agent stops signalling then the carer agent will return to foraging. If the infant agent continues signalling then the carer agent will move until it is in contact with the infant agent.

The presence of food objects extends the scenario ontology but does not extend the infant agent's ontology. This is because the food objects are invisible to the infant agents. The infant agent's general pattern of behaviour in the SECURE-BASE scenario is similar to the way the infant agents behave in the PARK scenario. The infant agent possesses three goals: exploration; socialisation and security. These vary in activation level. In the PARK scenario the carer agents were stationary, so to gain security the infant agents had to move towards the carer agent until the carer was within their Safe-range distance. A difference in the SECURE-BASE scenario is that the infant agents can gain their own security by moving, and in addition the infants signal for the carers to move back during foraging activities. Modes of evaluation include graphical displays of the running simulation and analysis of trace output. The graphs in figures 3.2 and 3.7 demonstrate how the different architecture fulfil the SECURE-BASE scenario at different levels of caregiving sensitivity.

3.3 The GLA and the GLS architectures

3.3.1 Extending the GS architecture with sensors and memory

The **GLA** and **GLS** architectures share the same goal selection mechanism as the **GS** architecture described in the last chapter. Hence infant agents act to explore, socialise or gain safety when the goals for these actions are selected in a winner-take-all mechanism. In all three architectures it is the Safe-range parameter that determines the distance at which activation of the goal of gaining safety is initially triggered. The **GLA** and **GLS** architectures differ from the **GS** architecture because they adapt the value of the Safe-range distance up or down depending upon the responsiveness of the carer. The **GLA** updates the Safe-range distance every time the infant agent signals to gain safety, and the **GLS** architecture updates the Safe-range distance every time the infant agent signals to socialise. The Safe-range distance therefore acts as a measure of the confidence of the infant in its carer's responsiveness. High Safe-range distances mean that the infant regards the carer as providing security. Low Safe-range distances mean the opposite, namely that the infant has insecurity in its relationship with its carer.

The amount that the Safe-range parameter is increased or decreased is always small compared to the magnitude of the Safe-range parameter and is given by functions that are discussed in more detail in section 3.3.4. The mechanisms and structures that these architectures possess, and which the **GS** architecture does not possess, include: a sensor mechanism that notes when bids for carer attention have been made; a memory system that stores this information until a response is sensed; and a system that changes the value of the Safe-range parameter in reaction to the latency of carer response.

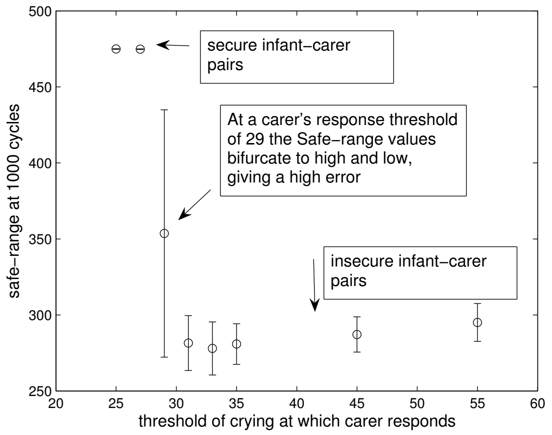

Figure 3.2: A snapshot of the infant Safe-range found at 1000 cycles for different values of carer responsiveness. Each data point is a mean from 20 simulations.

Changing the Safe-range parameter changes the pattern of the infant's behaviour. The Safe-range parameter controls the point at which the security goal is activated. The higher the Safe-range parameter, the further away from the infant the carer can travel without the security goal in the infant being activated. If the Safe-range parameter is small then the carer can only go a small distance from the infant before the security goal rises sufficiently for the infant's behaviour to switch. The infant then turns and starts signalling to the carer. The Safe-range parameter is only updated on the small proportion of cycles in the simulation where the carer returns from being outside the current Safe-range distance. This means that simulations can run for thousands of time cycles and only include dozens of separate learning experiences.

3.3.2 Experiments with the GLA architecture

Computational experiments have been carried out to assess the performance of the **GLA** architecture in varying conditions of carer responsiveness. Figure 3.2 shows the results of some of these experiments. Each of the data points in figure 3.2 is formed from averaging the results of twenty simulations. In all conditions, each of the individual simulations was run for the same number of cycles and started with infant agents in the same initial state. In each simulation the starting value of the Safe-range parameter was 320 distance units. Any differences in the end state of the infants either results from the effect of interacting within different caregiving environments or small random effects (or some interaction of both types of effect).

Infants start signalling to the carer as soon as the security goal is activated. Carer sensitivity is represented in the simulation by the threshold of infant signalling at which the carer will respond.

The value of the Safe-range parameter after a thousand cycles in environments with extreme levels of carer sensitivity is predictable. If a carer responds almost instantaneously then it will always be assessed as prompt and the Safe-range value will always rise. Infant agents with very sensitive carer agents will always end up being in secure pairings in the virtual world of the **GLA** architecture. These type of results can be seen in the top left corner of 3.2, where carer sensitivity values of 25 and 27 have resulted in infant Safe-range values of 480. This is the value of the visual range for the infant, and thus acts as an upper bound for the Safe-range parameter. When the carer is extremely insensitive the reverse result occurs. The infant always lowers the value of the Safe-range parameter, which consequently ends up very low. So low carer sensitivity invariably ends up with insecure infant-carer pairs, as shown by results at the bottom right of figure 3.2.

The experiments presented in figure 3.2 were exploratory, with the purpose of finding out what level of carer sensitivity would constitute an intermediate level of response. The intermediate level of carer sensitivity seems to be at a value of 29, which means that the carer waits until the infant is signalling at a level of 29 signalling units before it starts to respond. This is a somewhat arbitrary figure. If the speed that the carer agent moved was changed, or the objects that the infant explored were slightly more or less interesting, then the carer response threshold which gave an intermediate Safe-range value in the infant agent would be different. The average confidence level for infant agents who experienced caregiving with a response threshold of 29 was a Safe-range of 368. However, this average value of 368 is gained from the results of 20 simulations, 6 of which possessed Safe-range values above 400, and 4 of which possessed Safe-range values less than 300. It appeared that when a carer response threshold is set at 29 the security of infants in their carers has spread widely between high and low levels, though their carers continue reacting to the same level of signalling throughout the simulation. These results show that at the particular set of starting parameters in this experiment, the carer response threshold of 29 signifies a critical point (or perhaps critical region). At this critical point the infant Safe-range values bifurcate between high and low values with very few at an intermediate level.

The results in figure 3.2 show snapshots of simulations, all taken at the same time. In particular the results where the carer threshold is 29 and the infant results are widely spread out do not say anything about how the distribution of results might change given further time in simulation. The results for ten simulations carried out with a carer sensitivity of 29, carried out for ten thousand cycles, are shown in figure 3.3. These experiments illustrate the transitional nature of the

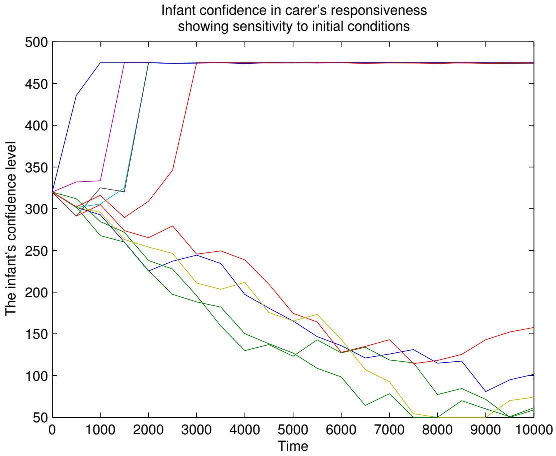

Figure 3.3: A graph showing the results from 10 experiments which were all over 10000 cycles duration, for the carerB and babyB with the carer's parameter of when to respond to crying set at 29.

results shown in figure 3.2.

When the results are viewed longitudinally, as in figure 3.3, it is apparent that the infant agents which possess intermediate levels of the Safe-range parameter at one thousand cycles come to possess either high or low levels at ten thousand cycles. Five of the initial ten infant agents turn out to be Secure and five end up Insecure. It is worth repeating that they all started in identical states and their carer agents possessed identical behavioural policies. What has differentiated the development of the infant agents are small random effects early in the simulations. These random effects appear throughout the duration of the simulation. For example, the positions of the food objects are random, so that each time the carer agent starts to forage the distance that it has to move away from the infant agent will be different. Each time the infant agent signals to the carer agent with the signal level above the carer agent's response threshold it triggers carer action to return to the infant, and for each of these occasions the carer will be a varying distance from the infant. The carer agents always move at the same speed, sometimes they will be close and return promptly. On other occasions they will be further, and the time they take to return will be assessed as being slow to respond and overdue. When random effects appear early in the simulation they can push the simulation into forming Secure or Insecure infants, depending on how they affect the carer agent's time to return. This is because once the Safe-range has been increased or decreased the probability of the next response being assessed as prompt or tardy has been changed. For example, if the infant agent gets a late response,

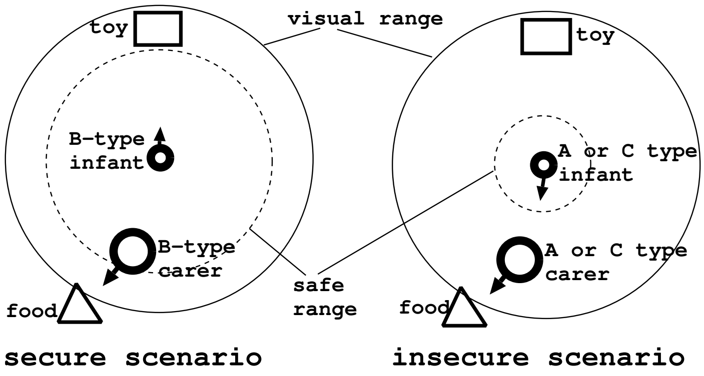

Figure 3.4: A secure infant moves towards a toy as its carer moves away towards food. When an insecure infant experiences the same event it signals to, and moves towards, the carer.

the Safe-range distance is decreased. It is now more likely that the carer will go over the Safe-range distance more often, there being less area marked out by this distance. If a number of decreases in the Safe-range distance occurred without any intervening prompt responses, the Safe-range distance would become increasingly small. This means the infant agent is becoming chronically untrusting of the carer agent's performance. The probability of further responses being viewed as taking too long will only increase further. The reverse obviously holds true for carers that carry out a series of prompt returns. This state of affairs is demonstrated in figure 3.4. In summary, the ten simulations in figure 3.3 do not remain on identical trajectories because positive feedback acts upon small random effects.

What effect does possesion of a small Safe-range value have on an infant agent's perception of its carer agent's behaviour, and therefore its own consequent actions? Figure 3.4 shows how contrasting infant agent behaviours result from situations where carer agents are currently making identical relative movements. In the secure scenario, the carer agent is moving towards food but is still within the infant's Safe-range so the infant is not signalling and is moving towards an unexplored toy object. The distance that the carer has to travel from the edge of the safe-range to the food is short. If the carer is called back it can cover this distance quickly. In the insecure scenario, the infant's past experiences have given rise to a smaller Safe-range. The carer has already crossed this boundary and the infant has switched goals from exploration to bringing the carer back within its Safe-range by signalling and moving towards the carer. In this scenario the carer may be the same distance from the infant and the same distance to the food. However, the distance from the edge of the safe-range to the food is greater. When the carer is called back it will take longer, even if it has the same response threshold as the carer in the secure scenario.

To investigate the behaviour of the **GLA** architecture close to the carer agent response threshold of 29 two further sets of longitudinal experiments have been carried out. The next two figures show the sensitivity of the simulation to small changes in the carer's responsiveness. In figure 3.5 the carer is marginally more sensitive, responding to infant cries when they reach a level of 28. This shows all infants becoming maximally confident. After an initial period, all infant-carer pairs become secure.

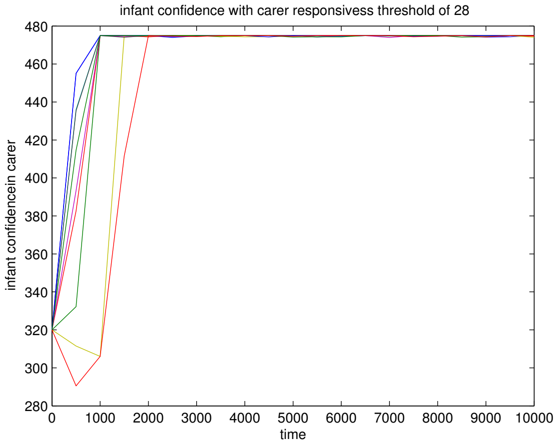

Figure 3.5: A graph showing the results from 10 experiments which were all over 10000 cycles duration, with the carer's response level set at 28.

Figure 3.6 shows that when the carer is marginally less sensitive, responding at a level of 30 all but one of the infants become minimally confident and therefore Insecure. Part of the reason that the simulation gives such a sharp transition from parameters that give complete security to nearly complete insecurity is that each carer agent's behaviour is rigid and unvarying within each simulation.

From the graphs in figures 3.3 to 3.6 we can see that the learning rule causes gradual decreases in confidence when it is decreasing, but much less gradual increases in the safe-range when it is increasing. It might be argued that this is not how this type of learning occurs in reality. An alternative is that confidence only increases gradually, but decreases abruptly when expectations are violated. The learning rule is discussed in more detail in section 3.3.4.

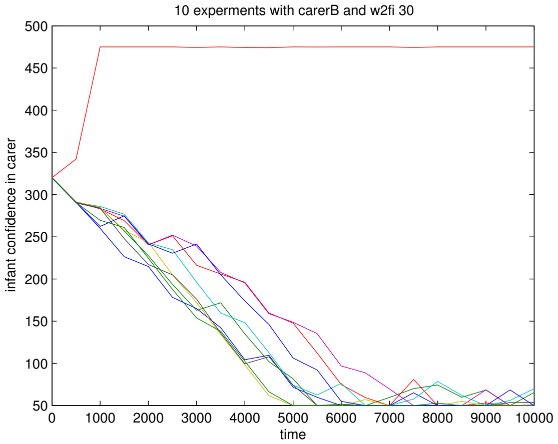

Figure 3.6: A graph showing the results from 10 experiments which were all over 10000 cycles duration, for the carerB and babyB with the carer's response parameter set at 30.

3.3.3 Experiments with the GLS architecture

In the previous section the **GLA** architecture was initially investigated by carrying out experiments whereas carer responsiveness was varied in order to find the level at which carer responsiveness would give intermediate values of infant security. The same experiments were repeated with infants possessing the **GLS** architecture. Figure 3.8 shows the results. These experiments show a much less pronounced critical region because there is no positive feedback due to the changing Safe-range parameter. The attachment status of the infants varies from secure to insecure with less clustering. This is because the learning of the Safe-range distance is independent of the infant's behaviour in episodes of anxiety. Despite the absence of a feedback loop similar to that in the **GLA** architecture the behaviour of the **GLS** architecture remains complex and has not been fully analysed. Both the **GLA** and the **GLS** architectures represent pure forms of learning. Real infants are likely to be able to learn about their carer's effectiveness from anxious and sociable episodes. Architectures that possess just one of these mechanisms, with no ability for the learning in other contexts are not very plausible. However, it may be that real infants possess learning biases that greatly favour learning from particular contexts, such as anxious episodes, and it is this bias that gives rise to certain statistical properties of the distribution of attachment styles that are discussed in section 3.4.1. The balance of the mechanisms which real infants possess is an active research question.

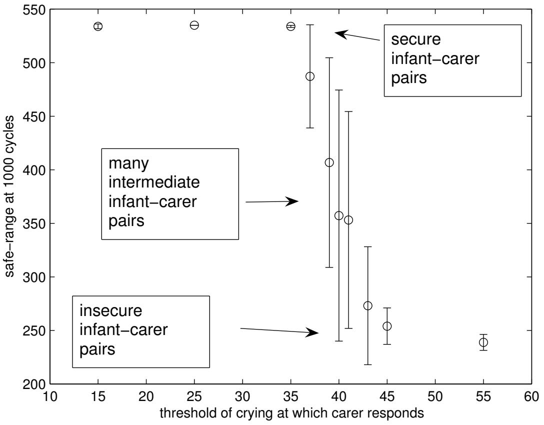

Figure 3.7: A snapshot of the infant Safe-range found at 1000 cycles for different values of carer responsiveness Each data point is a mean from 20 simulations.

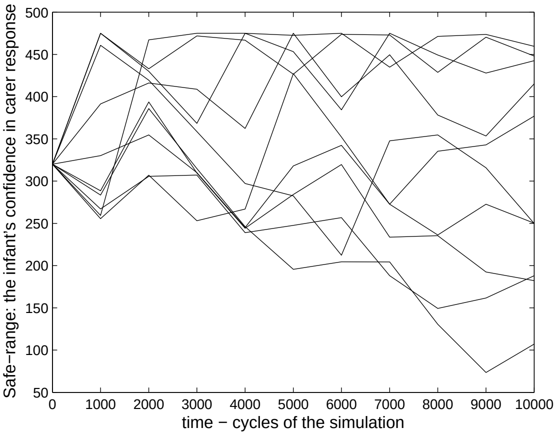

Figure 3.8: Infant Safe-range changing over time, with infants learning about carer response only from interactions with no anxiety.

3.3.4 A description of the novel learning mechanism

Reinforcement learning in living organisms fits within an evolutionary framework. Complex organisms such as humans are not provided with a single behavioural policy which is set at birth to last through their lives. An initial behavioural policy may be set, but a human spends the rest of its life customising and fine tuning this policy, often in response to feedback. For example, infants are born with a sucking reflex, but quickly learn how to adapt that reflex to better fit with the nature of the breast or bottle from which it regularly feeds. Policies can change by an infant choosing between currently available alternatives, or change because the infant develops abilities which enlarge the available behaviours in the policy. We can view infant agents in the simulation as possessing a behavioural policy that describes how they will behave in situations relevant to attachment. The **GLA** and **GLS** architectures both allow the infant to change its policy, not by introducing new types of behaviour, but simply by changing the value of the Safe-range parameter. For an infant possessing the **GLA** and **GLS** architectures to explore different attachment policies means exploring the effects of setting the Safe-range parameter at different levels. The problem that human infants faced in the EEA was to set their behavioural policies so that they optimised the potential for learning about their environment, whilst accounting for the reliability of the security provided by their carer. In the **GLA** and **GLS** architectures this means finding a value of the Safe-range parameter that matches the reliability of their carer agent.

The learning problem faced by the infant agents is to use feedback from the responses of their bids for attention to set the Safe-range level appropriately. This problem is intended to mirror the problem faced by real infants. Part of the problem for both real and virtual infants is that a usually reliable carer will sometimes produce unreliable responses and the reverse will also occur for unreliable carers. In addition, both real and virtual infants do not have the cognitive ability to reason that any given response is atypical or representative. Real infants have to take a statistical approach to the problem, aided by the possession of instinctive cues to danger and cues to safety. In the simulation produced in this work, infants have been given an instinctive ‘cue’ to associate security with carer proximity. Infant agents assess that when they are in proximity to their carer they are safe. When they assess their carer is outside their measure of proximity they are unsafe. The choice of reinforcement learning algorithm for the infant agents is constrained in the type of processing requirements, such as memory capacity, that are plausible for a human infant to possess.

The learning mechanism in the **GLA** and **GLS** architectures is based upon the infant being able to record how long the carer takes to respond to its request for action. Both the GLA and GLS architectures translate response latencies into reward signals. For example, very prompt responses give high reward signals, very tardy responses give punishment (negative reward) signals, and intermediate responses may not give a reward or punishment. When the infant receives a response from the carer with a given latency, that latency is used to calculate the size of the reward.

Appendix C presents a formal description of the reinforcement function. The novel aspects of the learning function include that both positive and negative reinforcers are present, and the positive reinforcement signal is produced by a discount function that possesses beneficial aspects of existing finite horizon and geometric discount functions.

3.4 Theoretical implications

3.4.1 The origins of the categorical classification of attachment differences

Mary Ainsworth and co-workers developed a classificatory system for infant attachment behaviour as they constructed the procedure for the Strange Situation Experiment. Prior to the study reported by Ainsworth *et al.* (1978), a classification system had been developed that sorted infants into groups according to the ways in which their attachment related behaviour was organised. Initially it had been thought that the key behaviours to use in discriminating individuals with differently organised approaches to attachment were the responses to separation. After further analysis it became apparent that it was the reunion behaviours that gave the clearest discrimination between the three different types of attachment relationships. That is, it was the responses to reunion which gave the clearest delineation of infant responses into three clusters. The categorical form of classification devised by Ainsworth *et al.* (1978) has been used in hundreds of studies. Recently, attempts have been made to use current statistical techniques to analyse the validity of classifying attachment behaviour into A, B and C categories rather than assessing infant behaviour along continuous dimensions. This issue is relevant to the results described in this chapter because the computational experiments with the **GLA** architecture suggest distinct categories of attachment behaviour can result from initially identical carer-infant pairs.

3.4.2 A modern taxometric analysis of attachment differences

Fraley and Spieker (2003) have carried out a statistical analysis of Strange Situation behaviour. They used statistical techniques to assess whether attachment behavioural patterns in the Strange Situation are better described with categorical or continuous forms of classification. Fraley and Spieker (2003) suggest that categorical forms of classification are justified if the categories being analysed have the properties of 'taxons'. The term 'taxon' is used to refer to "naturally occuring types," "natural kinds," or "nonarbitrary classes". According to Fraley and Spieker (2003):

> *"if security* [of attachment] *represents a natural taxon, then there is something qualitatively different about the way the attachment system operates for secure children and insecure children, something that elicits, probabilisitically, behaviours indicative of security (e.g. contact seeking when distressed, being easily soothed by an attachment figure) among children who belong to this taxon and that fails to elicit, probabilistically, these behaviours among children who do not belong to the taxon.*(page 389, Fraley and Spieker, 2003)

The theoretical underpinning of Fraley and Spieker's analysis is based upon techniques of taxometric analysis developed by Meehl (1999). Ideally the presence of a taxon would be determined by showing that all group members of that taxon *"share a common discrete source of influence on their manifest characteristics"* (Fraley and Spieker, 2003). Here we will term this the causal definition for a taxon $taxon_{causal}$. An easy case for ascribing taxonicity according to the $taxon_{causal}$ criteria is classifying the sex of most species of animal where genetic tests can be made. Male and female animals, including humans, form taxons because of a biological-genetic source of influence. More difficult cases to determine if taxons exist according to the $taxon_{causal}$ criteria might arise when considering humans of different religions or members of political parties. These may form different taxons if within their groups they share a tightly knit ideology or possess different goals or belief structures that act as common discrete sources of influence on their manifest characteristics. However, it is not always possible to ascribe taxons by tracing back from observable differences to find the ultimate causes of those observable differences. For example, one may only have cross-sectional data on a population available when attempting to assign taxonicity. Meehl notes,

> *"Taxometrics is a statistical procedure for determining whether relationships among observables reflect the existence of a latent taxon (type, species, category, disease entity). A formal-numerical definition is needed because intuitive, commonsense notions of "carving*

> *nature at its joints" or "identifying natural kinds" cannot resolve disagreements as to taxonic reality for hard cases. Specific etiology (e.g. major gene, germ, traumatic event) is often unknown and is not appropriate in nonmedical domains. Lacking an infallible criterion, the taxonic inference relies on the internal configural relations among the conjectural fallible indicators. ... Common misconceptions are that the taxon must be "sharply" distinguished, quantitative indicators must be bimodal, the causal origin must be biological, emergence of a large dimensional factor refutes taxonicity, and adopting a taxon is merely a matter of convention or preference.*(page 165, Meehl, 1999)

Meehl's statistical analysis does not make any claim to the actual causal processes that have brought about observable differences between groups. It is an alternative to causal analysis according to the $taxon_{causal}$ criteria, and occurs where causal analysis is intractable or impossible. It gives rise to a different kind of statistical criterion for a taxon, as Meehl explains:

> *"I would like to offer a general rigorous definition of the work "taxon" but after trying for many years, I have been unable to generate one. Fortunately, we need not define all terms explicitly or operationally, as is misleadingly told to beginning psychology students. When teaching, I provide a formal-numerical implicit definition and call it "$taxon_{PM}$" meaning taxon as defined by Paul Meehl's method. If a population (patients, honeybees, stones, daffodils) is characterised by several quantitative observables such that statistical relations among these observables satisfy certain taxometric criterial derivable from a postulated latent structural model, then the situation is taxonic."*(page 167, Meehl, 1999)

Therefore, in addition to the $taxon_{causal}$ criteria, we now have the $taxon_{PM}$ statistical criteria. The results discussed previously in this chapter suggest a third criteria for taxonicity. Agent based simulations provide a dynamic definition of taxon that is concerned with how taxons may arise where members of categories do not "share a common discrete source of influence on their manifest characteristics" (Fraley and Spieker, 2003). The positive feedback loop in the experiments with the **GLA** architecture demonstrates that the $taxon_{causal}$ and $taxon_{PM}$ criteria do not cover all interesting cases of taxonicity. Therefore in addition to the $taxon_{causal}$ and $taxon_{PM}$ criteria for taxonicity we can add the $taxon_{emergent}$ criteria[1]. Groups are assessed to be taxons according to this criteria if the dynamic

---

[1] This is really just another kind of causal criterion, where the cause is external processes in the environment rather than a mechanism in the child. Causes can be internal or external.

properties of the systems in which they exist include feedback loops that will drive the formation of differentiable categories from initially similar states. The results for the architectures discussed earlier in this chapter show that the **GLA** architecture gave rise to $taxon_{emergent}$ categories, whilst the **GLS** architecture did not. The concept of $taxon_{emergent}$ categories is of broader interest than attachment theory. Any system that possesses the potential for strong positive feedback loops that can act upon small differences in initial conditions, or in small random effects that occur after initialisation, may produce outcomes that appear to be discrete categories but were initially indistinguishable.

Sensitivity to initial conditions is sometimes termed the 'Butterfly effect' (Gleick, 1996). This term captures a property of weather systems such that a small change in initial conditions, which might be produced by a butterfly flapping its wings, might have large changes in weather conditions like causing a tornado to occur. The Butterfly effect is commonly linked with Chaotic systems. However, for a system to be chaotic it also needs to possess a property of recurrence, which is the approximate return of the system towards its initial conditions.

Waters and Beauchaine (2003) criticise Fraley and Speikers's work because the latters' taxometric analysis was only applied to existing categories. An answer to the question of which categories might be the subject of taxometric analysis may come from simulations such as those described in this chapter. As was discussed in the introductory chapter, when faced with complex data, simulations can give simpler models than the kind of models which analysis of complex data can provide on its own. This is because complex patterns may emerge from simple models. The categories that are suggested by synthesis can then be validated by taxometric analysis of real empirical data.

Whether attachment categories can be viewed as examples of $taxon_{emergent}$ categories is an empirical question. If attachment categories were found to be $taxon_{emergent}$ (or have a high element of differentiation due to positive feedback loops) it might have implications for clinical practice, if only because a carer having knowledge of these mechanisms may help that carer to break current patterns of behaviour and gain a more reflective caregiving practice.

3.5 Evaluation of the GLA and GLS architectures

Computational models need to distinguish between core theoretical assumptions, peripheral theoretical assumptions and implementational assumptions. If mechanisms are held to be peripheral rather than of core importance it means that the same behaviour might be realised in other equivalent ways, or the precise detail of the assumption is not critical in the performance of the model.

The reinforcement learning algorithm possesed by both the **GLA** and the **GLS** architectures is an example of a peripheral aspect of the theory being proposed in this chapter. This is because the important aspect of the reinforcement algorithm is that it can distinguish between prompt and tardy responses. This thesis conjectures that the results seen in chapter are likely to be largely independent of the precise details of this algorithm. The confirmation of this conjecture is one of the aspects of the future work arising from this project.

The infant agent's internal measure of security in their carer agents is another example of a peripheral assumption. Bowlby (1969) initially described infant security solely in terms of proximity. However, Sroufe and Waters (1977a) later introduced the term 'felt-security' to emphasise that infants are not merely measuring a one dimensional distance between infant and carer. Other factors, such as carer attentiveness, are used by the infant to measure security. We can view the Safe-range parameter in two ways. Either it is a direct representation of the proximity component of felt security, and therefore the simulation is lacking factors such as carer attentiveness. Alternatively we can view it as as an abstract representation of a much richer measure of security. This richer representation includes more features than distance alone. This does not change the nature of the findings from the simulation, whereby incidences of high or low carer responsiveness change the criteria by which the level of future security is assessed.

A significant limitation in the scenario is that it only requires a single dimension for affective tone. The basic cry is associated with basic needs such as hunger, and the angry and pain cries are more intense. The scenario requires that a signalling infant be distinguished from a silent infant, and that a smiling infant be distinguished from a crying infant. However, the scenario doesn't require that infants should have different ways of crying. In reality infants do have different kinds of cry. Real infants produce three types of cry (Durkin, 1995). The basic cry, the angry ("mad") cry and the pain cry can all be distinguished by spectrograms and by human audiences. Justification for only having one dimension of crying is that the designs do allow the affective tone to be varied, so that infants do sound more 'upset'. What the infants don't possess is a qualitatively different kind of sound.

In studies that are more closely based upon the Strange Situation, there are differences in the behavioural measures that have been recorded and analysed. The initial attachment study suggested that maternal sensitivity was the key factor in determining an infant's attachment status. More recently, other contributors have suggested that it is the quality of general social interaction between mother and infant that matters (Schore, 1994). Wolff and van IJzendoorn (1997) found that *"sensitivity is an important but not exclusive condition of attachment security"*(page 571Wolff and van IJzendoorn (1997)). The other causal factors for differences in attachment style included the variation in general affectional interaction along a dimension of warmth to coldness. This distinction is highly salient in the task of designing agents. This is because the attachment system and affectional system are distinct domains of behaviour. Macdonald (1992) provides four reasons for making a distinction between systems underlying intimacy in human relationships and the systems underlying the propensity for fear in the absence of an attachment figure:

- Positive and negative emotions appear to result from different biological systems.
- Attachment is found in many species of mammals and birds and is ubiquitous in primates. Pair bonding is rarer, only 17 % of primates show pair bonding, in most primates mating is promiscuous.
- Attachment can be compatible with lack of warmth, *"even in the face of punitive and abusive behaviour by the carer"*.
- Intimacy shows sex difference and attachment does not.

Macdonald (1992) also suggests that the attachment and affiliative systems carry out different evolutionary functions:

> *"the function of the attachment system is to provide security in the face of threat, the human affectional system functions to facilitate cohesive, psychologically rewarding family relationships and paternal investment in children"* (page 756, Macdonald, 1992)

The next chapter is concerned with simulating the Strange Situation Experiment. This chapter has produced two infant architecture which measure different aspects of the carer agent to infant agent relationships. The **GLA** architecture measures responses to infant agent bids for attention when the infant agent is anxious and the **GLS** architecture measures responses when the infant agent wants to socialise. Which of these two architectures is the best starting point for a simulation focusing on the Strange Situation? Thompson (1997) suggests that carer sensitivity when a child is fearful, anxious, or distressed may be more prognostic for attachment status as originally defined by Bowlby (1969) and used by Ainsworth *et al.* (1978). Therefore the **GLA** architecture will be the preferred architecture to use in the next chapter.

3.6 Conclusion

The research objective for this chapter was to design infant architectures that would form Secure or Insecure states depending on the responsiveness of the care that they received. For an infant, learning how reliable their carer is in ensuring their safety is of great importance. Infants rely upon their carers for nearly all their needs, but infants are also driven to explore objects and learn about other agents. In the EEA an infant architecture could perform optimally when it possessed an accurate measure of its security. This means it could maximise its learning opportunities whilst keeping within an acceptable level of risk. Two infant architectures are proposed that both satisfy the requirement to balance learning and security. Both the **GLA** and **GLS** architectures represent the infant agent's confidence in the carer agent's responsiveness with a single numerical parameter termed the Safe-range distance. This parameter is a less rich variant of the concept of felt security, but is expected to act similarly. Infant agents that possess both the **GLA** and **GLS** architectures learn that if their carers provide a lot of security they can allow the carer to move further before they attempts to regain proximity. We can view these infant agents as leaving their security in the hands of their carers. An infant agent that has learnt that its carer is unreliable needs to be proactive in maintaining its own safety, leaving fewer resources for learning about its environment. An infant that has low confidence in its carer will have a low Safe-range distance.

Since both the **GLA** and **GLS** architectures both reproduce the requirements specified in the SECURE-BASE scenario, can we select between them to find a preferred model? The **GLS** architecture is a better representation of affiliative interaction, whereas the **GLA** architecture is a better candidate for inclusion in the simulation of the Strange Situation which is described in the next chapter. This is for two reasons. Firstly, the **GLA** architecture possesses the potential for positive feedback loops to operate on it which may explain the distribution of infants into three clusters of attachment types which is found empirically. In addition, the Strange Situation Experiment puts infants into an environment that is expected to activate anxiety. Thompson (1997) makes the point that empirical studies that measure carer responses to bids deriving from anxiety are likely to be more prognostic of attachment style. Since the **GLA** measures exactly these type of events for infant agents it is the preferred architecture to simulate the Strange Situation in the next chapter.

Chapter 4

Different routes to safety

> *"If the organisms carries a 'small scale model' of external reality and if its own possible actions within its head, it is able to try out various alternatives, conclude which is the best of them, react to future situations before they arise, utilise the knowledge of past events in dealing with the present and future, and in every way react in a much fuller, safer, and more competent manner to the emergencies which face it"* (Craik 1943)
>
> *"The avoiding reactions tend to spread ... It can be carried to such a point that the individual is not only 'steeled against' the appeal and suffering of others, but he actually dreads appealing to their sympathy, and may, for example, conceal illness for fear of making a 'fuss' or 'scene'."* (Suttie 1935)

4.1 Introduction

This chapter is concerned with explaining how and why the three main infant attachment styles that can be observed in Strange Situation studies arise. In particular, alternative explanations are considered for the Avoidant pattern of attachment behaviour in reunion episodes of the Strange Situation Experiment. Two novel architectures are described which have been implemented in agent simulations. Each of these architectures is intended to reproduce the patterns that have been observed in infants during the laboratory and earlier home observations recorded in Strange Situation studies. These architectures vary in whether and how the infants represent themselves in their environment. As the quotes above suggest, this chapter is concerned with how infants represent the nature of their relationship with their main carer. This is an important issue. How we represent our first relationship may act as a template for relationships throughout our lives[1].

The structure of this chapter follows a similar methodology to that in the two preceding chapters:

Stage 1 - **The problem is established**: behaviour in the initial, exploratory Strange Situation study is abstracted and formed into the REUNION scenario. This scenario duplicates the structure of the empirical studies by having two stages, a long home stage and a short terminal testing stage based upon the Strange Situation Experiment. An important aspect of the REUNION scenario is that it describes three different strands of infant development from initially identical infant agents.

Stage 2 - **The function of behaviour is analysed from an evolutionary perspective.** Studies used to help assess the function of the behaviour of the three main attachment styles include those which describe: physiological measurements; cross-cultural differences in attachment; and the prevalence of infanticide in primate species, traditional societies and modern societies.

Stage 3 - **Two solutions are proposed and implemented**: both of these architectures are extensions of the **GLA** architecture described in chapter four.

- The **reactive-action-learning (RAL)** architecture behaves similarly to architectures described in previous chapters. However, it possesses two additional features. First is a perceptual subsystem that senses the quality of interaction during close physical contact and activates the goal of contact avoidance. In addition, this architecture possesses a selection mechanism that is not ‘winner-take-all’. Rather, if two highly activated goals are incompatible with each other a third, and less salient, action may be activated.
- The **hybrid-action-reasoning (HAR)** architecture possess two subsystems that can activate actions in the motor system. One is a reactive goal selection system similar to that found in previous architectures.

---

[1] In a longitudinal study Sroufe *et al.* (1999) found that by preschool age, Secure infants tend to have become more popular, less isolated and assessed as more competent by teachers. Goldberg notes that *“While 4 year olds with secure histories were rarely either victims or victimisers in bullying relationships, victimisers most often had avoidant histories and victims were most often in the resistant group”* (page 177, Goldberg, 2000). In adulthood: attachment security is linked to being an emotionally autonomous parent who values intimate relationships and the effects of those relationships; ambivalence is linked to being emotionally preoccupied and entangled in the details of early experiences; whilst avoidant emotional responses in infancy are linked to emotionally dismissing adults, similar to, though mostly less extreme than the type of person that Suttie refers to in the second quote above (Goldberg, 2000; Hesse, 1999).

The second route to actions being activated within the **HAR** architecture is from an additional deliberative subsystem which can present new actions and inhibit the existing actions resulting from the reactive selection mechanism. The **HAR** architecture does not possess an additional goal activator for the goal contact avoidance. However, it does possess a perceptual system that provides the same information about previous experiences of close contact to the deliberative subsystem.

Both of these solutions should be capable of forming the three different types of infant response, depending on the training that the infant agent receives before it undergoes its test separation. The motivation for inclusion of two architectures, where one (the **RAL** architecture) is purely reactive and the other (the **HAR** architecture) includes deliberation, is partly derived from consideration of LeDoux's theory of emotional architectures. LeDoux (page 164, 1996) characterises two different routes to emotional appraisal from the sensory thalamus to the amydgala. These are labelled as a 'low road' and a 'high road'. The low road is a 'quick and dirty processing pathway' that bypasses the cortex and only provides a crude representation. The high road is a slower route to action via the cortex. Although it is unlikely that the same circuits that LeDoux details are what direct infant attachment behaviours, the **RAL** and **HAR** architectures do provide contrasting 'markers' that direct further development. Future work, described in section 5.3, surveys intermediate cases between the **RAL** and **HAR** architectures.

Stage 4 - **The two architectures are evaluated**: Due to time constraints a comprehensive exploration of the design space of the two architectures has not been undertaken. The principal form of evaluation is to assess how these architectures satisfy the REUNION scenario and how they would support a variety of types of infant competence not required by the REUNION scenario, such as how infants search for objects and problem solve.

4.2 Framing the problem

4.2.1 Selecting the source material

The REUNION scenario described in this chapter is concerned with capturing patterns of attachment behaviour that can be observed with infants, at home, and in the laboratory conditions of the Strange Situation experiment. The original

Strange Situation study observed the laboratory behaviour of infants who had previously undergone 72 hours of home observations over the preceding year. Analysis of many subsequent Strange Situation studies that followed this pattern have shown that *"the maternal impact on the infant-mother attachment relationship has been shown to be much larger than the impact of child characteristics such as temperament"* (page 252, van Ijzendoorn and Bakermans-Kranenburg, 2004). Therefore, the REUNION scenario has been simplified by assuming that the carer's behaviour is the single source of individual differences in the infant agent's home and laboratory behaviour. This format means that the carer agent's behaviour is treated like an Independent variable in a controlled experiment. This is because in computational experiments using the simulation, it is the carer agent's pattern of behaviour that is varied by the programmer running the simulations. The infant agent's home and laboratory behaviours are treated as dependent variables. In the simulation, setting the carer agent's initial behaviour leads to the patterns found in the infant agent's laboratory behaviour and in the infant agent's home behaviour. This arrangement is clearly a simplification as it ignores any effect that the infant might have on the carer, but van Ijzendoorn and Bakermans-Kranenburg (2004) suggests that it captures the major causal factor.

As was noted in the last chapter, what an infant experiences over the eight short episodes of the Strange Situation is intended to activate and intensify infant attachment behaviours. Each infant taking part experiences approximately the same environment, and the infant's responses are recorded by video through a two way mirror. The eight episodes, which each last three minutes, are:

Episode 1 - mother and infant introduced to unfamiliar room;

Episode 2 - mother is nonparticipant while infant explores;

Episode 3 - a stranger enters;

Episode 4 - mother leaves but stranger remains, (the first separation episode);

Episode 5 - mother returns and stranger leaves, (the first reunion episode);

Episode 6 - mother leaves infant on its own, (second separation episode);

Episode 7 - stranger returns;

Episode 8 - mother returns and stranger leaves, (second reunion episode).

The unfamiliar surroundings are an initial instinctive cue to danger. The presence of a stranger may elicit further stress. A major jump in anxiety occurs during the first separation, when the mother leaves the infant in the room in the presence of the stranger. The second separation episode raises infant anxiety levels higher still, where the infant is left on its own in the room. It is at this point, after anxiety levels have been raised and then dampened and then raised again, that the mother reappears in the second reunion episode.

Meins (1997) sketches the range of behaviours that can occur in separation and reunion episodes of the Strange Situation experiment:

> "*A mother and her child are in an unfamiliar room together. The child is about a year old and plays with some toys on the floor as the mother reads a magazine. An unfamiliar woman enters and starts to chat to the mother and then tries to play with the child. After a short time, the mother gets up and leaves her child alone with this stranger. What will the child do now? What should the child do now? If the mother returns to the room after a few minutes how will the child respond?*" [After the separation from its mother the child might:] "*cry, throw a tantrum or become anxious, attempt to follow the mother out of the door or carry on as if nothing had happened. In observing infants in naturalistic settings and in this laboratory-based Strange Situation, researchers have witnessed all of these reactions; but taken in isolation, the infant's reactions to such a separation tell us remarkably little. The crucial issue for attachment research is how children react on being reunited with their mothers.*"(page 1 Meins (1997))
>
> "*Consider children who cry when their mothers leave the room: some of them will be consoled immediately by the mother's presence, others will need to be picked up and hugged before they can be comforted; but some children will cry more angrily when their mothers return and when they are picked up to be comforted may even strike out and stiffen or squirm on their mother's arms. Children who attempt to follow their mothers may also react in a number of ways on reunion: they may show a desire for physical contact and closeness by approaching her or clambering onto her knee; alternatively, they may be content with merely greeting her and continuing to interact with her at a distance. Even children who seem oblivious to their mothers' leaving may surprise us in their reactions on being reunited. Some will carry on playing in much the same way, paying as little attention to the mother's return as they did to her exit, but others will immediately approach and want to be picked up.*" (page 2 Meins (1997)).

What we want to do is understand the architectural mechanisms by which these behaviours come about, and to form a theory (or number of theories) explaining the purpose, if any, of these behaviours for the infant. We are interested in how patterns of behaviour are formed by the interaction of the infant's cognitive architecture and its caregiving environment. The close matching of behaviour in the Strange Situation and the home studies has allowed the Strange Situation procedure in the laboratory to act as an indicator of the quality of the mother-child relationship that exists outside of the laboratory (Goldberg, 2000). However, the separation behaviours are not a clear guide in this regard. This is because separation behaviours in the Strange Situation laboratory setting are not fully predicted by the carer's and infant's behaviour in the home environment. But importantly, reunion behaviour in the Strange Situation is strongly predicted by the home behaviour of the mother and the infant.

Regardless of how they reacted in separation;

- the infants whose response to their mothers on reunion in the Strange Situation was: to not seek direct contact; to avoid their mother's gaze and avoid physical contact with her, *are described as insecure-avoidant and labelled type A*. These children return quickly to play and exploration but do so in a desultory fashion, with less concentration than secure children. Often this play or exploration will be in moderate proximity to their mother. For example, they may play with a toy close to the mother, but with their back turned to her. Whilst playing they stay close to and keep an eye on their carer. They received care at home which can be summarised as being consistently less sensitive. In comparison with average levels across all groups: A type mothers were observed at home being less emotionally expressive and having a greater aversion to close physical contact; they left infants crying for longer durations and provided more physical contact of an unpleasant nature; and at home these infants were more angry, they cried more and were observed to 'sink-in' less during physical contact.

- the infants whose response to their mothers on reunion was: positive, greeting, approaching, making or accepting contact with, or being comforted by her, *are described as securely attached and labelled type B*. These children returned to play and exploration in the room quickly. They received care at home which can be summarised as being consistently more sensitive. In comparison with average levels across all groups: B type mothers were observed at home being more emotionally expressive and having a smaller aversion to close physical contact; they left infants crying for shorter durations and provided less physical contact of an unpleasant nature; at home these infants were less angry, they cried less and were observed to 'sink-in' more during physical contact.

- the infants whose response to their mothers on reunion was: not being comforted and overly passive or showing anger towards their mothers, *are described as insecure-resistant/ambivalent and labelled type C*. These children do not return quickly to exploration and play. They received care at home which can be summarised as being less sensitive and particularly inconsistent. In comparison with average levels across all groups: C type mothers were observed at home being more emotionally expressive and having a smaller aversion to close physical contact; they provided physical contact which was unpleasant at a level intermediate between A and B carers and left infants crying for longer durations; at home these infant's were more angry, they cried more but were observed to 'sink-in' more during physical contact.

- the infants whose response to their mothers on reunion was: totally disorganised and confused, *are described as insecure-disorganised and labelled type D*. The home environment of behaviour for this very small proportion of infants has been found to be dysfunctional, often with depressed mothers or with maltreatment of the infant (Meins (1997), Ainsworth *et al.* (1978), and Weinfield *et al.* (1999)).

4.2.2 What needs to be explained?

The work in this chapter is intended to merge with the previously described architectures to form an integrated whole.

There are two respects that the simulations described in this chapter are partial, in that they do not incorporate all possible available data observed on infants. One aspect of modern Strange Situation results that this work will not be simulating is the presence of the disorganised category of infant attachment response. As was mentioned above, this category of infants forms a very small proportion of infants in non-clinical populations. In addition, it was altogether absent from the study by Ainsworth *et al.* (1978). Since this study is being used as the empirical foundation of the scenarios the simulation of disorganised category of attachment response is being left for future work. Another aspect is that the simulations in this chapter (in common with previous chapters) do not take into account how infant variables might determine attachment status.

The explanation of the Strange Situation that is described in this chapter therefore builds upon the work described in previous chapters, and ties up loose ends left in those chapters. However, this chapter will not duplicate the results of previous chapters and therefore does *not* need to explain elements of the infants behaviour in the Strange Situation that include:

- *how infants switch between competing goals:* in the Strange Situation all infants at some time elicit most of the types of behaviour that are recorded by the observers. Therefore it is likely that all infants cry and play at some time during the eight episodes. Even within each three minute episode, infants will show a range of different behaviours. What is important is the relative frequencies that particular behaviours are produced by the infants. This switching of behaviours is not a principal focus of explanation in this chapter because the **GS** architecture described in chapter three will support these patterns of behaviour.

- *how infants learn about the promptness of their carers in responding to their signals:* the **GLA** and **GLS** architectures provide differing accounts of how infants might assess promptness of their carers in responding to their signals and then use this information in estimating how much security their carers provide. Both the **GLA** and **GLS** architectures solve the SECURE-BASE scenario, producing infant agents with attachment styles that vary along a single dimension, from Secure to Insecure. When infants possessing the **GLA** and **GLS** architectures become Insecure they exhibit an increased tendency to protest at home *and* in unfamiliar environments.

What the SECURE-BASE scenario does not include, and what the **GLA** and **GLS** architectures do not form, is two types of Insecure infant. The Strange Situation categories include two main types of Insecure infant, but in the implementations previously described in this thesis there is only one type of Insecure infant. Therefore the focus of the problem that is set out in this chapter is how and why two categories of Insecure infant arise in Strange Situation studies. The kinds of source material that are of interest are those that help to differentiate the abstract generic insecurity found in the SECURE-BASE scenario into the two main Insecure categories of Avoidant and Ambivalent infants.

Differences in carer behaviour towards the infant are central in explaining how infants come to possess different styles of insecurity. As was noted earlier, the REUNION scenario treats the home caregiving experience of infants as an Independent Variable in the simulations. At home the caregiving experience of Avoidant infants resembles that of Ambivalent infants, they both share a similar general pattern of insensitive responses from their carers. The responsiveness of the caregiving that Avoidant and Ambivalent infants receive differs because the carers of Avoidant infants are consistently insensitive, whereas the sensitivity of carers of Ambivalent infants is inconsistent. In addition, the home behaviour of the carers of Avoidant infants is more often rejecting and close physical contact between Avoidant infants and their carers is more often unpleasant.

The REUNION scenario requires that both the home and laboratory behaviour of infants is caused by the type of caregiving they receive. The home behaviour of infants that later undergo the Strange Situation experiment are valuable sources of data because of critical continuities and discontinuities between home and laboratory behaviours in the three infant groups. Secure infants show a strong continuity between patterns of home and laboratory behaviour. For example, Secure infants show minimal levels of anger at home and in the laboratory. The behaviour of Avoidant and Ambivalent infants is similar at home, with both groups of infants crying more than Secure infants, and both showing more anger and resistance at home than Secure infants. It is in the reunion episodes of the Strange Situation that Avoidant and Ambivalent infants demonstrate really distinct patterns of behaviour. The behavioural patterns of Ambivalent infants might reasonably be predicted from their home behaviour. The high frequency of crying and anger at home is translated into high levels of anger and resistance in the Strange Situation. Many aspects of Avoidant infants' behaviour show continuities between home and laboratory. For example, Avoidant infants show the highest levels of avoidance in the Strange Situation. They also show high levels at home of measures related to avoidance such as: showing a negative response to being held, and low levels of measures that are the opposite of avoidance, such as sinking in and active contact behaviours. However, Avoidant infants exhibit a striking discontinuity between home and laboratory. Whilst the most angry at home, these infants are the least angry in the Strange Situation. Avoidance in the Strange Situation experiment is a public limiting of emotion, but is linked to strong emotions such as anger:

> *"Observations of infants and mothers in other settings show that ... avoidance is highly associated with the mother's anger, her emotional inexpressiveness, and her rejection of physical contact with the infant. Avoidance on reunion is also highly predictive of the infant's social and emotional behaviour in other situations. Infants who strongly avoid the mother and show no anger following separation are, for example, more likely to attack or threaten the mother in other settings."* (page 34 Main and Weston, 1982)

One interpretation of the general continuity between the home and laboratory behaviour of Ambivalent infants is that these infants use the same strategy of coercive signalling at home and in the laboratory conditions of the Strange Situation. This interpretation is consonant with how infant agents possessing **GLA** and **GLS** architectures become Insecure in simulation. When infant agents with the **GLA** and **GLS** architectures have their bids for a response from the carer ignored, they simply bid for attention more insistently. A similar process may be what occurs in real Ambivalent infants. However, this is not what occurs in real Avoidant infants.

When the carer re-enters the laboratory in the final reunion episode of the Strange Situation we might expect Avoidant infants to signal in an angry manner because this is how they often signal at home. That Avoidant infants do not signal angrily is a puzzle. One solution to this puzzle is that Avoidant infants are simply less stressed in the reunion episode than the other categories of infant. There are studies which have that have measured the physiological state of infants undergoing the Strange Situation experiment.

Any solution to the puzzle that relies upon Avoidant infants being less stressed is dispelled when results from these studies are considered. This is done in the next section, which is concerned with analysing the function of infant behaviour from an evolutionary perspective.

4.2.3 Analysis of function

4.2.3.1 Ruling out several hypotheses

A number of studies have attempted to measure the psychophysiological state of infants undergoing the Strange Situation (Spangler and Grossman, 1993; Hertsgaard *et al.*, 1995; Fox and Card, 1999). One of the things that these studies have attempted to investigate is how physiological systems act as correlates of observed behavioural responses (Fox and Card, 1999). It might be argued that Avoidant infants are using a different strategy to realise their goals in the Strange Situation. It may be that they actually possess different goals to Ambivalent infants. Rather than attempting to gain safety they may simply be less anxious and be temperamentally inclined to explore more. This explanation is weakened by results from studies which have made physiological measures of stress for the three infant attachment groups during and after they underwent Strange Situation Experiments. Heart Rate (HR) and cortisol studies of infants undergoing the Strange Situation show that Avoidant infants are at least as stressed at infants in the other attachment groups, despite the fact that they appear less stressed. For example Sroufe and Waters (1977b) found that:

> *"secure, avoidant, and resistant infants all showed increased HR's upon separation, which remained elevated during reunion. Secure infants' HRs recovered on average after less than 1 minute of contact with their mothers. After the secure infants were put down, they showed HR deceleration when they returned to play and attended to objects. Resistant infants, in contrast, requested to be put down before their HRs recovered to the preseparation level; then, after being put down with their HRs still elevated, they reached up to be held*

> *again. Avoidant infants showed an increase in HR from the beginning of separation until long into the reunion session, even though they outwardly appeared to be unaffected by the separation."* (page 231, Fox and Card, 1999)

In addition to physiological measures, very close observation of Avoidant infants has been undertaken that give some clues to the internal state of the infant. The difference between Avoidant and other infants in the reunion episodes of the Strange Situation is stark. There are more subtle differences in other environments. For example, avoidance in the home is linked to gaze avoidance and a pattern of slow closing of eye lids that does not occur in other infants. Details of these 'micro-behaviours' are not included in any of the scenarios within this thesis, but their exclusion is relevant to the issue of the structural accuracy of the simulation, which is discussed in section 5.2.5. Further details of 'micro-behaviours' are described in appendix D.4.

How should these results by interpreted in the context of assessing the function of infant behaviours from an evolutionary perspective? One curb to viewing this data as resulting from adaptations is that the patterns of behaviour may instead be cultural artifacts and not deriving from the fundamental species-specific architecture of human infants. This alternative explanation is undermined by studies that show that the three main attachment styles occur with varying frequencies across many different cultures (van Ijzendoorn and Kroonenberg, 1988)[2].

4.2.3.2 The evolutionary function of the Secure pattern of behaviour

Section 3.2.2 discusses an evolutionary analysis of Secure patterns of infant behaviour in more detail. A summary of the hypothesis put forward there is that parent-offspring conflict gives rise to selective pressures that are ubiquitous in the evolution of sexually reproductive animals, though it can sometimes be very well hidden, and often beyond the conscious awareness of the parents or offspring it affects. Evolving in conditions of parent-offspring conflict means that sometimes the infant's and carer's interests will not coincide. A carer may want to expend its time and other resources in alternative ways than caring for an infant. According to the hypothesis put forward in section 3.2.2, infants possess mechanisms that detect if their carer is not attending adequately to their safety. These mechanisms have evolved in the EEA. If the carer is assessed and adequately responsive then the infant will be predisposed to explore more and signal less. These infants are termed Secure infants, and they have been represented in the last chapter in the

---

[2] see appendix D.5 for details of cross-cultural research

SECURE-BASE scenario and reproduced in simulation with the **GLA** and **GLS** architectures.

4.2.3.3 The evolutionary function of the Ambivalent pattern of behaviour

Alternatively, if the carer is assessed as not adequately responsive then one solution is for the infant to signal louder and more insistently, perhaps with a qualitatively different signal that includes angry protest. This pattern of behaviour matches real ambivalent infants at home and in the reunion episode of the Strange Situation. It also matches the variant of Insecure infants specified by the SECURE-BASE scenario and implemented with the **GLA** and **GLS** architectures.

Cassidy and Berlin (1994) review a large number of studies in an attempt to summarise data and theory concerning infants possessing an Ambivalent style of attachment. They describe the function of an Ambivalent style in the following way:

> *"In response to a parent who is minimally or inconsistently responsive, the insecure/ambivalent infant is thought to develop an understandable strategy of increasing his bids for attention. The insecure/ambivalent is viewed in this scheme as having a coherent strategy of exhibiting extreme dependence on the attachment figure. Main and Solomon (1986) state that "in its heightened display of emotionality and dependence upon the attachment figure, this infant successfully draws the attention of the parent" (p. 112)"*
>
> *"This strategy of heightened attachment behaviour can be viewed in both positive and negative lights. It is adaptive because it served to increase or maintain proximity to the attachment figure, which in turn increases the infant's chances of survival (Bowlby 1969). This strategy, however, is not wholly satisfactory. One drawback is that it requires the infant to take on "more than optimal responsibility for maintaining contact" (Bretherton, 1985, p. 11), and thus prevents the infant from attending to his own developmentally appropriate tasks. The infant's subjective experience is also clearly less than optimal. For instance, fearfullness may result from limited familiarity and success with the environment. Moreover, the infant who must resort to extremes of affective signalling with unpredictable success in gaining the parent's response is poorly equipped to understand and organise his affective experiences."*

Ambivalent infants rely upon their carers to provide comfort and security to a greater extent than do Secure or Avoidant infants. Cassidy and Berlin (1994) observe that a strategy of emphasising immaturity and reliance on an ‘other’ to fulfil one’s goals is a general strategy that can be observed in human contexts beyond infancy, such as in adult romantic relationships.

4.2.3.4 The evolutionary function of the Avoidant pattern of behaviour

Ambivalent and Avoidant infants show the same pattern of increased signalling and angry protest at home. Since Ambivalent infants also protest in the Strange Situation, a simple explanation for Ambivalent patterns of behaviour is that Ambivalent infants have learnt a general policy of increasing protest. However, this account does not explain Avoidance in reunion episodes in the Strange Situation. There is an additional factor that is causing Avoidant infants to not signal as they would at home, but instead to signal less and move themselves to a closer proximity to their carer. This kind of avoidance is not represented in the SECURE-BASE scenario and cannot be produced by the **GLA** and **GLS** architectures without modifying them.

There exists a body of evidence that suggests that the rationale of parent-offspring conflict impacts upon present-day caregiving in some modern technologically advanced societies in the form of abuse, neglect, sub-optimal parenting and infanticide in the form of abortion and in modern hunter-gatherer societies as infanticide of newborn infants (Hrdy, 1999). Soltis (2004) reviews a number of sources of cross-cultural evidence to demonstrate that parent-inflicted infanticide occurs overwhelmingly when the child is unlikely to survive. Soltis has also conducted an evolutionary analysis of excessive crying in early infancy and concludes that the function of the mechanism for which excessive crying is a likely side-effect is the infant proving its vitality and thus avoiding infanticide. Soltis presents the rationale for infanticide from the carer’s perspective using two differing ways of talking about the affect of evolutionary factors on behaviour:

> *“When the chances for infant survival are low, parents may increase overall reproductive output by terminating parental investment in offspring (i.e. engage in adaptive infanticide), which should be taken as shorthand for Natural selection has favored in parents a psychological mechanism that results in the withdrawal of care from their own offspring, this mechanism is activated under conditions in which the infant is unlikely to survive, and selection has favored this flexibility in parental behavior because in past environments the selective withdrawal of care under inauspicious circumstances resulted, on average, in the production of more offspring over a lifetime.”* (page 445, Soltis, 2004)

If we make the assumption that similar conditions prevailed for nine to twelve month old infants in the EEA, this would have provided a selective pressure for the evolution of infant strategies to minimise the probability of infanticide or other negative consequences resulting from carer abuse or negligence. When conditions were harsh so that an infant was unlikely to survive, infanticide and neglect probably increased in likelihood. Infants that could adapt their behaviour to mitigate the risks they were under would be better adapted and show better inclusive fitness. Hrdy (1999) devotes a chapter to the issue of *"How to be an infant worth rearing"* (pages 459-474, Hrdy, 1999), which details many examples, from modern hunter gatherer societies and historical records, of the calculations that are made by carers when human infants are abandoned or killed.

It is important to remember that this line of argument is not suggesting that any widespread patterns of present day caregiving are predictive of infanticide. The point is merely that infants may possess mechanisms that assess modern patterns of behaviour this way. This hypothesis is not that type A carers are likely to kill their infants because they hold the infants awkwardly and unpleasantly. The hypothesis is that infants possess mechanisms that were selected for in an environment when these kinds of behaviour did predict harm. In the EEA there was a good match between unpleasant holding and harmful parents. In the present-day there is still infanticide but forms of holding are not a good prediction. Why should forms of holding lose their predictive power, for the infant? Another way of asking the same question is: why are forms of holding no longer correlated with forms of maternal behaviour? There are a number of different possible causes which are discussed in appendix E.2.

4.2.4 Two theories that explain distinctions in Insecurity

The principal cause of Insecurity in infants is maternal insensitivity. However, what are the causes that differentiate Insecure Avoidant infants from Insecure Ambivalent infants?

According to Ainsworth *et al* (1978), Avoidant infants act avoidantly in reunion episodes to avert close physical contact because they receive more physical contact of an unpleasant nature at home and are more likely to have bids for attention rejected at home. Observations that are statistically significant ($p < 0.05$) which support this hypothesis include that the carers of Secure infants are significantly more tender, affectionate, responsive to crying and responsive in reunions, and less inept in close contact than the carers of Avoidant and Ambivalent infants. In many measures of close contact the carers of Avoidant and Ambivalent infants cannot be distinguished. However, the carers of Avoidant infants do provide significantly more unpleasant and interfering contact, and are more averse to close contact and show a greater lack of emotional expression and emotional rigidity than the other two types of carer. Avoidant infants are significantly more angry and resistant than Ambivalent infants, who are in turn significantly more angry and resistant than Secure infants at home[3].

Cassidy and Berlin (1994) suggests that a key factor in infants becoming Ambivalent is the inconsistency of their caregiving:

> *"It is thought that the Strange Situation behaviour of insecure-ambivalent infants reflects a working model of the caregiver as inconsistently available, where even on reunion these infants continue to act as if she were unavailable. This uncertainty about maternal availability is thought to result neither from consistent maternal availability nor consistent maternal unavailability, but rather from* **inconsistent** *maternal availability"* (page 927, Cassidy and Berlin, 1994)

This is an interesting hypothesis because it ties in with learning theory, which suggests that partial reinforcement schedules result in greater resistance to extinction because the conditions under which an animal learns are similar to the conditions of extinction (Anderson, 1995). Interpreting this finding for the Strange Situation would suggest that ambivalent infants always respond with higher levels of protest because this *may* trigger a response from the carer. However, basing a simulation upon this theory is limited by the fact that:

> *"the theoretically based proposition that inconsistent maternal availability is associated with insecure- ambivalent attachment has not been tested directly. The only research examining maternal behaviour has focused on means levels of behaviour. Thus, mothers who are sometimes available and sometimes unavailable will appear, on average, less available than those who are consistently available"* (page 927, Cassidy and Berlin, 1994)

---

[3] All 82 measures of infant and carer behaviour are presented in section D.1. Section D.2 presents further details of statistical relationships between home and Strange Situation carer and infant behaviours particularly relevant to the distinctions made in the REUNION scenario. The four coding measures that provided the clearest discrimination between attachment groups in reunion episodes were: proximity seeking, contact maintaining, avoidance, and ambivalent behaviour. Details of the scoring criteria for these measures of infant behaviour are presented in D.3

4.2.5 Scenario formation

The REUNION scenario is intended to capture the aspects of an infant's home experiences that cause it to develop Avoidant, Secure or Ambivalent styles of attachment behaviour. Each simulation is split into a training period, corresponding to the lengthy home observations, and a test period, corresponding to the much shorter Strange Situation assessment. The style of behaviour that the carer is set to reproduce can be thought of as an Independent variable in an experiment (or as input data to a function). The Dependent Variables (or output data from a function) are the infant's behaviour, at home and in the Strange Situation stage of the simulation. A mapping between the input data and output data emerges from the dynamic interaction of infant and carer agents in a similar 2-dimensional virtual world to that in which the previous simulations occured. In the training period infants explore the space of behaviours open to them, attempting to optimise their behavioural 'policy' in terms of plausible adaptive benefits, and these learnt 'policies' are carried forward to the Strange Situation stage of the simulation.

The scenario has a conditional structure. If carer agents act in one of the three main forms of caregiving then the infant agent's behaviours is required to match.

In the previous chapter the **GLA** and **GLS** architectures both reproduced Secure and Insecure forms of attachment. The REUNION scenario is based upon using the **GLA** architecture because it was this architecture that provided the best explanation of the differentiation in to secure and insecure categories that is observed empirically (Fraley and Spieker, 2003).

The starting point for abstracting source material from Strange Situation studies was the existing ontology of the SECURE-BASE scenario. In common with the SECURE-BASE scenario, the REUNION scenario also incorporates a **response-threshold** within the care which means that the carer only responds to the infant when the infant signals above this level. Which objects, properties or relationships need to be added to the ontology of the SECURE-BASE scenario to capture all the aspects of behaviour that are needed for the REUNION scenario? Additional abilities include the carer having the ability to make different styles of close bodily contact and the infant having the ability to sense the different styles of close bodily contact performed by the carer.

All the previous scenarios are prerequisites for the REUNION scenario. The empirical studies of the Strange Situation show that in reunion episodes nearly all infants do some crying and in separation episodes nearly all infants do some exploration. Behaviours can switch quickly and may oscillate. This state of affairs needs to be reproduced as a required outcome. The scenarios will not be in the form of scripts that describe precise strings of actions. It is the frequencies and durations of actions and intentions over an extended period that have to be matched.

The REUNION scenario does not capture the most subtle aspects of avoidant behaviour, such as the slow closing of eye lids mentioned earlier (the infant agents do not possess eye lids!). The importance of these omissions are dealt with in section 5.2.5 of the conclusion.

An important aspect of the REUNION scenario is that it does not describe differences between the different groups as completely stark 'all or nothing' differences. All infant agents and all carer agents can show all the different types of behaviour at some time. What the REUNION scenario requires is that infant and carer agents belonging to the different groups produce particular behaviours in the approximately correct relative proportions.

A summary of the Avoidant patterns of carer agent and infant agent behaviour is:

- **If** at home, the carers of Avoidant infants consistently reject infant signals indicative of a desire for closeness and when they do make physical contact it is more often unpleasant.
- **then** at home Avoidant infants show more angry protest, they cry more frequently and for longer durations.
- **and then** in reunion episodes of the Strange Situation, Avoidant infants show little angry protest

A summary of the Secure patterns of carer agent and infant agent behaviour is:

- **If** at home, the carer of an infant responds in a timely fashion to more communications from the infant; and provides physical contact of a pleasant nature
- **then** at home, the infant will communicate with less intense negative tone, show less angry protest, crying less frequently and for shorter durations, and be more rewarded by close physical contact.
- **and then** in reunion episodes of the Strange Situation, the infant will show some distress but get back to play quicker and with more attention in its play.

A summary of the Ambivalent patterns of carer agent and infant agent behaviour is:

- **If** at home, the carer of an infant frequently reject signals indicative of a desire for closeness, but does so inconsistently; and this carer varies in the quality of physical contact they provide, but does not provide the most unpleasant contact.
- **then** at home the infant will show more angry protest, they cry more frequently and for longer durations. This infant will be more rewarded by close physical contact than are infants who receive consistent contact of a more unpleasant nature.
- **and then** in reunion episodes of the Strange Situation this infant will show particularly resistant and angry behaviour.

A big problem with the data on the Strange Situation is that we don't know how many times the infants have experienced similar situations to the Strange Situation previously. There is some evidence from cross-cultural comparisons. For example, studies have looked at Japanese, Italian and German infants that have different home experiences and found that the frequencies of the three main attachment styles change according to home experience (van Ijzendoorn and Kroonenberg, 1988). This evidence shows that in cultures where caregiving tends to be closer and more intense, with less separations and more bodily contact, the proportion of avoidant infants drops. In samples where separations are more frequent the proportion of avoidant infants is higher. However, we don't know how many reunions in unfamiliar surroundings the infant has previously undergone. There is evidence that most infants that undergo the Strange Situation have not undergone similar experiences. In particular, the strongest evidence for this is that the Strange Situation cannot be repeated within a period of a few months. If it is repeated infants become very distressed. From this we can deduce that infants had not experienced lots of situations in their normal daily life that were very similar to the Strange Situation, otherwise they would show more distress in the Strange Situation. However, we simply do not now how many marginally similar situations the infants face. Therefore the REUNION scenario will make the somewhat arbitrary decision that it should be able to respond appropriately whether it has undergone similar separations in unfamiliar surroundings, or whether it has not.

The REUNION scenario can be viewed as requiring the infant to perform a kind of problem solving. Therefore an obvious kind of integration is to extend the REUNION scenario to incorporate other forms of problem solving that infants perform at this age. As with previous scenarios the modes of evaluation include graphically observing the running simulation and producing trace print outs.

4.3 The solutions

4.3.1 Rich interpretation vs. deflationary accounts of the REUNION scenario

In the study of cognitive development there is a dispute which Munakata *et al.* (2002) describes as between *"Rich interpretation vs. deflationary accounts"* of infant behavioural phenomena. Rich interpretation involves *"casting simple responses in terms of overly complex abilities and cognitive processes"*, whereas deflationary accounts involve *"underestimating the abilities of preverbal and nonverbal populations ... mistakenly casting thoughtful behaviours in terms of simplistic processes"* (page B43, Munakata *et al.*, 2002). Bowlby's (1969) account of how the information processing related to attachment behaviour develops predates this dispute and therefore does not take a position in it. Bowlby described the development of processing related to attachment in childhood in four stages:

- The reflex phase of attachment: birth to a few months.
- The fixed action pattern phase of attachment: a few months after birth towards the middle and end of the first year.
- The goal corrected phase of attachment: middle to end of the first year to approximately 3 years.
- The goal corrected partnership phase of attachment: from approximately 3 years to adulthood.

This analysis seems clear, however Bowlby's description of goal corrected (goal directed) mechanisms is open to either rich interpretation or a deflationary account. This is because goal directed behaviours may occur by trial and error, or by means-end analysis[4]. This work will term a trial and error mechanism a reactive solution. If a mechanism allows any form of means-end analysis it

---

[4] The distinction between the types of information processing requirements of trial and error learning and means-end analysis has also been described by other infant psychologists and considered by a number of contributors to AI and Cognitive Science. Infant psychologists include: Piaget *transitional behaviour vs. intentional behaviour* (Piaget, 1953, 1955); Willatts *transitional behaviour vs. means-end reasoning*, Munakata *trained behaviour vs. means-end reasoning*, (Munakata *et al.*, 2002), Baillargeon *reactive vs. planful* (Baillargeon *et al.*, 1990). Contributors to AI and Cognitive Science include: Dennett - *Contrasts Darwinian (reflexive) and Skinnerian (Associative) organisms with Popperian organisms 'that by criticising their theories, let their theories die in their stead'. (Dennett, 1995, 1996; Popper, 1997))*; Sloman, - *reactive vs. deliberative processing* (Sloman, 2001, 2002); Newell - *preparation vs. deliberation* (Newell, 1990)

will be termed a deliberative solution[5] In either case there is a goal, so both cases fall under Bowlby's definition of goal directed. In both reactive learning and deliberative analysis infants may possess representations of their initial state and the goal state that they desire. The distinction is that infants using reactive learning don't need to internally represent action operators that move from one state to another in some form of mental model or mental simulation. Infants using reactive mechanisms can try an operator, either at random or according to some inherent ordering of responses. Then if this operator succeeds they can store this for future use in similar circumstances.

Two architectures have been implemented: the **RAL** and **HAR** architectures. Qualitative aspects of these architectures behaviour will be discussed in the next two sections. Then in section 4.3.4 some preliminary quantitative results for the two architectures is presented.

4.3.2 The RAL architecture

The first implemented solution is a reactive architecture labelled the **RAL** architecture, and which is based upon the **GLA** architecture. This has been implemented using explicit representations of the goals at which it is directed. However, in a real infant the representation of goal states may only be implicit representations of internal states that are used in controlling actions, but do not endure and cannot be used for more than one purpose, nor survive after their immediate triggering[6].

Key aspects of infant behaviour in the REUNION scenario to be explained by the **RAL** architecture include: why Avoidant infants switch from their normal high level of protest at home to a minimal level in reunion episodes of the Strange Situation. In addition, any architectural solution must explain why Avoidant infants in reunion episodes return to exploration but do so in a desultory fashion. It was noted earlier that several contributors to Attachment Theory have compared patterns of avoidance to a phenomenon that ethologists have termed 'displacement activity' (Hinde, 1983). The most noticeable aspect of examples of displacement activities in animals can be seen when an animal does not produce expected salient behaviours, such as fleeing or fighting, when confronted by an aggressive con-specific. Instead animals which produce displacement activity carry out behaviours that seem out of place, such as preening. The **RAL** architecture provides a high level explanation of avoidance as a displacement activity, where the Avoidant infant agent's desultory exploration and lack of protest can be

---

[5] However, it should be emphasised that simple means-end analysis sits at one end of a spectrum of deliberative mechanisms that become more complex.

[6] Sloman unpublished report - *'What the brain's mind tells the mind's eye'*

compared with an animal preening in the face of an aggressor.

The **RAL** architecture is based upon the **GLA** architecture, and in common with the **GLA** architecture it possesses a number of perceptual subsystems which activate goals. However, to support displacement phenomenon the **RAL** architecture needs to possess components that are not possessed by the **GLA** architecture. These are an Avoid-pain goal activation module, and a more sophisticated resource allocation and arbitration module. The additions allow a greater repertoire of possible actions than the **GLA** architecture allows.

Figure 4.1 shows that each goal activator module has in input derived from three sources. These are a component of its activation set by external perception of the world; a component of its activation from internal sources (simulating the effects of residual activation and the effects of hormones); and information about the external target object or agent to which the goal is directed. The arrows in figures 4.1 represent continually updated information flow that can be viewed as similar to transfer of electrical current. The avoid pain subsystem operates in a similar manner to the other goal activator modules. The key difference between the selection mechanism in the **GS** architecture and the **RAL** architecture is that, in the **RAL** architecture there is not a winner-take-all mechanism. The effect of the winning goal is modulated by the effects of other goals, if these are also high in activation.

Figure 4.2 shows an instance of time during a reunion episode for an Avoidant infant agent that possesses **RAL** architecture. Both the goals of attachment (labelled the **attach** subsystem) and pain avoidance (the **avoid** subsystem) produce high levels of activation that ‘cancel each other out’ leaving a more weakly activated goal of exploration to direct the next action. An infant that possessed a **RAL** architecture in the state represented in figure 4.2 would have experienced a history of unpleasant close physical contact to cause the **avoid** subsystem to become predisposed to producing high activation. The avoid pain subsystem’s activation is sensitive to the distance that the carer is from the infant. In separation episodes its activation is zero. When the carer is far its activation is minimal and it may not trigger displacement behaviour. When it is close it has a greater inhibiting effect on infant signalling than on infant movement, so that the infant may continue moving towards the carer without signalling. The **avoid** subsystem also operates in the training stage of the simulation that represents the home stage of the actual Strange Situation studies. The operation of the **avoid** subsystem in the infant agents supports avoidance in the training stage of the simulation that is comparable to data for real avoidant infants showing high levels of negative response to being held, and low levels of measures that are the opposite of avoidance, such as sinking in and active contact behaviours. When the infant is very close to the carer it produces activation that is high enough to inhibit signalling

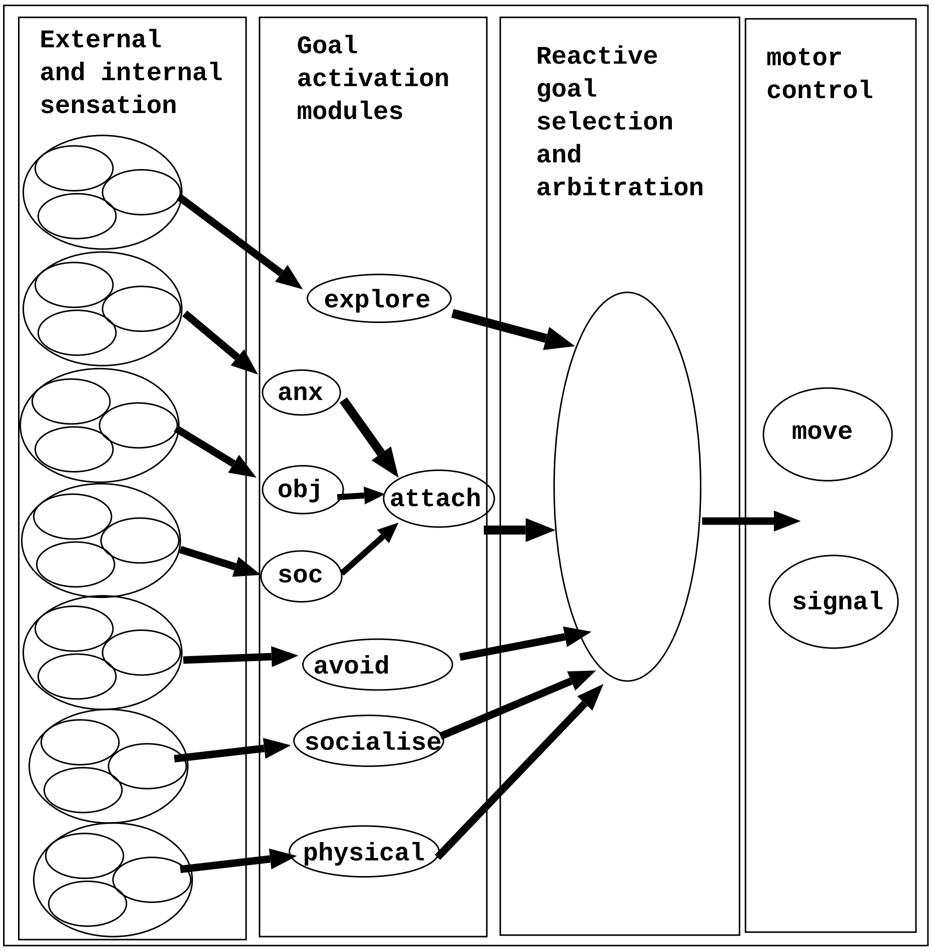

Figure 4.1: The **RAL** architecture. There are seven perceptual subsystems that input to goal activators. These correspond to the goals for exploration (explore), anxiety (anx); object wariness (obj); social wariness (soc); socialisation (socialise); and physical need (physical). The reactive goals of anxiety, object wariness, and social wariness are all combined in a single goal of attachment (attach) before they are considered for selection. The reactive goal selection mechanism is not a 'winner take all' mechanism. The goal activators with highest activation may mutually in inhibit each other leaving a less highly activated goal to direct behaviour.

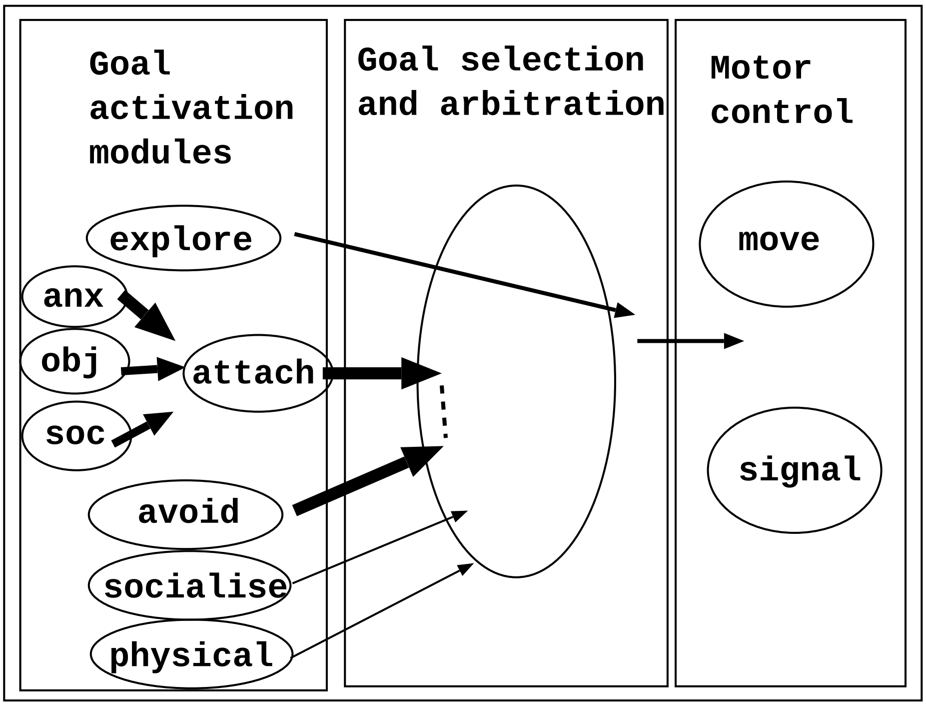

Figure 4.2: A snapshot of the **RAL** architecture within an infant that has learnt to be Avoidant. The dashed line between the **attach** and **avoid** arrows signifies mutual inhibition which is activated because these goals are both highly activated and they make mutually exclusive demands on the motor system. (Some details of the perceptual system omitted)

and moving to the carer. The inhibition of signalling and movement to the carer leads to displacement activity in the form of exploration. Figure 4.2 illustrates the moderate activation of the exploration goal winning out and directing the infant agent's motor system.

Figure 4.3 shows an instance in time during a reunion episode for a Secure infant agent that possesses the **RAL** architecture. At the very beginning of a reunion episode a Secure infant agent would activate the **attach** goal and move and signal to the carer agent. However, during most of of the reunion episode the goal to explore is most highly activated. This state of affairs would come about when in its previous home-training experiences the infant agent received a high level of carer responsiveness so that its confidence in the carer's ability to provide safety, in the form of its safe-range parameter, is very high. In addition, the infant agent would not have received close physical contact of an unpleasant nature, so that the avoid pain subsystem would provide minimal activation. When the carer returns

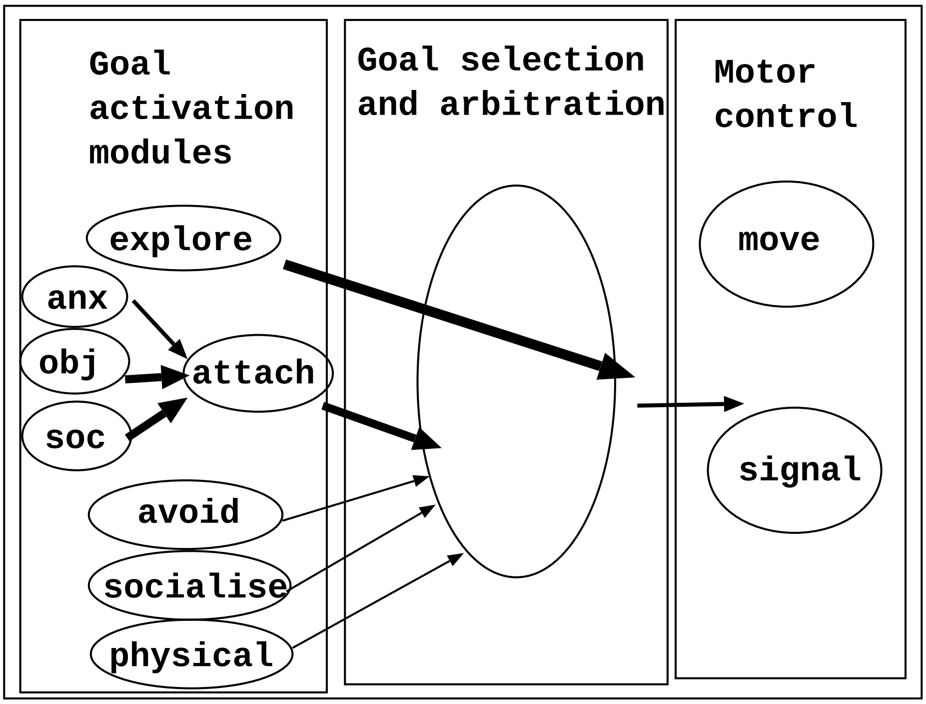

Figure 4.3: The **RAL** architecture when it has learnt to be Secure. At the very start of the reunion episode, the attachment goal (attach) is high. After reassurance by the carer this goal activation lowers, leaving the goal of exploration to activate the motor system. (Some details of the perceptual system omitted)

after the separation episode the infant's anxiety level quickly subsides because of its high confidence level. At no time does the **avoid** subsystems's activation effect other processes.

Figure 4.4 shows an instance in time during a reunion episode for an Ambivalent infant agent that possesses the **RAL** architecture. For much of a reunion episode, including the beginning, only the **attach** goal is highly activated. Ambivalent infants have experienced lower levels of carer responsiveness, hence the high level of anxiety. However, they also experience lower levels of unpleasant close physical contact, hence the low level of activation for the avoid pain subsystem.

The **RAL** architecture explains why the behaviours in the reunion episodes are indicators of the quality of the infant-carer relationship beyond the laboratory. Reunion episodes provide Anxiety activations that are just low enough for the balancing effect of the Avoid-pain subsystem to be seen. Insecure Avoidant infants agents differ in their reunion behaviour from Insecure Ambivalent infant agents

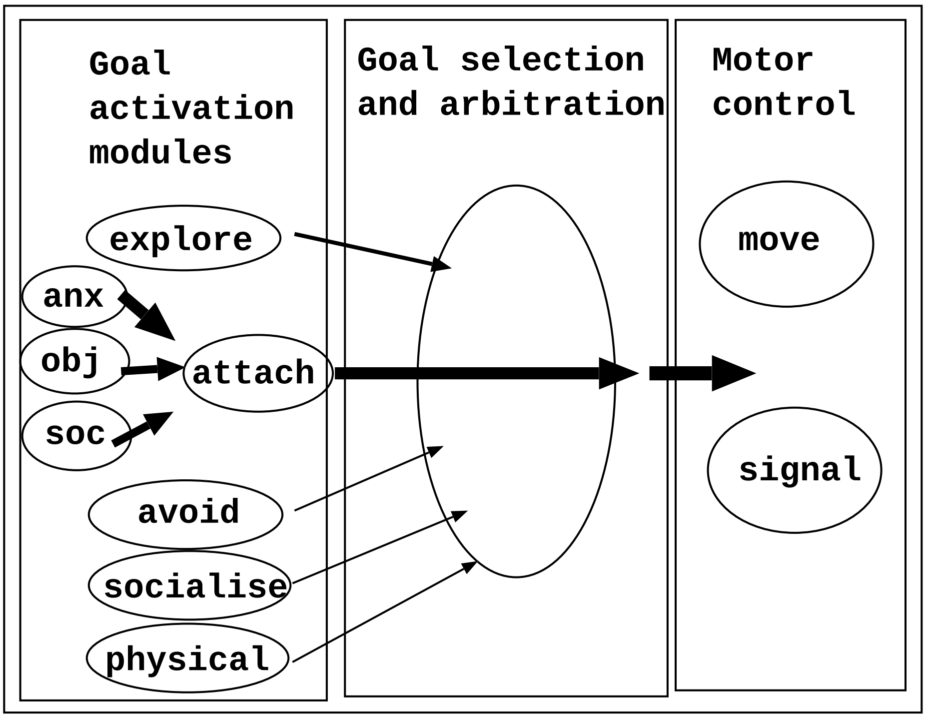

Figure 4.4: The **RAL** architecture when it has learnt to be Ambivalant. Throughout a reunion episode this architecture activates the goal of attachment (attach). The goal of contact avoidance (avoid) is not highly activated because previous experiences with the carer have not resulted in very high levels of unpleasant close contact.

because at close distances to their carers the goal of avoiding close contact inhibits active expression of **attach** goal behaviour, leading to behaviours linked to the exploration goal being activated as displacement behaviours.

The **RAL** architecture can be viewed as a psychological theory. However, the core parts of the theory need to be distinguished from peripheral theoretical elements and implementational aspects of the simulation that are not parts of the psychological theory. The principal core aspect of the **RAL** design that sets it apart from the other designs described in this chapter is that the two motives that mutually inhibit each other in the **RAL** architecture are of the same type. They are activated in the same way, are represented in the same way, and compete along side each other for activation. It is this competitive process that results in mutual inhibition of the two most active goals.

There are a number of limitations that the **RAL** architecture, and the carer agent with which it interacts, which are shared with the **HAR** architecture. These are described next in this chapter and are dealt with in detail in section 4.4. The key limitation which the **RAL** architecture does not share with the other architectural responses to the REUNION scenario is that the **RAL** architecture suggests that infants behave the way they do in the Strange Situation without any cognitive, thoughtful or reasoning process directing behaviour. There exist numerous accounts of behaviour for infants this age which do seem to show intentional action and even simple deliberative processes (Diamond, 2002; Willatts, 1999; Munakata *et al.*, 2002; Baillargeon *et al.*, 1990). The **HAR** architecture suggests that some kind of cognitive, thoughtful or reasoning process is involved with how some behaviours are produced in the Strange Situation.

4.3.3 The HAR architecture

This section describes a hybrid architecture called the **hybrid-action-reasoning** (**HAR**) architecture. In addition to a reactive level similar to the reactive architectures already described, this hybrid architecture possesses a planner implemented as a production system. The plans are very simple, and allow the infant to 'look-ahead' and reason about the immediate outcomes of its small set of available actions on its environment.

Figure 4.5 shows a more detailed view of the **HAR** architecture. The following figures present simplified views. There are two types of information flow represented in this figure, the arrows in the deliberative part of the diagram have a different meaning to those in the lower reactive part. In the deliberative processing the information transfer is on-demand, in comparison with the continually updated information transfer, which is more like electrical current, in the lower reactive processes. The deliberative processes also involve buffering of information, memory retrieval and inhibition.

Aside from being differentiated into reactive and deliberative parts, subsystems of the **HAR** architecture can also be distinguished according to whether perceptual, central or motor actions are occurring. As in previous architectures based upon the **GS** architecture, each of the six goal activator subsystems has three types of input: activation from external sense data, activation from internal sense data (simulating the effect of residual activation and hormones) and a target for the goal. In addition to the goal activator subsystems there is a perceptual system that inputs information about the comfort of close physical contact to the deliberative system.

The reactive central processing system carries out reactive goal selection. If reactive goals are not inhibited then they control the motor actions of the agent architecture. The infant agent can move and signal towards any target object or agent (or not move and/or not signal if that is appropriate). In the previous architectures described in this thesis, if reactive perception gives rise to reactive action through a reactive selection process then all representations may be implicit, in the sense that they do not survive after their immediate triggering. The current state of the infant is not explicitly represented in totally reactive architectures, it is only represented in the nature of the goals that have been activated. Something very different happens in the **HAR** architecture. When reactive goals are activated they are also represented explicitly in the deliberative component of the architecture. This means that the **HAR** architecture can represent the effects of performing actions towards those goals without having to actually perform those actions. This is because the current state of the infant agent and its active goal are represented in a form that endures and can be reused, perhaps for many purposes.

When the goal of attachment is activated by the reactive selection process because the carer is too far away and two data items are input to the deliberative subsystem. The current state of the relevant aspects of the infant agent's world would be represented by the item:

$$\textbf{[[carer isat far]]} \tag{4.1}$$

The current goal would be represented as:

$$\textbf{[goal [carer isat near]]} \tag{4.2}$$

When the carer agent returns within the infant agent's visual range the actions that can be carried out by the reactive motor control subsystem are to signal or move towards specified object or agent targets. At the deliberative level there are three action operators:

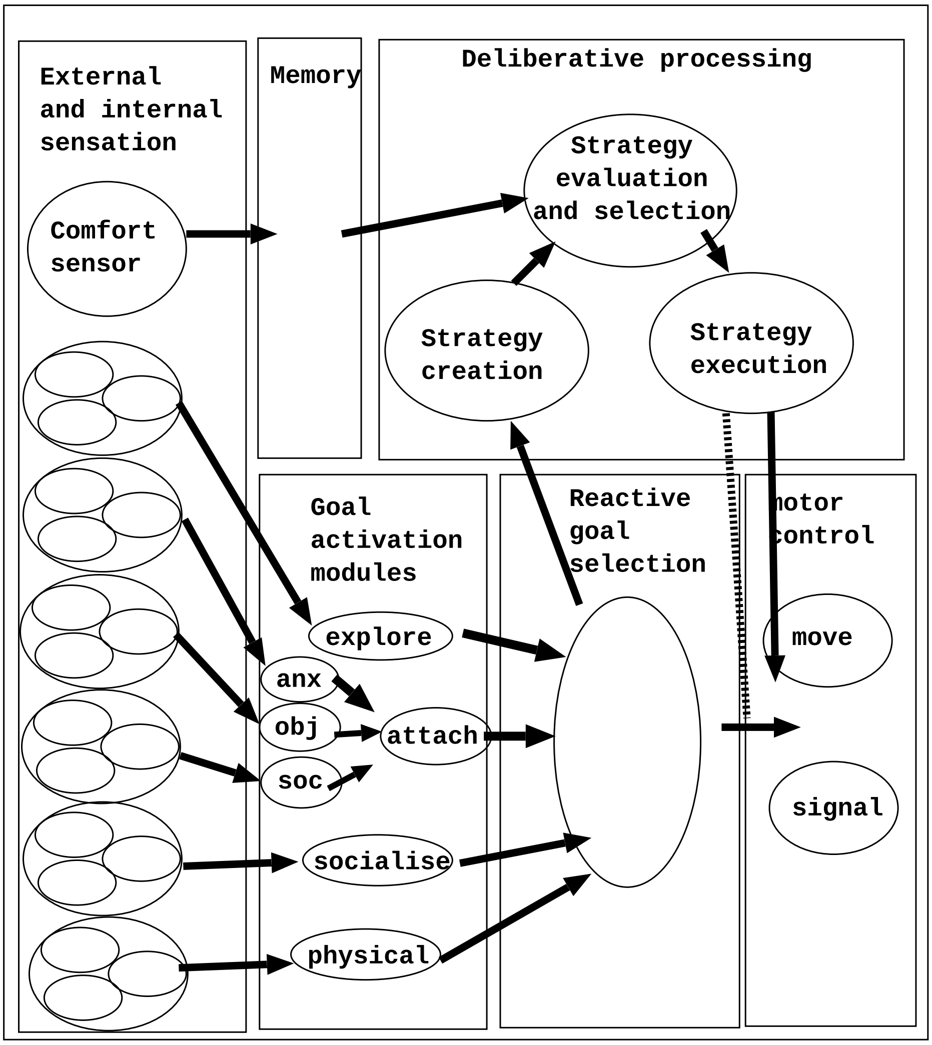

Figure 4.5: The Hybrid Design **HAR**. Deliberative mechanisms provide a secondary route to action, activated as a result of interrupts to reactive selection and arbitration and imposing different actions. The reactive goals are: exploration (explore); anxiety (anx); object wariness (obj); social wariness (soc); socialisation (socialise); and physical need (physical). The reactive goals of anxiety, object wariness, and social wariness are all combined in a single goal of attachment (attach) before they are considered for selection. The comfort sensor measures contact pleasure and represents this information symbolically.

[move-only] (4.3)

[signal-only] (4.4)

[move-and-signal] (4.5)

The item 4.3 represents the infant agent changing its coordinates in the 2D world so that it moves closer to its current target object, which is the carer agent. The item 4.4 represents the infant agent broadcasting a signal that can be sensed by the carer agent. The item 4.5 represents both these actions occurring simultaneously. What the deliberative system as a whole does is apply action operators to the perceptually derived item that describes the infant agent's current state and attempt to transform the item so that it matches the the goal state. The plans are represented as lists which can be 'read' from left to right. The initial part of the plan is the current perceived state of the world. Operators are fixed to the right of this statement about the world, and a new statement about the state of the world is then appended to the right most position of the plan. This process may be undergone a number of times before the left most state descriptor matches the agents goal state descriptor. A generalised three link plan format is:

[[current state of the world][operator to be tried][predicted world state]] (4.6)

Item 4.6 shows a three link plan format with a perceived state of the world being transformed into a state of the world that doesn't currently exist by means of a single operator. If the predicted state of the world matches with the goal state then the planning process is stopped, else the mechanism would add an another operator to the plan and describe a further imagined state of the world.

The operators are applied one at a time, and are ordered so that the move and signal operator (item 4.5) is always applied first. This means that at present the order is treated as an innately specified part of the architecture where moving and signalling is a prepotent deliberative response. However, the architecture might be developed so that it is open to learning by reinforcement so that different operators were considered first. This would only have an effect when the selection and evaluation subsystem has been developed to allow interrupts of deliberation. The results of the strategy creation process are plans such as:

[[carer isat far][move-and-signal][carer isat near]] (4.7) and

[[carer isat far][move-only][carer isat near]] (4.8)

If the infant agent currently possesses the goal described in data-item 4.2 (which represents the goal of attachment) then both the plans represented in data-items 4.7 and 4.8 would be created to satisfy this goal. Both plans shown in items 4.7 and 4.8 show the same current perceived state of the world, namely that the carer is too distant for the infant to assess it is secure. Both plans also possess the same imagined state of the world that matches the goal state, namely that the carer is close enough for the infant to assess it is secure. Where they differ is in the operators that are used to get from perceived state to imagined goal state. In the current implementation both these plans are created serially and then passed to the Strategy evaluation and selection subsystem, which only evaluates plans when no further plans can be created. This is a weakness of the architecture because there are likely to be many instances when after evaluating one plan it is better to immediately follow that plan rather than wait for another to be produced. The current arrangement where evaluation only occurs when all operators are exhausted should be developed to allow interrupts before all operators are exhausted.

Figure 4.6 presents the first ten time steps for a reunion episode, showing externally observable behaviour, reactive processing and deliberative processing in separate columns. This figure shows that at the very start of a reunion episode all infant agents are in a similar state, and the reactive goal of proximity is activated in all infants. The signalling action is activated without the movement action being activated. At the moment the carer agent returns within the infant agent's perceptual range, the reactive system for all attachment groups quickly activates the action of movement towards the carer agent. The same initial deliberative steps are also undergone by all infant agents. All infant agents build the same initial plans, which are represented in items 4.7 and 4.8. After these plans are evaluated the behaviour of the infants diverges. After a short period Secure infants will stop activating the goal of attachment and start exploring. Ambivalent infants will positively evaluate the plan of signalling and moving and the reactive and deliberative systems will activate the same actions of simultaneously signalling and moving. Avoidant infants will negatively evaluate the plan of signalling and moving. The plan of moving without signalling will initially be positively evaluated whilst the carer is distant from the infant. The deliberative system will therefore inhibit the reactive actions and substitute the action described by the accepted plan.

The Strategy evaluation and selection subsystem first evaluates the plans by

Time cycle	**Externally observable events**	**Internal reactive processing**	**Internal deliberative processing**
1027	Carer beyond infant's visual range, infant signalling but not moving	attachment goal active, no data on position of carer, signalling activated	previously adopted plan [signal without moving] directs behaviour
1028	Carer returns to within infant's visual range	attachment goal active, carer's position is set as target for movement and signalling actions	new reactive goal triggers new deliberative process
1029	Infant continues signalling and turns to face carer and moves towards carer	goal and actions maintained	data items for proximity goal, and current position of carer placed in working memory
1030	Infant continues signalling and moving	goal and actions maintained	new root for plan structure placed in working memory
1031	Infant continues signalling and moving	goal and actions maintained	signal and move operator and proximity goal matched
1032	Infant continues signalling and moving	goal and actions maintained	signal operator and proximity goal pushed onto plan
1033	Infant continues signalling and moving	goal and actions maintained	plan stored for evaluation
1034	Infant continues signalling and moving	goals and actions maintained	new plan started from root
1035	Infant continues signalling and moving	goal and actions maintained	move-only operator and proximity goal pushed onto plan
1036	Infant continues signalling and moving	goal and actions maintained	plan stored for evaluation
1037	Infant continues signalling and moving	goal and actions maintained	operators exhausted so plans evaluated
1038	EITHER Infant continues signalling and moving OR infant moves without signalling	EITHER new deliberative action matches reactive OR it inhibits reactive and activates new action	inhibitory signal only maintained while same deliberative goal remains active

Figure 4.6: Showing the external behaviour and internal processes moment by moment in a short period at the start of a reunion episode. Reunion occurs at time step 1028. Both Avoidant and Ambivalent infants will produce two plans in their deliberative subsystems before inhibiting reactive actions (this requirement is a limitation of the current architecture). Secure infants may become close enough to the carer during deliberation that the goal of proximity is deactivated and substituted with an exploratory goal.

matching them against any available data items that have been produced by deliberative level perception[7]. The comfort sensor in figure 4.5 is one such sensor. The data items produced by the comfort sensor are of the form:

$$\textbf{[[carer isat near][valence-measure]]} \tag{4.9}$$

where the valence measure is a cumulative measure of the positive or negative experiences of close physical contact with the carer. If one of the plans is selected then it is passed for execution, which involves activation of the first operator that it is listed in the plan and inhibition of the current active action from the reactive system.

4.3.3.1 Representing Avoidant patterns of behaviour

Figure 4.7 shows an instance in time during a reunion episode for an Avoidant infant possessing the **HAR** architecture. When the Strategy creation subsystem is activated it receives as input the goal of attachment (as in item 4.2). At the same time, the perceptual system inputs to the deliberative system the external state of the infant, which is represented by the data-item 4.1. The process of strategy creation involves the production of multiple strategies in a serial process. These plans match the goal state and so are considered by the Strategy evaluation and selection subsystem. When these strategies are passed for evaluation and the infant agent is still relatively distant from the carer then a plan that directs the infant to move to the carer without signalling is selected. When the carer agent is closer no plan is selected because all match the data-item from the comfort sensor which possess a strongly negative valence measure. Therefore all the actions in the plans are inhibited, leaving the highest reactive actions left to direct external behaviour. This may be a reactive action to explore, resulting in desultory exploration.

The **HAR** architecture has to explain the same data as the **RAL** architecture. The big challenge is to explain discontinuities in behaviour. Both architectures have to show why in one context infant agents produce the highest level of angry protest and in another context why they produce the lowest level of angry protest. The **HAR** architecture represents the hypothesis that Avoidant infants don't protest in the reunion episodes of the Strange Situation because they reason with simple mechanisms about the likely consequences of this action and decide to gain their goals either by moving towards the carer agent, or by 'doing nothing'

---

[7] Deliberative level perception does no need to be deliberative in the sense of enabling look-ahead, it simply needs to package its output in a form of representation that is used by the deliberative system

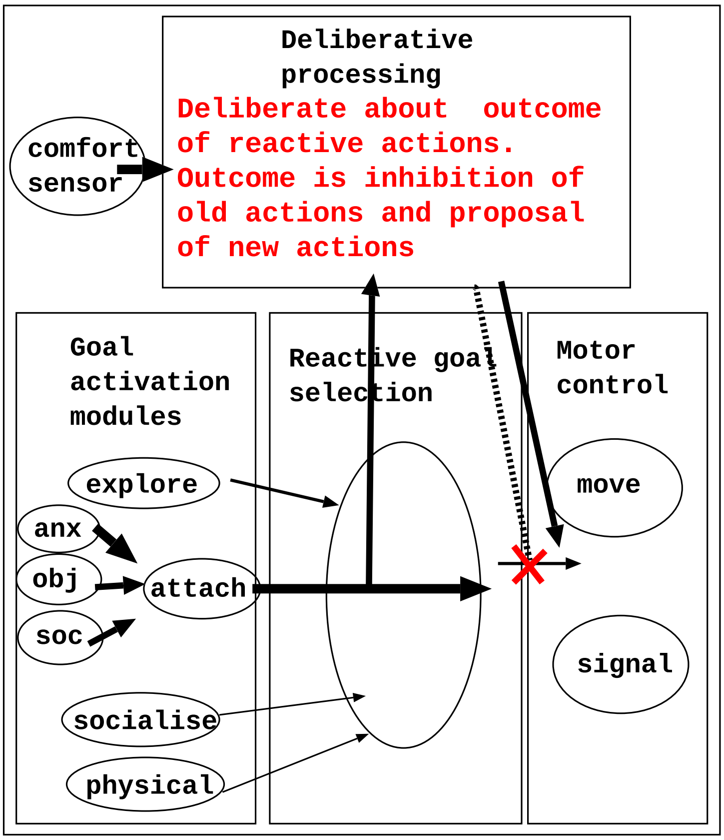

Figure 4.7: A snapshot of the **HAR** architecture in a reunion episode when it has learnt to be Avoidant. Note that there is not need to invoke a reactive goal of avoidance. The dashed line from the deliberative subsystem to the action subsystem signifies inhibition

and committing ‘default’ reactive behaviours.

The ultimate external cause of avoidance in the **RAL** and **HAR** architectures is essentially the same. This is that previous unpleasant experiences of close physical contact mean that in reunion the avoidant infant agent both wants proximity to gain safety and doesn’t want proximity so that it can avoid unpleasant contact. The mechanisms by which this conflict occurs are different. In the case of the **HAR** architecture a reactive prepotent response is inhibited after a deliberative process has ‘looked-ahead’ and found a better action. The reactive subsystem of Avoidant infants activates signalling but the deliberative subsystem inhibits this prepotent response and instead activates the action of moving without signalling or not moving or signalling. If the carer is at an intermediate distance then the infant agent will move without signalling. If the infant agent is very close then the infant agent will neither move nor signal to the carer, but will explore a nearby toy object instead.

Currently the **HAR** architecture only considers goals that have won-out in the reactive selection mechanism. This means that it cannot reactively explore without deliberating about exploring, which means that an infant agent cannot ‘pretend’ to explore as a means to keep safe. However, the architecture might be adapted so that the functions of reactive goal activation and deliberative motive activation might be split into two types of more independent process.

4.3.3.2 Representing Secure patterns of behaviour

Figure 4.8 shows an instance in time during a reunion episode for a Secure infant agent possessing the **HAR** architecture. Secure infant agents with the **HAR** architecture show low levels of protest as both their reactive and deliberative components become focused on exploration. Apart from the beginning of a reunion episode, a Secure infant agent does not possess a highly activated goal for attachment. In a room full of novel toys it is the exploration goal that is most commonly activated. The **HAR** architecture allows an infant to reactively and deliberatively explore.

In the cases of goals for exploration and socialisation, the goal includes a target that is a variable that is instantiated upon a specific toy or agent in the infant agent’s visual range. The goal of exploration is represented by the item:

$$\textbf{[goal [explore target-object]]} \tag{4.10}$$

and the goal of socialisation is represented by the item:

$$\textbf{[goal [social target-agent]]} \tag{4.11}$$

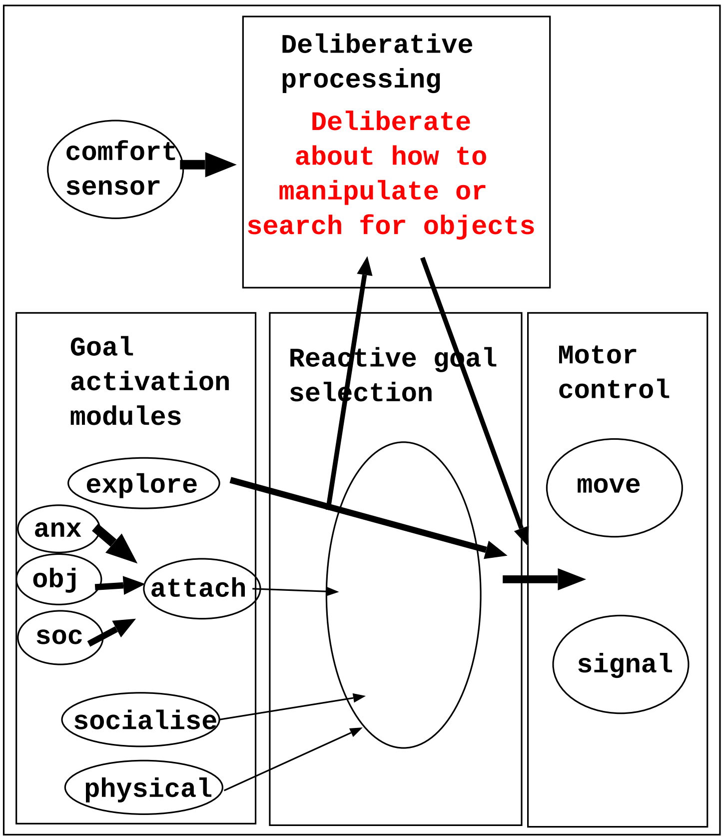

Figure 4.8: A snapshot of the **HAR** architecture in a reunion episode when it has learnt to be Secure. For much of the three minutes the infant behaviour may result from reactive and deliberative processes concerned primarily with the goal of exploration.

In items 4.10 and 4.11 the target objects or agents are set by the reactive goal activator module when the goal is activated for consideration by the deliberative subsystem. The representations of both socialisation and exploration in reactive and deliberative subsystems, and in the manner in which their actions are carried out, are shallow and should be considered as 'dummy' or 'mock-up' components.

4.3.3.3 Representing Ambivalent patterns of behaviour

Figure 4.9 shows an instance in time during a reunion episode for an Ambivalent infant possessing the **HAR** architecture. In the **HAR** architecture deliberative consideration of what to do when a reactive goal of attachment is activated is mandatory. That is, Ambivalent infants will always deliberatively consider using the 'signal and move to the carer' strategy, even is that is what they are already doing. This plan is always accepted but does not need to be executed as it is already being carried out.

4.3.3.4 Comparison with Crittenden's (1995) theory

Crittenden (1995) explains infant attachment behaviour with a theory that is similar to the account provided by the **HAR** architecture. Crittenden's account explains Avoidant infant behaviour in reunion episodes without recourse to the idea of displacement activities. However, it doesn't specify the types of information processing enough detail for direct implementation in a simulation. For Crittenden, the key difference between Avoidant and Ambivalent infants is their carers' differing responses to emotional signalling. If an infant's emotional signals result in rejection, then emotional signalling is in effect punished. Avoidant infants *"have learned to organise their behaviour without being able to interpret or use affective signals; that is they have made sense of cognition but not affect"*(page 371, Crittenden 1995). In contrast, the inconsistent responses that Ambivalent infants receive means they have been reinforced for affective behaviour, learning *"the temporal association of desire and its satisfaction with anger, uncertainty, and fear "*(page 371, Crittenden 1995). The **HAR** account is different to Crittenden's account because in the **HAR** architecture all infants use both reactive and deliberative processing. In the **HAR** architecture in the Avoidant case the deliberative subsystem inhibits the reactive, and in the case of the Ambivalent the reactive is undisturbed by deliberation. The **HAR** architecture could be adapted to conform to Crittenden's theory by adding a mechanism that stopped Ambivalent infants from deliberating, and left them to just rely on reactive 'prepotent' responses.

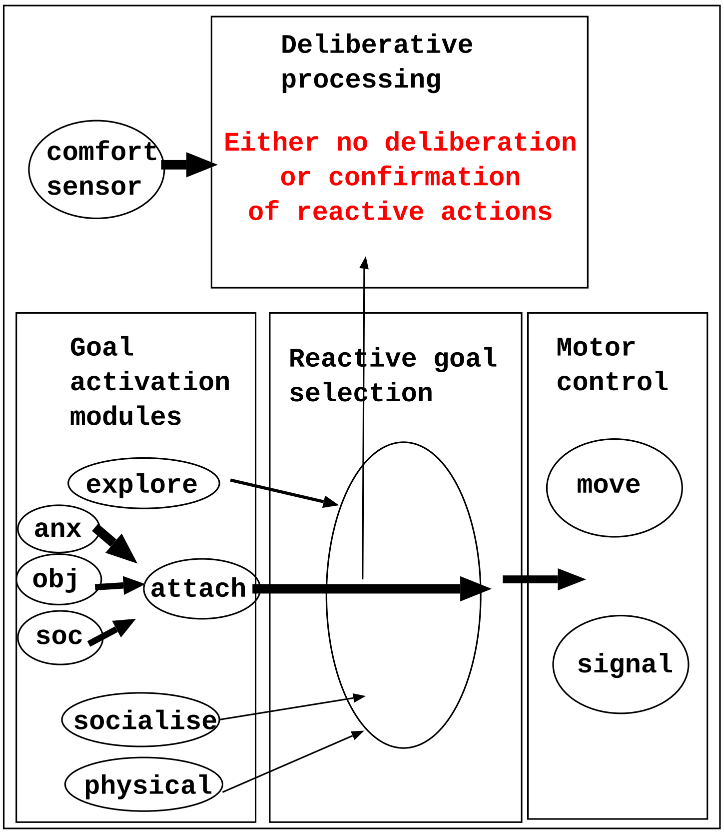

Figure 4.9: A snapshot of the **HAR** architecture in a reunion episode when it has learnt to be Ambivalent Infant behaviour may mostly result from reactive processes, with either minimal or no contribution from deliberative processes.

4.3.3.5 Comparison with evidence from studies of cognitive development in infancy

Figure 4.8 shows that two forms of object exploration are open to infants when they have activated the goal of exploration. There are a number of studies that corroborate this aspect of the account provided by the **HAR** architecture. Diamond (Diamond, 1985, 1991a,b; Diamond *et al.*, 1994; Diamond, 2000, 2002) has authored a number of studies that provide evidence that several infant tasks involve the interplay of two types of process. These processes are a fast reactive prepotent response and a signal from the maturing prefrontal cortex that inhibits the fast reactive prepotent response. These tasks include the A not-B tasks and object retrieval tasks. Diamond has found that:

> *"This effect of reinforcement on a response is evident in infants soon after birth and in the simplest organisms. It is early developing (in phylogeny and ontogeny) and robust, capable of surviving considerable neurological insult. A more fragile and later developing ability is the capacity to resist a predominant response, whether it is innately strong or has been strengthened by reinforcement... Although instinctual and habitual responses are very strong, even in humans, we are capable, with effort, of breaking a habit, whereas organisms without frontal cortex have no such option. The ability to resist the strongest response of the moment endows humans with extraordinary flexibility and freedom to choose and control our actions. It gives us the option of not being creatures of habit."* (page 363-364, Diamond, 1988)

The ability to inhibit reaching for objects becomes possible at about the same age that the goal corrected phase of attachment starts. Although inhibition of signalling and moving to the carer is a different kind of action in a very different context, it may be that attachment styles become distinguishable partly because a similar ability to inhibit protest actions matures within the infant:

> *"Between 8 and 12 months of age infants first become able to inhibit predominant response tendencies, that is, they first become able to resist the strongest response of the moment (a response tendency can be inherently predominant, such as reaching straight for a visible goal: If you see what you want the tendency to go toward it does not have to be learned. Indeed, it requires effort and discipline to resist this tendency when a more circuitous route is appropriate. A predominant response can also be acquired or learned, e.g., on the*

> *basis of reinforcement experience.) Inhibition of the dominant response response depends upon maturation of dorsolateral prefrontal cortex."* (page 68, Diamond, 1991b)

Willatts (1999) and Munakata *et al.* (2002) describe related experimental setups that investigate the kind of processes that occur when four to nine month old infants use a towel to pull an out of reach toy to within their grasp. Willatts discerned when a truly intentional response occurred by careful observation of eye movements and Munakata distinguished reactive from deliberative processes by seeing how responses varied when learning was available and when it had not occured because the experimental apparatus was new to the infants. Both Willatts and Munakata's results suggest that infants develop the ability to intentionally reason about how to grasp out of reach toys by using a towel to pull them closer by nine months of age. Comparisons between studies of problem solving of the type described by Willatts (1999) and Munakata *et al.* (2002) and behaviour in the Strange Situation is limited by the fact that different types of information process may occur in 'hot' versus 'cold' emotional contexts (Kerr and Zelazo, 2004).

4.3.3.6 Limitations of the HAR architecture

There are a number of limitations of the **HAR** architecture which are shared with the previously described **RAL** architecture. These limitations are dealt with in detail in section 4.4. There are also limitations specific to the **HAR** architecture, which are described in this section.

4.3.3.6.1 It is incomplete The **HAR** architecture possesses shortcomings that could be rectified with further time spent on its development. These shortcomings reflect the constraint of preparing the architecture for a deadline for PhD submission rather than a limitation in the underlying theory. Incomplete aspects of the **HAR** architecture include that there is no caching of deliberative plans that have been successful. Nor is there interruption of deliberation. Future work might allow for deliberative strategies to be evaluated as they are serially constructed. If a plan were positively evaluated then deliberation would stop and the plan would be passed for execution. In the case of Avoidant infant agents, deliberation would continue, whereas no deliberation would occur in Ambivalent infants.

4.3.3.6.2 It is too advanced A key potential limitation which the **HAR** architecture does not share with the **RAL** architecture is that it may rely upon structures and mechanisms that are too advanced for the target stage in development. The **HAR** architecture introduces a whole range of advanced structures and processes in a single agglomerate package that over estimates the capabilities of nine month old infants (Diamond 2002, Willatts 1999). In principle, the **HAR** architecture has sufficient mechanisms to compute many complex plans. It is only limited by its impoverished data-base and set of operators.

Spelke (2000) has theorised that infants possess a 'core knowledge' system whose building blocks emerge early in human ontogeny and in phylogeny. Spelke hypothesises that these systems do involve forms of representation, but the representations are specific to a class of entities for a particular set of purposes. Ultimately, these domain-specific representations may be combined, but this may not occur until the development of natural language. Until the development of natural language an infant's available forms of representation may be more similar to other animal species than adult humans (Hauser and Spelke, 2004). If this is the case then forms of representation in the **HAR** architecture may be too advanced because they allow domain general reasoning.

We are only concerned with representation within infants in the particular circumstance of them finding themselves in the reunion episode of the Strange Situation. We are therefore not interested in problem solving that requires large numbers of intermediate states. To reproduce the Avoidant pattern of behaviour seen in reunion episodes of the Strange Situation only a subset of these deliberative capabilities is needed. To carry out the simplest forms of deliberative 'look-ahead' to see the result of their actions infants need to be able to represent the effect of performing actions instead of merely performing actions. A minimal set of representations and operators would include the infant being able to represent the initial state of deliberation, that is the world as they now perceive it. They also need to represent the goal state, that is the state or features of the world that they would like to bring about. Lastly they need to possess simulative action operators to move their initial representation to see if it reaches the goal state (Baillargeon *et al.*, 1990; Willatts, 1997). With only these operators perhaps it would be more accurate to term the reasoning that results 'proto-deliberation' (Sloman and Chrisley, 2005; Arbib, 2002).

4.3.3.6.3 It proposes multiple completely-independent modules that carry out selection. There are two places in the **HAR** architecture that selection processes occur. Firstly there is reactive selection, goals of different activation levels are selected for immediate execution by the motor system. Secondly, the deliberative component of the architecture also carries out selection tasks as it chooses which operators to add to plans and then selects between completed plans. This division of selection into completely independent reactive types and deliberative types of selection is a core assumption of the **HAR** architecture. However, there is evidence that there is not such a clear cut distinction in the brain. It is suggested that the Basal Ganglia is an organ that carries out selections from lots of different largely independent parallel processing functional units (Redgrave *et al.*, 1999; Prescott *et al.*, 1999, 2002). Some of the functional units competing for selection direct movements, and disorders of the Basal Ganglia are implicated in Parkinson Disease and Huntington Disease (Kandel *et al.*, 2000). As was noted earlier, the Basal Ganglia is also implicated in the operation of Contention Scheduling (Rumiati *et al.*, 2001). It has also been proposed that the Basal Ganglia acts to select cognitive actions (Kandel *et al.*, 2000), and in this role has been modelled as a production system in the ACT-R architecture (Anderson and Lebiere, 1999). Different function roles are undertaken by different parts of the Basal Ganglia. Functionally segregated circuits for skeleto-motor, oculo-motor, associative, and limbic processing are linked in parallel through the same layers of the Basal Ganglia. Each of these separate functions is linked from the cortex through the striatum, pallidum/nigra and thalamus back to the cortex (Redgrave *et al.*, 1999). Though each functional circuit is organised to be segregated and in parallel, there may exist important constraints or dependencies between different functional sub-components.The limitation of possessing multiple centres of completely independent process selection is discussed further in section 5.3.2 in the next chapter.

4.3.4 Preliminary quantitative results for the RAL and HAR architectures in three boundary caregiving contexts

A preliminary set of quantitative experiments with the **RAL** and **HAR** architectures have been carried out by placing initially identical infants in three boundary environments that reflect extremes of caregiving:

- low sensitivity and high discomfort (type A)
- high sensitivity and low discomfort (type B)
- low sensitivity and low discomfort (type C)

Figure 4.10 shows results. Ten experiments were run for each attachment type with each of the two infant architectures giving sixty experiments in total. There were two independent variables that were manipulated in these experiments, carer response threshold and carer discomfort level. The response threshold was the level of infant signalling that would trigger a carer response. So a level of 50 represents an insensitive carer that ignores the infant agent's signalling for a long period. In contrast a carer response threshold of 5 represents a very sensitive carer. Regarding the carer discomfort levels, a level of 0.9 represents a maximal level of discomfort and a level of 0.1 represents a minimal level of discomfort.

Attachment type	**Carer response threshold (IV)**	**Carer discomfort level (IV)**	**Safe-range (DV)**	**infant measure of discomfort (DV)**
RAL architecture				
A	50	.0.9	238.87 (SD 13.92)	0.583 (SD 0.031)
B	5	0.1	479.8 (SD 1.50)	0.0665 (SD 0.0017)
C	50	0.1	254.4 (SD 6.12)	0.0657 (SD 0.0021)
HAR architecture				
A	50	.0.9	267 (SD 7.44)	0.598 (SD 0.012)
B	5	0.1	479.9 (SD 1.49)	0.0649 (SD 0.0016)
C	50	0.1	243 (SD 11.01)	0.0588 (SD 0.0019)

Figure 4.10: Quantitative results for the **HAR** architecture.

There were two output or dependent variables measured. The infant agent's safe-range value and its internal measure of discomfort that it has learnt from its experience. The safe-range values had an upper bound of 480, and all infant agents in both **RAL** and **HAR** architecture conditions approached this limit in the B type caregiving condition. The lower bound on the safe-range was 160, and no instances were found of infant agent's meeting this bound. The starting safe range value was set at 320 and so infant agents in all A and C type conditions for both architectures show a decrease in safe-range and all B type infants in both conditions show an increase in safe-range. The algorithm by which infant agents learn a weighting for an association with caregiving discomfort in close contact is discussed in more detail in section 4.4.0.1. Figure 4.10 shows that all infants in B and C type conditions for both architectures show a low association weight for discomfort, and all A type infants for both architectures show a high association for discomfort.

How should these results be evaluated? They are only preliminary and merely act as an initial proof of concept for the claim that both architectures can capture the same patterns of behaviour. For a more substantive evaluation to be made then a more systematic and comprehensive set of experiments would need to be carried out that simultaneously varied carer sensitivity and discomfort from extreme to intermediate values.

4.4 Evaluation of assumptions common to the RAL and HAR architectures

4.4.0.1 Shared aspects of close contact learning in the RAL and HAR architectures

Learning about the nature of close contact with their carers occurs in different components of the **RAL** and **HAR** architectures. However, aside from this fact, the learning mechanisms that have been implemented in these architectures are alike. In both mechanisms the infant agent perceives a quality of close contact discomfort with the carer, and uses this value to learn an association. Values of discomfort that the infant agent can perceive have been set to be within the range 0 to 1. A value of 0 represents fully comfortable close contact and the value of 1 represents a maximum bound for discomfort. The association value is represented as a single number and is learnt with a variant of the neural network supervised learning delta rule (Dawson, 2004). The change in the association (δw) is given by:

$$\delta w = \alpha(t - w) \tag{4.12}$$

The α term represents a learning parameter, currently set at 0.1. The t term represents the target discomfort value that the infant is trying to learn and the w term represents the current value (a weighting) of infant agent's association with discomfort in close contact. The learning rule in 4.12 behaves similarly to how animals learn by reinforcement. This is because the patterns of learning that is created by the rule described in 4.12 parallels the pattern of learning seen in the Rescorla-Wagner reinforcement learning model (Pearce, 1997; Dawson, 2004). In the case of equation 4.12, the perceived current discomfort acts as a supervision signal.

Limitations of this approach include that the the carer agent response is fixed. Therefore, the infant agent quickly learns its carer agents discomfort value. In reality even carers that provide great discomfort on some occasions will provide holding with little or no discomfort on other occasions. Therefore in reality the infant would be attempting to learn a 'moving target' that would never settle and could only be approximated.

In addition, this mechanism is limited because it over simplifies discomfort by representing it as a single number. A richer representation would include the context of any holding, so if the infant and carer agents were engaged in boisterous interaction what would constitute discomfort would be very different from discomfort in quiet moments. Also, perception of discomfort is likely to show individual differences between infants, and this kind of infant difference has been absent throughout this thesis, but may be introduced in future work.

4.4.1 A fixed carer predisposition for close contact

A number of sources of evidence point to the conclusion that carer effects are larger than infant effects in deciding what attachment style an infant will come to possess at one year of age (van Ijzendoorn and Bakermans-Kranenburg, 2004). What this means is that, for most infants, how a carer behaves has a stronger influence on the ultimate nature of the infant-carer relationship than the sum of initial infant predispositions. However, infant effects cannot be discounted. One infant effect that neither the **RAL** nor **HAR** architectures currently allow is that an infant agent that is initially averse to tender close contact may trigger an equal aversion to such contact in a carer agent who might otherwise not have experienced an aversion so strongly. Another, perhaps more interesting extension to the **RAL** and **HAR** architectures might be that infant agents only develop a minor aversion after experiencing moderate initial carer ‘stiffness’, but that the negative sensations of each agent cause positive feedback. This kind of feedback would be different to that demonstrated in the last chapter because both carer agent and infant agent might start with intermediate levels of aversion and then both would change as a result of each sensing the others emotional retreat.

Simulative experiments have been carried out where the available energy in the virtual environment was varied. If energy levels are abundant then the carers of Avoidant infants behave more like the other carers. However if energy is scarce the reverse happens, and all carers seem like carers of Avoidant infants. This accords with an empirical study that found low levels of support to caregivers predicted a higher likelihood of infants receiving an Insecure classification (Crockenberg, 1981).

4.4.2 The role of consistency and the ‘Sports-hall Situation Experiment’

Section 4.2.4 noted two different explanations for how Avoidant and Ambivalent styles of attachment come to be differentiated. The **RAL** or **HAR** have been designed to differentiate Insecure styles of attachment according to the nature of the carers close physical contact. An alternative hypothesis, based upon the idea that consistency of caregiving is what brings about infant differences, has not been implemented. This is partly because Ainsworth *et al.* (1978) does not publish statistical measures of dispersion for the measures of carer behaviour observed at home, so a measure of consistency was not readily available. However, a number of contributors have noted that consistency in responsiveness is a key discriminant for the carers of Avoidant and Ambivalent infants (Cassidy and Berlin, 1994; Crittenden, 1995). Currently neither the **RAL** or **HAR** architecture possesses mechanisms that would allow them to differentiate Avoidant from Ambivalent styles of caregiving based upon consistency data. This is because the mechanisms that these architectures do possess merely keep a running total of the quality of contacts that have been previously experienced.

Work in progress has involved developing the carer agents so that their architectural design can support the notion of panic. This is implemented by a parameter that sets the level at which the carers' behaviour switches to a different and less efficient mode. The carers of Secure and Avoidant infants have a very low probability of 'panic' being triggered, so that the other currently specified parameters provide a good summary of their behaviour. The carers of Ambivalent carers often 'panic', and this means that their behaviour is patchy, unpredictable and overall much less responsive than that provided by the carers of Secure infants. The notion of 'panic' is not a good representation of the ordinary use of this term, but does reflect the qualities of the carers of Ambivalent infants found in studies that have used the Adult Attachment Interview (AAI) to assess the attachment status of infant caregivers (Hesse, 1999). The AAI analyses the discourse properties of adult carers talking about attachment relations and describes the carers of Secure infants as autonomous and 'free to evaluate' and the carers of Avoidant infants as dismissing. The carers of Avoidant infants are described as preoccupied, enmeshed, and often engaged with angry struggles with their own carers. It is a pattern of caring which results from the preoccupied states of mind that the 'panic' behaviour is trying to capture.

Future work may involve developing the perceptual apparatus of the **RAL** and **HAR** architectures so that information about consistency can be used to differentiate caregiving style. An important part of evaluating a design based study of infant phenomena is whether new predictions can be made and then empirically tested. The contrasting close contact and consistency hypothesis do make different predictions about the behaviour of infants in particular circumstances. For example, if the Strange Situation was adapted so that it was undertaken at one end of a much larger room, perhaps in a sports-hall, then the distance of the carer when the carer first appears at reunion could be varied. If the whole procedure was carried out at one end of the sports-hall then it may result in similar outcomes to the normal Strange Situation. If the carer and infant are placed at one end of the sports-hall but the carer is returned to the other end then infants would be reunited with a carer that was still a long distance away. In these circumstances the **RAL** and **HAR** architectures as they now stand would predict Avoidant infants to behave less avoidantly until the carer got closer. This is not a straightforward suggestion. As Ainsworth *et al.* (1978) note, in the development of the Strange Situation Experiment it was found that infant behaviours were very sensitive to the relative positions of toys, furniture and the door. Since the consistency hypothesis has not been implemented we cannot be sure how this theory would act, but it may be that it would predict a different pattern of behaviour. Therefore the 'Sports-hall Situation Experiment' can be put forward as an extension of the Strange Situation Experiment to potentially test hypotheses that have arisen out of this work.

4.4.3 Lower level description of physical contact needed

Both the **RAL** and **HAR** architectures are described at a high level of description. Two areas of these simulations that suffer from this decision are: the representation of comfort in contact; and the addition of temporal constraints on different types of processing.

The Avoid-pain subsystem in the **RAL** and the comfort sensor in the **HAR** architecture both 'solve' the problem of Avoidant reunion behaviour but both create a new problem of how to represent physical contact in this high-level simulation. The literature does possess descriptions of a perceptual system that incorporate the effects of close contact, such as that reviewed by Polan and Hofer (1999). Rat studies have uncovered mechanisms related to physical contact and attachment, that are termed 'hidden regulators', and which are believed to be the physiological basis of the state of 'felt-security'. These mechanisms are believed to produce low level expectations regarding comfort and safety in rats and human infants. Their presence is even thought to modulate physiological aspects of social phenomena such as separations and losses in adulthood. The links between specific aspects of physical contact and expectations of 'felt-security' are mediated by the release of substances like opiates, and the addition of these mechanisms may form part of the future development of these simulations.

4.4.4 Introduce temporal constraints on processing

The **RAL** and **HAR** architectures have not been matched to neural or other physiological structures or mechanisms. This means that temporally constraining the processes in each architecture is not possible. Lists and arbitrary POP11 procedures and non-atomic operations are used in their implementation. This means that the timing of the internal processing in the implementation is not limited by known neurophysiological constraints. Timing is important because we can expect temporal differences in the performance of deliberative and reactive processes. However, just because we might expect differences doesn't mean that they will occur, and any differences may be too slight to measure, particularly in the context of the Strange Situation.

4.4.5 Representing anger

Currently this work represents infant anger simply by more angry infants signalling with a more intensely negative tone. As was noted in the last chapter, infants produce different types of crying, which include a distinguishably angry cry (Durkin, 1995). We are not concerned with the specifics of the appearance of emotions but instead with the underlying function. However, the appearance of emotions are believed to perform functions. There exist several types of states which have been put forward as candidate descriptions of anger states that are described in appendix F.1 and might form part of the future development of this work.

4.4.6 Explaining attachment in term of four types of question

This work has focused on using agent architectures to explain infant attachment behaviour. Are there aspects of infant attachment behaviour that this work has failed to explain because of this focus? We can attempt to explain biological behaviour in terms of four questions (Tinbergen, 1963; Hinde, 1983):

- What are the immediate causes of the infant's behaviour, in terms of environmental triggers and architectural mechanisms and structures that act upon these triggers?
- What are the functions of the infant's behaviour?
- How were the long term predispositions for the infants behaviour formed during ontogenetic development?
- How did architectures that support the range of behaviours seen in human infants come to evolve during phylogenetic development?

These four questions are closely linked, with many interactions between them. For example, immediate causes of behaviour are distinguished from ontogenetic causes by the time-scales that these causes act over.

The **RAL** and **HAR** architectures have been chiefly concerned with providing answers in terms of immediate and longer term architectural causes of behaviour. An ontogenetic, evolutionary and dynamic systems conception of attachment style is to view it as a way of setting life-long predispositions for the form of social and emotional interaction that the individual will undertake. Perhaps evolution has discovered the benefits of a mechanism for controlling the behaviour of infants that becomes 'set' to predispose the infant for certain kinds of social interaction. The **GS** architecture allowed all types of behaviours to be carried out at some point. Any behaviour that had not been carried out for a long time would gain in internally provided activation until it was carried out. The **GLA** and **GLS** architectures show that the predisposition for different types of behaviours may change to maintain conditions, to act conservatively. The **RAL** and **HAR** can act more radically, if certain types of behaviour do not bring reward to the infant, such as emotional intimacy, then they are dropped. By the end of infancy an infant's architecture may have 'decided' what style of social interaction to follow and this style may be perpetuated throughout the individual's life (Chisholm, 1999; Belsky, 1999; Simpson, 1999).

However, these architectural solutions have not attempted to deal with ontogenetic questions related to representational change and emergence of new forms of representation. In principle these architectures should be amenable to supporting this kinds of explanation. For example, future work might involve deliberative rule compilation or goal activator compilation.

The period of infancy that the **RAL** and **HAR** architectures are attempting to represent is also the period when infants start to show evidence of understanding the concept of agency, that others have intentions and can provide knowledge. Russell (1996) suggests that the development of concepts may derive from human infants being able to experience a sense of their own agency. A necessary requirement for experiencing this type of agency is executive functions that allow an infant to control its own thought processes. Once an infant can do this it can represent its 'self' as separate from their environment. It may be this ability that allows an infant to become aware of itself as a causal agent. Once an infant conceives as itself as a causal agent it is a perhaps a short step to then conceive of others as causal agents too, giving rise to the development of joint attention (Carpenter *et al.*, 1998; Tomasello, 1999b,a). Autonomous agents of the type implemented in this thesis may be able to provide a proof of concept for theories that posit that psychological phenomena arise from executive functions. This is because autonomous agents allow a researcher to construct information processing architectures of the complexity needed to support the kinds of executive functions that Russell (1996) describes.

Understanding agency is a prerequisite for cultural learning. In appendix F.2 Minsky describes a type of cultural learning which he characterises as attachment learning which the current **RAL** and **RAL** architecture cannot support, but which future development may allow.

The questions of the possible function of infant attachment behaviour and the phylogenetic developmental causes of infant attachment behaviour have, to a great extent, been treated together in this thesis. It is possible that some form of validation for the evolutionary assumptions of function made in this work may be derived from multi-generational evolutionary simulations. A start has been made on this. An evolutionary algorithm has been partially implemented and is described in appendix F.3.

4.4.7 No support for longer term phenomena

There are two longer term phenomena that are important and which none of the architectures in this thesis support. These are:

- Repetition of the Strange Situation. If infants go back to the laboratory where they took part in a previous Strange Situation experiment they will often show greatly heightened distress.
- Stages gone through following long term separations. Bowlby (1969) described three stages that all infants go through following long term separation of an attachment figure. These are distress, anger and then detachment.

None of the architectures in this thesis possess the ability to support these particular phenomena. These are deficiencies in the core assumptions of both architectures.

4.5 Conclusion

This chapter has described the implementation of the **RAL** and **HAR** architectures. These architectures were both designed to satisfy the REUNION scenario, which is based on an abstraction of the first exploratory Strange Situation study. The main research question in this chapter is: what are the caregiving experiences, and what are the internal architectural mechanisms, that give rise to the two different Insecure types of attachment? This research question has been refined by consideration of physiological studies which rule out the hypothesis that Avoidant infants undergo minimal stress in the Strange Situation. Cross-cultural studies are also assessed to rule out the hypothesis that the types of behaviour seen in the Strange Situation are cultural artifacts. The function of the different patterns of behaviour are analysed from an evolutionary perspective. This work follows the hypothesis that all infant patterns of behaviour in the Strange Situation are involved in the infant gaining security.

Why have two architectures been chosen to act as solutions to the REUNION scenario? The **RAL** and **HAR** architectures represent two sides of an active debate in infant psychology between the virtues of more or less advanced mechanisms for explaining infant behavioural phenomena. Although these architectures were designed with the intention of creating contrasting solutions to the problems represented by the REUNION scenario, they actually have many assumptions in common. For example, both the **RAL** and **HAR** architectures are described at a similar level of abstraction and both architectures differentiate Insecure attachment styles according to the quality of close bodily contact.

The process of implementing the **RAL** and **HAR** architectures has highlighted design issues that may feed into empirical studies of real infants. The contrast between causal models based upon close contact or consistency suggests an adaptation of the design Strange Situation to a research design that has been labelled the 'Sports-hall Situation Experiment'. In this type of experiment reunion episodes would not occur in the confines of a small office, so there would be no immediate prospect of close bodily contact. In addition, several different models of anger have been suggested, and distinguishing which (if any) of these occur in real infants might also be empirically studied.

Chapter 5

Conclusion

> *"Once we are faced with a problem, we may begin to work on it. We may do so by attempts of two kinds: we may proceed by first attempting to guess or to conjecture a solution to our problem; and we may then attempt to criticise our usually somewhat feeble guess. Sometimes a guess or conjecture may withstand our criticism and our experimental tests for some time. But as a rule, we soon find that our conjectures can be refuted, or that they do no solve our problem, or that they solve it only in part; and we find that even our best solutions - those able to resist the most severe criticism of the most brilliant minds - soon give rise to new difficulties, to new problems. Thus we may say that the growth of knowledge proceeds from old problems to new problems, by means of conjectures and refutations."* (page 258, Popper, 1972)
>
> *"All organisms are constantly, day and night, engaged in problem solving; and so are all those evolutionary sequences of organisms - the phyla which begin with the most primitive forms and of which the now living organisms are the latest members"* (page 242, Popper, 1972)

5.1 Introduction

The research in this thesis captures both the senses of problem solving expressed in the quotes above. As an example of scientific research this work is promoting the construction of theories that are guesses as to how things work in reality. These guesses are sufficiently detailed and precise to be the basis of working simulations. A claim of this thesis is that producing simulations of complete agents is a valuable way of doing research into developmental psychology. The simulations produced in this thesis are conjectures, to be evaluated, refuted and hopefully improved. The subjects of study - infants interacting with their carers and exploring their environments, are also problem solving in the sense Popper uses the term. The particular problems of infancy that this work is simulating are related to how infants deal with trade-offs such as how to balance security and exploration or to how an infant might gain security from a carer that it would otherwise want to avoid.

The aim of this project has been to build infant and carer software agents that reproduce the differing patterns of behaviour found in studies of attachment behaviour in our target time period. These simulated agents have been designed in a manner that increases our understanding of how and why the patterns are formed in reality. Another way of saying this is that in this project there has been a requirement to not only reproduce observable patterns of behaviour at a superficial level, but also to better understand the causal processes that give rise to the observable behaviour. The simulations have been designed to be as simple as possible whilst being powerful enough to represent possible mechanisms and causal structures that underlie the patterns of behaviour. Since no previous theories describing attachment behaviour have been implemented in complete agents, the theories described in this thesis are only putative 'first approximations' that will require further development.

How might we frame the architectures that have been produced in this thesis so they can be easily compared with other psychological models? The first part of the chapter characterises the architectures in this thesis against seven criteria against which a model can be measured (Webb, 2001). Webb notes that *"No specific position in the space of models thus defined as the* **only** *correct one, but a good methodology should be explicit about its position and the justification of that position"* (page 1033, Webb, 2001). Webb's seven dimensions are:

Section 6.2.1 - Relevance: whether the model tests and generates hypotheses applicable to biology.

Section 6.2.2 - Level: the elemental units of the model in the hierarchy from atoms to societies.

Section 6.2.3 - Generality: the range of biological systems the model can represent.

Section 6.2.4 - Abstraction: the complexity, relative to the target, or amount of detail included in the model.

Section 6.2.5 - Structural accuracy: how well the model represents the actual mechanisms underlying behaviour.

Section 6.2.6 - Performance match: to what extent the model behaviour matches the target behaviour.

Section 6.2.7 - Medium: the physical basis upon which the model is implemented.

The second part of the chapter is a description of future work that might be undertaken. This work is expressed in the form of architectural mechanisms that might be added to the previously described architectures, and that might remedy some of the limitations of the architectures described in this thesis. The final section of this chapter concludes this thesis with remarks concerning the benefits of doing research on infant attachment by designing autonomous agents.

5.2 Placing the contributions made in terms of Webb's seven dimensions

5.2.1 Relevance

The target system for this thesis is clearly biological. This does not mean that the only application of the simulations described in this thesis is in narrowly directed scientific research. The simulations might also be adapted to act as educational tools, allowing students of psychology to interact and modify different theories of psychological phenomena. Through 'playing' with the simulations these students may gain insight into requirements for psychological theorising.

In addition, results from chapter four show that the **GLA** architecture can interact with the carer within the virtual environment as a dynamic system. Positive feedback loops magnify small initial differences so that ultimately the population of infants appears as if it possesses discontinuous categories. This aspect of the model might be extended, for example by seeing if the simulations could be adapted to produce other novel examples of dynamic systems behaviour. However, Webb cites several researchers who warn against investigating models for their own sake. Hoos (1981) terms it *'modelitis'* when researchers are *"more interested in the model than the real world and studying only the portions of questions that are amenable to quantitative treatment"* (page 42, Hoos, 1981). A similar point is made by Bullock (1997) who suggests that when *"simulations are presented as 'artificial worlds' worthy of investigation for their own sake ... this practice is theoretically bankrupt, and such statements have no scientific currency"* (page 457, Bullock, 1997).

5.2.2 Level

This work has been directed at a high goal oriented level. The basic units of behaviour in the simulation are not the actions of hands or eyes, but are instead more abstract actions such as movement or signalling towards particular objects or agents[1]. This state of affairs has been assessed as a limitation in several previous sections. However, it was not easily remediable because of the nature of the source material that the scenarios were based upon. Ainsworth *et al.* (1978) used some types of frequency and percentage measures of behaviour that would be more readily adapted to lower level scenarios with more concrete basic units. In addition Ainsworth *et al.* (1978) also made coded measure of behaviour that abstracted away from concrete behaviours that might be simply coded in frequency or percentage measures. Ainsworth *et al.* (1978) notes that:

> *"One of the major difficulties in relying upon frequency assessments of behaviours implicated in interpersonal interaction is that ... behaviours that are superficially quite different may have a certain degree of interchangeability. If this is the case, then assessments based on specific behaviours separately considered may well fail to reflect the true state of affairs.*
>
> *Furthermore, in any naturalistic situation and even in our laboratory situation, in which the behaviour of the adults (mother and stranger) was only partially controlled through instructions, the behaviour shown by the infant toward another person can scarcely be assessed without considering the context provided by that person's behaviour. Such considerations wreak havoc with comparisons across individuals or across situations, unless one can find some way of taking into account the contingencies of interchangeability of behaviour and of reciprocal behaviour (or lack of it) in the partners in the interpersonal transactions in question."*

What this means is that the abstract level of description in the simulations have ultimately been inspired by Bowlby's Behaviour System explanation of attachment phenomena (Bowlby, 1969), because it is this framework that informed the coding procedures used by Ainsworth *et al.* (1978). Recent research has uncovered many lower level neural and other physiological attachment mechanisms (Fox and Card, 1999). Incorporation of this lower level material is a possible future direction of research for this work.

---

[1] Appendix B.2 compares the high goal oriented level demonstrated in this thesis with various other levels of description

5.2.3 Generality

How many real world situations does this work apply to? The research in this thesis is quite closely directed at attachment behaviours in the nine to twelve month period of infancy. Natural extensions of this work might include investigation of architectures to support a variety of different kinds of attachment phenomena, and a variety of different kinds of infant behavioural phenomena. Other attachment phenomena and attachment theories that this work might be extended to simulate include work on:

- Later infancy - (Main *et al.*, 1985; Main, 1991)
- Adolescence - (Main *et al.*, 1985; Main, 1991; Kobak *et al.*, 1993)
- Adult Attachment - (Hesse, 1999)
- Cross-cultural studies - (van Ijzendoorn and Kroonenberg, 1988; Main, 1990; van Ijzendoorn and Sagi, 1999)
- Cross-species studies - (Kraemer, 1992; Suomi, 1999; Hofer, 1995; Polan and Hofer, 1999)
- Evolutionary and Life-history theories of attachment - (Chisholm, 1999; Belsky, 1999; Simpson, 1999)

Other infant behaviour that the work described in this thesis might be extended to simulate include:

- The A not-B task (Diamond, 1985, 1991a,b; Diamond *et al.*, 1994; Diamond, 2000, 2002; Munakata, 1998)
- Examples of simple means-end reasoning, such as the intentional use of a tool to bring a toy within reach (Willatts, 1999; Munakata *et al.*, 2002)
- Examples of Joint Attention - (Carpenter *et al.*, 1998; Tomasello, 1999a,b; Butterworth, 2004)

5.2.4 Abstraction

Webb (2001) suggests that a model becomes more abstract the greater number of elements and processes that are present in the target and which have been removed from the model. We can extend this working definition by providing three subtypes of abstraction. Firstly there is ontological abstraction: the things that exist in the simulation are an abstract subset of things that exist in reality. For example, the infant agents do not possess elements such as eyes and hands and they do not carry out vision or manipulation of objects in the same manner as real infants. Secondly there is intricacy abstraction: the things that exist in the simulation are the same as those that occur in reality, but in less complex arrangements. This type of abstraction was dealt with in chapter three where Simon's parable of the ant on the beach was used to explain why infants agents in simply arranged environments move in uncomplicated patterns. Thirdly there is variable scaling abstraction. For example the length of time that time-cycles of the simulation represent varies between the training and test stages of the REUNION scenario.

The scenarios in this thesis present behaviour at different levels of ontological abstraction. The SECURE-BASE scenario is based upon meta-analysis of dozens of studies. It only includes a description of the common aspects of all these studies and therefore doesn't include many concrete details of behaviour. It describes changes in behavioural predispositions that occur over many months of an infants life. Other scenarios, such as the PARK, COY and WARY scenarios described in chapter three represent more concrete behaviours that occur over time spans of a few minutes. The REUNION scenario possesses a mixture of levels of abstraction, the long training stage is more similar to the SECURE-BASE scenario, and the terminal test stage more similar to the scenarios described in chapter three. In the REUNION scenario the observations in the home and test stages of the Strange Situation study were based upon different types of measure. The seventy-two hours of home observation concentrated on measures of infant-carer interaction, and didn't make detailed measurement individual differences in exploration. In contrast, the twenty four minutes of the actual laboratory procedure did include measures of exploration.

The simulation that reproduces the REUNION scenario also possesses variable scale abstraction because the time scales of the training and test stages map with different scales onto the timing of the empirical study (Ainsworth *et al.*, 1978). This is because the test stage only requires a number of time-cycles in the order of hundreds to reproduce the Strange Situation Experiment, with each time-cycle representing approximately a second in time. However, the training cycles only took ten thousand cycles, and ten thousand seconds is only a few hours. For the time scales to not vary between stages of the simulation the training sequences should have taken many millions of time-cycles. A way of thinking about the variable temporal-scale abstraction of these simulations is that in reality events where infants would learn about their attachment relationship would be spread out over a year. The 'filler' between these events that is not relevant to the issue of attachment is not represented in the simulations.

5.2.5 Structural accuracy

The architectures described in this thesis possess many limitations due to structural inaccuracies. These deficiencies have been assessed in more detail in previous sections. Three particularly important structural inaccuracies are:

- The architectures represent attachment in terms of proximity. However, the concept of 'felt security' is richer and more fully captures an infants representation of how it interprets the security provided by its carer (Sroufe and Waters, 1977a). This deficiency is not of crucial importance because proximity forms a large component of felt security and the other components, such as carer attentiveness, can be viewed as being captured in an abstract sense by the infant agent's representation of it Safe-range parameter.
- The **HAR** architecture does not represent the Basal Ganglia's central role in selecting between all actions from different levels. Extensions to the **HAR** architecture which are described in section 5.3.2 are a response to this deficiency.
- The architectures do not possess representations of many physical components, such as eyes, hands, fingers or even bodies with a volume and surface area. These absences are important because they limit the potential for the architectures as they now stand to be applied to other types of behaviour, such as research concerned with physical manipulation.

Simulating infants that are so abstract that they do not possess physical elements such as eyes or hands is a problem because many forms of evidence about what is going on inside infants rely upon mismatches between how the evidence of infant knowledge that eyes and hands transmit:

> *"Problems of inhibition are often problems at the output end. Instead of reflecting deficits in thinking, they reflect deficits in gaining control of one's behaviour so that it reflects what one is thinking. Infants and frontal lobe patients give some indication that they know more than they can express in their behaviour. Frontal patients give such information verbally; infants do so with their eyes."* (page 365, Diamond, 1988)

5.2.6 Performance match

How might we expect a theory to be developed from its beginnings over many iterative improvements? One approach is to conform to a hypothetico-deductive

Newton's theory of the Solar system	**Agent based theories of infant behaviour**
fixed point sun, single point planet	**single point agents with no limbs; errorless perception and action; carer with unchanging strategy; no innate temperamental differences in infant**
both sun and planets revolve round a single centre of gravity	**add limbs to infant; add noise to perception and action; add adaptive strategy to carer**
sun and planets are mass balls not mass points	**add innate temperamental differences in the infant**

Figure 5.1: Comparing Newton's theory of the Solar system with agent based theories of infant behaviour, as examples of positive heuristics in Lakatosian Research Programmes.

method where each iteration of the theories' development gives rise to predictions that are then empirically tested. The theory may be reformed as a result of the new empirical data and produce yet more predictions that drive further development of the theory. There are cases where this form of theory development is not possible. Lakatos (1970) describes how Newton developed his theory of the Solar system without making quantitative predictions that were used to reform and change the course of the theories development. This does not mean that earlier versions of Newton's theory would not provide some interesting predictions. What it means is that Newton's theory developed without any predictions that arose in earlier versions driving the progress of later versions.

Figure 5.1 presents the research in this thesis as a research program, and compares this program with the research program undertaken by Newton in his development of a theory of the Solar system. There are many ways that these two research programs do not match, for example Newton's work forms a major theory in the history of science and the theory of attachment presented in this thesis has much more modest ambitions. However, what the two research programs have in common is a period of development before which strong performance matching of the theory with empirical data can be expected. In addition, the work described in this thesis hasn't proposed a single theory that is being developed but several. Each of these theories is a putative explanation in the sense of being tentative and conjectural. Whilst much of Newton's earlier theoretical components were present in the final theory, there may be many more revisions and corrections in the development of agent based explanations of attachment.

5.2.7 Medium

Schlesinger (2001) reviews a number of agent based models focus of development and assesses that the majority of existing models focus on perceptual and motor aspects of information processing in cognitive architectures. Schlesinger (2001) organises the existing literature into three categories: studies which focus on organism-environment interaction, such as the emergence of tool-use in young infants (Willatts, 1989); studies which focus on sensory self-selection; and studies which focus on perception-action linkages. The natural medium for this kind of model is to implement it in a robot, so that the modeller doesn't have to construct complex agent-environment feedback, they just get it 'for free' in an embodied robot . There also exist a number of software based simulations using autonomous agents that are not focused on perceptual or motor aspects of processing (Tyrrell, 1993; Doran and Palmer, 1995; do Carmo Gadhano, 1999; Logan, 2005). The key difference between these types of simulation and the simulations that do rely upon perceptual and action is the temporal granularity of their interaction with the environment. The smaller the time-scale that needs to be simulated the more important is the inclusion of detailed perception and action mechanisms.

5.3 Future work

5.3.1 Developing new scenarios

The principal benefit of developing new scenarios is that the more behaviours an infant agent is required to reproduce the more constrained its architecture will become. An ultimate ambition for this work might be to produce a 'Unified Theory of Infant Cognitions, and as Newell states:

> *"Developing a unified theory of cognition is a move in the direction of taking more constraints into account. From one point of view, that will no doubt increase the difficulty and the puzzlement. From the point of view of constructing a theory, however, what is needed is more constraint not less. The issue is how to discover where each of the constraints has an impact on the nature of the mind."*

Possibilities for integrating different kinds of source material depend on the level of description of the observations and whether they can be abstracted into a single coherent virtual world. In short term future work each of the sources that might be chosen to be simulated should be a naturalistic set of observations. This is the case with the existing observations, such as the Strange Situation Experiment. This carefully controls the setting in which observations are made, such as the placement of toys and furniture within the laboratory setting. The behaviour of all adults is also carefully controlled. However, once furniture and other objects are placed appropriately the experimenters retire behind a 2-way mirror, and the infant is free to behave as it wants.

Exploring unfamiliar environments or meeting unfamiliar people are naturalistic behaviours that are directed by top level goals. These behaviours arise as a result of the workings of the whole infant architecture and cannot be explained by the workings of a single subsystem. It is the interaction of all subsystems that matters. This property of these naturalistic behaviours means that when adapted for simulation these form what can be termed horizontal microworlds. Clark (1997) contrasts horizontal and vertical microworlds:

> *"A microworld is a restricted domain of study: we can't solve all the puzzles of intelligence at once. A vertical microworld is one that slices off a small piece of human-level cognitive competence as an object of study. Examples include playing chess, producing the past-tense forms of English verbs, and planning a picnic, all of which have been the objects of past AI programs.*[...] *A horizontal microworld, in contrast, is the complete behavioural competence of a whole but relatively simple creature (real or imaginary). By studying such creatures we simplify the problems of human-level intelligence without losing sight of such biological basics as real-time response* [and] *integration of various motor and sensory functions"*

When a set of behavioural phenomena form a horizontal microworld it has the consequence that the natural way to simulate these phenomena is within complete agents, either in simulated worlds or in robots.

5.3.2 Developing new designs

5.3.2.1 Developing the theory by basing all selection processing on the Basal Ganglia

Section 4.3.3.6.3 described how the **HAR** architecture was limited as a theory of action selection because it wasn't closely constrained by what is known about the Basal Ganglia's role in mediating selection at a variety of levels which may

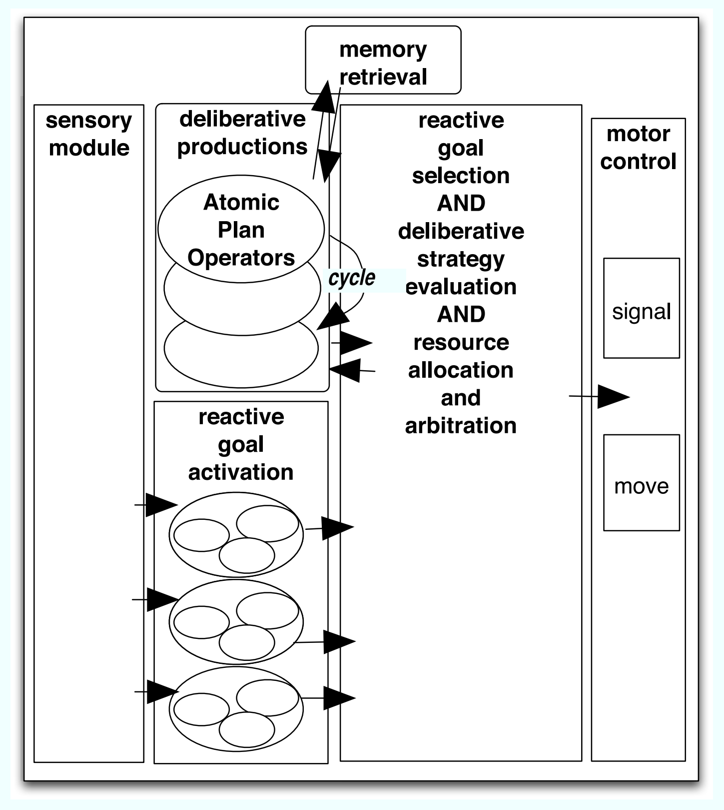

Figure 5.2: A proposed hybrid design that centres information processing around a single selection module that is comparable to the models of the Basal Ganglia put forward by Anderson et al (1998) and Redgrave et al (1999). The Basal Ganglia mediates all selection processes including those involved in reactive action selection and deliberative processes.

include: the behavioural (goal) level; the task/solution level; and the motor action level (Redgrave *et al.*, 1999; Prescott *et al.*, 1999, 2002). Where the **HAR** architecture possesses two subsystems that both carry out a selection function, it might be adapted to be an architecture (figure 5.2) that possesses a single subsystem that is based upon the Basal Ganglia. Figure 5.2 shows how the **HAR** architecture might be developed so that all selection occurs via the Basal Ganglia.

5.3.2.2 Developing the theory by implementing it at a finer and more consistent level of temporal granularity

Section 4.4.4 described how the processing in the **HAR** architecture has not been temporally constrainted. We might expect reactive selection of an action to take less time than deliberative processing which selects the same action, and this is what occurs in the **HAR** architecture. However, the difference in timing of the two systems is currently unprincipled and arbitrary. This limitation in the **HAR** architecture has been addressed during the development of the ACT-R architecture (Anderson and Lebiere, 1999). ACT-R is a cognitive architecture which solves the problem of how to represent selective processes by simulating the Basal Ganglia as a production system. This production system sits at the centre of cognitive operations and mediates all selection between alternative actions. The ACT-R production system selects what action to take according to what information it currently accesses from a number of buffers. These buffers hold information about the state of subsystems which describe the current goal, the contents of declarative memory, perceptual data and the current state of motor modules. The information in buffers is represented as chunks of information, which *"encode small independent patterns of information as sets of slots with associated values"* (page 23, Anderson and Lebiere, 1999). As the ACT-R architecture has developed, the chunks that it can manipulate have become simpler, so that now they are intended to represent 'atomic components of thought' that cannot be simplified any further without losing their power to represent the full capability of symbolic reasoning. Each discrete behaviour of the current ACT-R architecture must be controlled by a separate production, which gives an operational definition of the atomic grain size of cognition. Arbitrary lists or lisp procedures which existed in ACT-R 2.0 have been removed from the operation of production rules in ACT-R 4.0:

> *"Although, syntactically, chunks in ACT-R 4.0 are much the same as the chunks in ACT-R 2.0, they do display one major simplicification over the ACT-R 2.0 chunks. They do not allow for lists of elements to be the values of slots, ... it is not possible to place a coherent psychological interpretation on the list structures - the decision to jettison lists as slots values was key to the emergence of the chunk as an atomic structure"*(page 24, Anderson and Lebiere, 1999)

The ACT-R architecture has been further temporally constrained by the addition of theoretical assumptions for the timing of production firing. The default timing for a production action in ACT-R 4 is 50 msec.

5.3.2.3 Developing the theory by replacing Fregean with Analogical forms of representation

Spelke (2000) has suggested that the forms of representation used in early childhood are qualitatively different to those used later in development. However, Spelke has not specified the Core Knowledge thesis in enough detail for it to be straightforwardly implemented in an agent architecture. However, Spelke's work is consistent with the idea that there is not a simple development of very simple reactive mechanisms to full deliberative mechanisms[2]. What the Core Knowledge thesis suggests is that infants can carry out relatively complex processes, but that these abilities are domain specific. Later they may become accessible to linguistic domain general processes. Therefore, another way to extend the **HAR** architecture is to incorporate qualitatively different representations to the representations used in the **HAR** architecture.

The **HAR** uses Fregean representations to carry out deliberation. It manipulates these representations, which need to possess a number of important properties. The properties of the representations in the **HAR** architecture include:

- semanticity - the infant represents the carer, and properties of the carer with symbols which possess specific meanings, such as the symbol **[carer]** refers to an agent in the infants perceptual field, and the string **[carer is far]** refers to a distance between the infant and its carer;[3]
- arbitrariness - there is no connection between the perceptual data the infant receives on objects and agents, and the form of the symbols used to represent objects and agents[4];
- discreteness - none of the objects or agents are represented with continuous representations;

---

[2] Sloman (2006) describes possible intermediate states between reactive and fully deliberative architectures in: http://www.cs.bham.ac.uk/research/projects/cosy/papers/fully-deliberative.html

[3] It might be argued by a philosophically inclined readers that the infant (one of several entitites in the virtual machine that runs when the program runs) does not refer or understand. Only the writer of the program does, and that's why the 'baby' produces appropriate behaviour (Sloman *personal communication).*

[4] Aaron Sloman responded to the inclusion of this property by noting: "[Arbitrariness] *is often listed as important, but I think that's a complete mistake. It is not *required* that atomic symbols be arbitrary, but if they are that's not necessarily a problem. Making it a requirement rules out sensible things like using atomic symbols whose size, orientation, and shape correspond systematically with low level features in some sensory array."*. In the case of the **HAR** architecture, the sensory and perceptual mechanisms are too impoverished to support the inclusion of non-arbitrary Fregean representations, but in principle they do not need to be arbitrary

- displacement - the infant can represent things that are not currently present, for example when the carer is far, the infant can represent a state of affairs where the carer is close;
- systematicity - the infant's representations possess compositional semantics, so that for example, the plans **[[carer is far][signal-only][carer is near]]** and **[[carer is near][signal-only][carer is far]]** are not equivalent in the same way that **[man bites dog]** is not equivalent to **[dog bites man];**
- productivity - the infants only possess limited number of tokens for representing states of the world and actions to change the state of the world. However, if the number of tokens was increased the mechanisms that manipulate tokens that the infants already possess would allow the infant to represent an infinite number of 'never-before-represented' plans.

Spelke's Core Knowledge might be implemented in infant agent by the incorporation of analogical representations that bestow the ability to carry out limited forms of deliberative processing (Sloman, 1971, 1975). Analogical representations must still possess semanticity, but may not possess arbitrariness, discreteness and displacement. In addition, any compositional semantics may lead to more limited possession of productivity and systematicity. They accomplish the ability to 'look-ahead' and reason about future events without the Fregean forms of representation found in the **HAR** architecture. However, if they are not discrete, they cannot easily be used to explore branching alternative, unless some discretisation is imposed. Instead, an extended architecture might utilise analogical forms of representation that possess a mapping from structural properties of the internal representation held within the infant to the external world in which the infant exists. This mapping may be continuous rather than discrete, and forward simulation of this simulated world within the infant agent would allow the infant to assess the consequences of any of its actions. Analogical representations do not have to be isomorphous, which means that interpretation may not be trivial. In analogical representations there is a mapping from structural properties of the internal to external. For many forms of analogical representation, how the representation works, how it represents meaning, is hugely a function of the interpreter. Context can have a large effect in analogical representations that it may not have in Fregean i.e. the token 'Fred' is less sensitive to the context in which it is placed than some aspect of an image[5]. There are many subtypes of analogical representation.

---

[5] However, as Sloman (2006, *personal communication*) notes *"there are huge amounts of contextualisation in verbal descriptions, questions, instructions. Components of computer programs in modern programming languages, are hugely context-dependent, one aspect of which is the use of local variables: remove the local variable declaration, or extend its scope, and the interpretation of what the variables mean can change a lot."*

For example, analogical representations may possess precise metrics, metrical properties that are not absolute, non metrical but topological properties, or variable dimensionality. A number of recent examples of non-Fregean deliberative systems have been implemented (Shanahan, 2006; Cruse, 2003; Grush, 2004) . However a limitation of these methods is that these representations possess fixed sized vectors, so that any representation of future states has the same capacity as the initial representation. Another way of organising look-ahead is to have a variable structure such as a plan, sentence or other substructures that may each be expanded.

5.3.3 Making your own luck: top-down versus bottom-up research

This thesis has attempted to explain infant attachment behavioural phenomena, particularly behaviours observed in Strange Situation studies. The research approach taken in this thesis has been to gain understanding about infant behaviours by designing autonomous agents to reproduce those behaviours. Behaviours such as exploring a park, coyness with stranger people, wariness in strange environments and the particular behaviours shown by Secure, Insecure-Avoidant and Insecure-Ambivalent in Strange Situation studies have all been clarified. However, there exist contributors who question the importance of the Strange Situation and the attachment styles that are derived from this procedure:

> *"The claim* [that the Strange Situation acts as a shorthand for the quality of the infant-mother relationship] *fails the test of reasonableness. The mother and infant, who have been together for over a year, have experienced pain, pleasure, joy, and distress, and the infant's representations of and behavioural reactions to the mother must contain aspects of all these expriences. Is it reasonable to believe that a half-hour sample of behaviour in an unfamiliar laboratory room could reveal the history of all these experiences with the mother? Could any thirty-minute observation uncover psychological products created from over six thousand hours of interaction between these partners?"* (page 99, Kagan, 2000).

Kagan's criticism of the Strange Situation is rebutted by the results from this thesis. The architectures implemented in this thesis have shown that particular parameters within infant architectures, such as the safe-range parameter or the infants measure of comfort in close carer contact, may not be apparent in the infant's behaviour for much of the time. However, in particular circumstances, where the activation levels of contrasting goals such as proximity seeking and contact avoidance are finely balanced, these parameters can have dramatic effects on overall behaviour. The effect of these parameters is an example of Braitenbeg's claim in the introductory chapter that synthesising theories as simulations can give rise to simpler theories.

One of the most successful elements of the results produced in thesis are described in chapter four, where it was demonstrated that the GLA architecture interacted with the carer in positive feedback loops so that distinct categories of infant attachment behaviour emerge from initially identical infants and carers. Allowing for unexpected side effects is an important benefit of researching infant behaviours by simulating infant agents. A researcher can set the initial conditions of an system that is formed from interacting subsystems and then leave the subsystems to interact. Attempting to do research this way is a form of 'bottom-up' reverse engineering that is mirroring the process of evolution. As Dennett notes:

> *"The process of evolution is notoriously lacking in all foresight; having no foresight, unforeseen or unforeseeable side effects are nothing to it; it proceeds, unlike human engineers, via the profligate process of creating vast numbers of relatively uninsulated designs, most of which, of course are hopelessly flawed because of self-defeating side effects, but a few of which, by dumb luck, are spared that ignominious fate. Moreover, this apparently inefficient design philosophy carries a tremendous bonus that is relatively unavailable to the more efficient, top-down process of human engineers: thanks to its having no bias against unexamined side effects, it can take advantage of the very rare cases where beneficial serendipitous side effects emerge."*(page 684 Dennett (1994))

Designing agents to do 'bottom-up' research on infants can therefore be expected to *"improve the epistemic position of researchers by opening up different regions of design space"* (page 685 Dennett (1994)). Doing 'bottom-up' research can also guard against constructing internal subsystems that are assumed to be innately constructed and whose only role is to reproduce observed behaviour. This type of modelling is similar to the 'phenotypic gambit' that is made by come evolutionary biologists. The phenotypic gambit is the *"simplifying assumption that the relationship of the genotype to the phenotype is not especially important for understanding adapation"* (page 30, Chisholm, 1999). It is undermined by studies which show that observable behaviours can arise from non-obvious underlying mechanisms. For example, Belyaev *et al.* (1985, cited in Chisholm (1999)), showed that during artificial selection of foxes, Russian fur farmers selected for two strains of fox that were distinguished by their level of aggression. However, it was not a gene for aggression that had been selected for, but the length of the developmental window during which the foxes underwent primary socialisation. The shorter this window, the more aggressive the foxes became. Therefore the phenotype of aggressiveness was only indirectly linked to the genotype, there was no 'gene' for aggressiveness. Polan and Hofer (1999) also provides many examples of 'hidden regulation' in the control of attachment behaviours in rats. The architectures described in this thesis have many 'top-down' aspects, but future work may incorporate more 'bottom-up' mechanisms that allow development of mechanisms that utilise beneficial yet unexpected interactions within architectural subsystems.

5.4 Final thoughts about designing agents to understand infants

This thesis has described a number of empirical studies related to infant attachment behaviour which have been adapted into autonomous agent scenarios (sections 2.2.1, 3.2.1, 4.2.1). Future work may involve incorporation of a broader range of behavioural phenomena. These may include more detailed infant search behaviours, and attachment behaviours in older children and adults (section 5.2.3).

Attachment can be studied as a cross-species, cross-cultural, behavioural and physiological set of phenomena, that changes across the life-span and may be internally linked to a variety of information processing structures and mechanisms. Attachment phenomena can be studied statistically and as a dynamic system (section 3.4). One of the most significant contributions of this work is that it shows how these varying facets of Attachment Theory can be integrated within simulations. The inclusion of all these elements in simulations may allow evolutionary interpretations of infant attachment behaviour to be validated by running evolutionary simulations (section 4.4.6). The introductory chapter suggested two types of inter-disciplinary contribution: that this work may act as a Transferring Publication; and that this work may form an Interfield Theory. Acting as a Transferring Publication is a more modest contribution. Researchers in Attachment Theory or in Autonomous and Multi-agent systems may simply read this work and take back any lessons they have learnt to their own fields. If this work were to act as the foundation for an Interfield Theory then the design of agent architectures would need to act as a long-term framework within which different types of observation would be integrated. For example, agent architectures might allow physiological measures, such as heart-rate or cortisol level, to be integrated with observations of overt behaviour in a deep explanatory manner.

Several new empirical studies have been suggested by the process of creating simulated infant agents. These include adapting the Strange Situation so that reunion occurs at a variety of distances (the Sports-hall Situation Experiment, in section 4.4.2) and observing how consistency of home caregiving affects infant attachment style.

A number of different architectures has been designed and implemented. These include the **RAL** architecture which minimises the amount of representational complexity that infants may use to direct their behaviours and the **HAR** architecture which possesses more complex representations. Future work may involve implementing both these architectures at a lower level of description. In the case of the **HAR** architecture this might mean implementing it with a neural network that formed analogical representations, or with action selection carried out with a production system that was constrained to operate within the temporal bounds of the Basal Ganglia (section 4.3.3.6.3).

This work doesn't just show what agent design can do for the differing subfields of Attachment Theory. The domain of attachment, in humans and other species, and in infancy to adulthood, provides a valuable 'test-bed' for comparing the performance of different kinds of agent designs to explain a diverse set of observations. Autonomous agent architectures may provide a valuable framework for studying ontogenetic theories of infant representational change (Karmiloff-Smith, 1994), particularly those that link representational change to executive functions (Russell, 1996).

There are several ways that this work may be extended: it may be developed further as a psychological theory; it may be incorporated within robot control architectures; and it may become an educational tool. In the introduction the theories proposed in this work were described as putative, in the sense of being tentative and needing further development. It has many limitations, and is incomplete in many respects. The high level of description of the infant agents, which lack hands or eyes, facilitated rapid initial development of the simulations. These deficiencies now need to be overcome to extend this work. However, this work has demonstrated a capability to engage in original research and advance knowledge across several disciplines. Prior to this work theories that attempted to explain attachment phenomena were expressed in words and diagrams only. This work is the first realisation of Attachment Theory that is precise and detailed enough to be implemented in a computer program. This is a major step forwards because it allows rigour in hypotheses to be enforced, and this rigour reveals logical flaws or inconsistencies in theories that otherwise may not be apparent (page 44 (Elman *et al.*, 1996)). Designing and implementing theories of infant attachment behaviour is also a valuable pursuit because it changes the way we think about questions in the field of Attachment Theory. Simulations can stimulate empirical research by demonstrating how behavioural phenomena can emerge from dynamic systems.

An example of this are the dynamic feedback loops that occur in simulations of the **GLA** architecture (section 3.4). Emergent phenomena like these have been illuminated because this work did not merely conduct analysis, but in addition took a synthetic, design based approach.

Appendix A

The simagent toolkit

Section 1.7 notes that all the simulations described in this thesis have been produced in POP11 using the sim-agent toolkit[1]. The simagent toolkit was created for exploring information processing architectures of the type produced in this thesis. The infant, carer and stranger agents and food, toy and furniture objects have all been adapted from a set of base classes that are included within the toolkit library. The toolkit also possesses scheduling mechanisms for running simulations involving any number of these objects and agents. The internal mechanisms of the infant, carer and stranger agents are implemented as multiple rulesets or rulefamilies, using Poprulebase.

> *"Poprulebase* [is] *a very general, flexible and extendable interpreter for condition-action rules, written in Pop-11. It is based on the idea of a forward-chaining production system interpreter, but provides a collection of unusual facilities, including a smooth interface to neural net or other "sub-symbolic" mechanisms, through so-called "filter" conditions, a wide variety of condition types and action types (both user-extendable), and a variety of control options and higher level structuring facilities.*
>
> *Poprulebase allows rulesets to be combined into rulefamilies. Within a rulefamily, control can be transferred between rulesystems as the context changes. This also allows SOAR-like pushing and popping of contexts, as well as other sorts of transitions. Rulefamilies can be combined into rulesystems.* [...]
>
> *In simagent, each agent has its own rulesystem which may consist of a collection of rulesets and rulefamilies implementing a variety of mechanisms (perception, memory, motive generation, planning, self-*

---

[1] http://www.cs.bham.ac.uk/research/poplog/packages/simagent.html

> *monitoring, etc.). A feature of the toolkit is provision for differing resource allocation between components of an agent, and between agents. E.g. some agents may plan quickly relative to their perceptual processes, others slowly."*

Within each infant and carer agent is a rulesystem which consists of a collection of rulesets and rulefamilies. These implement a variety of mechanisms including perception, memory, goal generation and actions. In the case of the **HAR** architecture a reactive component is implemented alongside a deliberative subsystem. The deliberative subsystem is held within a rulefamily that can be run more quickly or slowly. The simagent toolkit allows the implementation of an architecture that requires the ability to repeatedly change the speeds of multiple subsystems during run-time.

Appendix B

Further analysis of the solutions

B.1 Comparison of the GS architecture with a model of Contention Scheduling

Section 2.3.1.5 describes the **GS** architecture in some detail. The status of the **GS** architecture in the broader field is here clarified by comparing it with the somewhat similar Contention Scheduling model of Cooper and Shallice (2000).

The Contention Scheduling mechanism described by Cooper and Shallice (2000) is not presented as a complete ‘stand-alone’ architecture, whereas the **GS** architecture is. The Contention Scheduling mechanism sits below a supervisory system that monitors progress and modulates the processes in the Contention Scheduling level. Both the Contention Scheduling system and the **GS** architecture are composed of subparts. The Contention Scheduling system is divided into three parts: a schema network, an object network and a resource network. Each of these networks is hierarchically layered.

The schema network consists of goal directed schemas and goals. Each schema is made up of partially ordered subgoals, and each goal is in turn composed of schemas, any one if which may be used to achieve the goal. The object network represents the state of perceivable objects in the external environment and the resource network represents the agent’s means of manipulating the objects in the external environment. The model of Contention Scheduling is implemented to carry out tasks in the *coffee preparation domain*, which might involve the model putting sugar, milk, grinds and water together to make a coffee. In this domain the top goal may be **Have a coffee**, which activates a top level schema which is *Prepare Instant Coffee*. The first goal in this top level schema might be **Sugar into Coffee**. which then possesses two alternative schemas, *Add Sugar from Packet* and *Add sugar from Bowl*. Each of these schema may possess a string of goals that needs to be achieved in the correct order, such as the goals: **Hold, Discard, Open, Transfer**. Each of these lower level goals will direct the base-level schemas, such as: *Pick-up, Put-down, Tear, Unscrew*. Whilst the schema network is directing base level actions it communicates with the Object and Resource networks so that is only attempts to manipulate objects and resources that are available.

In both architectures different goals are generated and compete in a winner takes all contest to see which goal becomes active and direct the external behaviour of the infant. So in both architectures the action schemas or goals have a state which is selected or unselected. Both architectures link perception to action, and the sensory and action mechanisms in the **GS** architecture are similar to the object and resource networks of the model of Contention Scheduling. However, the architecture for Contention Scheduling links perception to action via a hierarchically organised network of action schemas, which represent goals and multiple levels of subgoals. The **GS** architecture is flat, with one level of goals activating a single level of atomic actions. Selection of schemas and goals, and the consequent triggering of further selections go up and down the hierarchy in the model of Contention Scheduling, whereas activation merely passes across the **GS** architecture.The majority of the differences between the models derive from this key distinction. For example, the action schemas in the model of Contention Scheduling have a greater number of influences on their activation, in the form of top-down activation and lateral activation.

Monitoring goal achievement is more important in the model of Contention Scheduling than the **GS** architecture because the mechanism has to keep track of what actions have been taken, for example in the ‘making a sup of coffee’ domain the architecture has to flag when sugar has already been added. The breaking of this ‘flagging’ mechanism gives rise to action slips such as putting too many sugars in to a cup of coffee. In the **GS** architecture the actions don’t have this requirement for correct scheduling. Following the externally observed actions of the model of Contention Scheduling in the ‘making a cup of coffee’ domain we might see a string of actions such as *pick-up spoon, put in coffee granule container, put granules in cup, put down spoon*. The externally observed actions of the **GS** architecture would give a string such as *move to carer, move to toy t1, move to carer, move to toy t2*. The actions represented in the **GS** architecture are more abstract than those found in the model of Contention Scheduling. The model of Contention Scheduling is representing tasks such as making a drink in greater detail for a relatively short period of time.

The goal activators in the **GS** architecture possess more complex goal activators with heterogeneous structures than the individual goal nodes and schemas in the model of contention scheduling. Aside from the internal details of these goal activators, the **GS** architecture is much simpler. The simplicity of the goal representation in the **GS** architecture allows it to act as a more abstract representation that supports a simulation of the learning that occurs over the course of the infant's first year of life. This potential is exploited in the architectures described in the next two chapters.

B.2 Comparison of the level of description in the simulation with other work

We have chosen source material to adapt into a simulation. The behaviours are framed as a problem to be solved, so we need to find the right level of abstraction to represent the problem.The level of description taken by this work is established by comparing it to the three different levels of description detailed by Cooper and Shallice (2000). The lowest of these three levels is at the level of *motor response schemas* (Arbib, 1985). At this level operations such as 'movement by stepping' would be broken down into its component parts, such as 'lifting feet' and 'bending knees'. Next up in abstraction is the level that the model of *Contention Scheduling* is directed at (Cooper and Shallice, 2000). At this level operations such as cleaning one's teeth or starting a car can be represented. Representing movement at this level would differentiate between operations such as 'crawling' or 'stepping' even though both actions may be carrying out the same general function, such as changing position within a room. The highest of these three levels is at the level of *scripts* (Schank and Abelson, 1977) or *Memory Organisation Packets* (MOPs) (Schank, 1982). This level represents activities such as going to a restaurant or visiting a doctor's surgery. The sub-units of these activities can be separated by other activities such as reading a book or making a phone call.

For the particular patterns of behaviour being investigated in this thesis the lower level details of sensory modalities and motor actions are not relevant. What is needed in the design and implementation are mechanisms of sensation and action that fulfill the appropriate functional role, at the appropriate level of description. For example, the important aspects of the infant's movement are not how the infant moves from position to position, but how far from the carer or towards what other object the infant moves. Therefore in this project the source material has been interpreted at a level intermediate between that of the model of *Contention Scheduling* (Cooper and Shallice, 2000) and the higher level of *scripts* or *Memory Organisation Packets*. The consequences of this decision can be seen in the architectural solutions that are presented in this thesis. These are relatively flat, in that they do not possess a deep hierarchy of goals.

Appendix C

Formal description for reinforcement learning

To produce a formal description of the reinforcement function this work has followed an Artificial Intelligence (AI) approach to Reinforcement Learning (Sutton and Barto, 1998; Wyatt, 2002). In the training stage of the simulation there are repeated instances where the carer agent goes beyond the Safe-range limit, is called back by the infant agent, and then responds promptly or otherwise. Following an AI approach to Reinforcement Learning, the reward that the infant agent picks up at time t is defined as r_t. This is a measure of immediate reward utility/value/worth. Normally in AI style Reinforcement Learning, a notion of long term utility is also defined. Long term utility is called return ($\underline{r_t}$). Long term utility and return is what the infant agent is attempting to maximise by its behaviour. Return is often defined as a discounted sum of the future rewards.

$$\underline{r_t} = \sum_{k=0}^{\infty} g_k r_{t+k} \tag{C.1}$$

where g is the discount factor at time $t + k$.

In our case we are going to alter this problem specification in two ways:

- we will allow two types of immediate reward
- we will introduce a novel discount rule for each type of reward

C.0.0.1 Allowing two types of immediate reward

If we define r_t^- as the immediate negative reward (or punishment) at time t, and r_t^+ as the positive immediate reward, then the return can now be a function:

$$\underline{r_t} = \sum_{k=0}^{\infty} g_k^- r_{t+k}^- + g_k^+ r_{t+k}^+ \tag{C.2}$$

where g_k^+ is the discount function for the positive rewards and g_k^- is the discount function for the negative rewards. The benefit in having two discount rules is that we can independently vary the ‘pleasure’ that an infant agent experiences when it receives a response from the carer agent from the disappointment the infant agent experiences each time slice that it has signalled but received no response.

C.0.0.2 Introducing a novel discount rule for each type of reward

C.0.0.2.1 The negative discount rule There are two ways to think of the negative discount function. One way is that each time slice that the carer does not respond is a small punishment. Another way is that the infant agent updates the punishment signal at the same time that it updates the reward signal. There are two elements of the negative reward that we can vary. Firstly, there is the size of the actual punishment signal that the infant agents receive at each time slice. This might have been allowed to vary by being set differently in different contexts in the simulation. However, in the currently implemented simulation this is set at -1.

$$r_k^- = -1, \forall k \tag{C.3}$$

The next part of the punishment aspect of the reinforcement learning algorithm is how the reward is discounted. The method that has been chosen is:

$$g_k^- = \rho, \forall k \tag{C.4}$$

For any given timeslice, if the infant agent has previously signalled, it will experience the same ‘disappointment’. The increase in the punishment signal received going from ten to eleven cycles will be the same as that between twenty and twenty one cycles.

C.0.0.2.2 The positive discount rule Wyatt (2002) describes two common discount schemes used in Reinforcement Learning in the field of Artificial Intelligence. The most widely studied measure of return is an infinite horizon model termed the *geometric discount* model of return. This model is particularly used in work in learning from delayed reinforcement learning.

$$g_k = g^k, \text{where } 0 \leq g < 1 \tag{C.5}$$

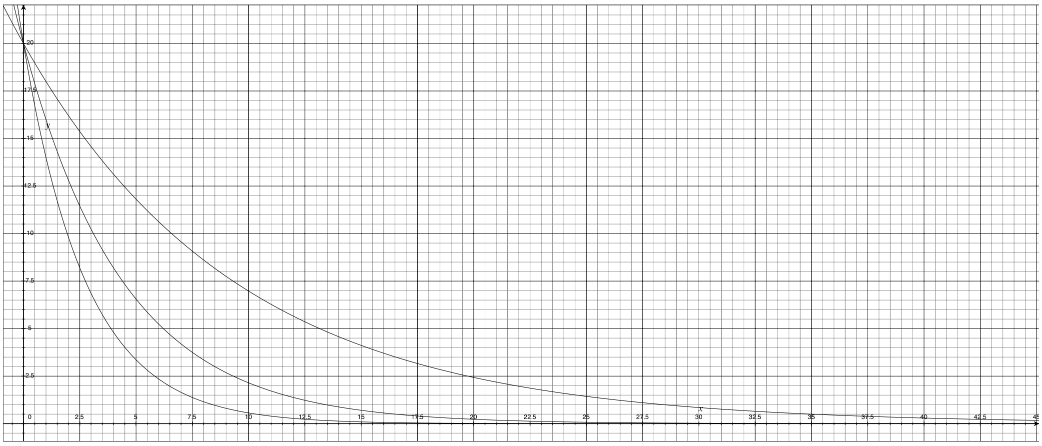

Figure C.1: The geometric discount scheme in equation C.5 operating with three different values of g. From highest to lowest lines the values of g are 0.9, 0.8 and 0.7. When the values for g are higher, long term reward becomes more important,

The infant agent in the simulation is required to learn from a delayed reinforcement so this discount model is initially attractive. If we first consider positive rewards, with this discount scheme the longer the infant agent waits for a response, the smaller the positive reward gets. The value of g determines the relative weighting of short term and long term rewards. Figure C.1 shows that as $g \rightarrow 0$ short-term reward becomes more important. This figure also shows that with this discount model, whatever the value of k, discounting starts immediately, and the greatest discount is between the first and second time slice. This is not the behaviour we require in a discount function. The duration of a timeslice in an autonomous agent simulation is somewhat arbitrary. It could represent tens of seconds or milli-seconds. Therefore we don't want to have the biggest discrimination always made between the first and second timeslices.

A discount model which overcomes the problem of steep initial discounting is the *finite horizon* model of return, where the horizon is a finite number *h* of steps into the future,

$$g_k = \begin{cases} 1 & \text{if } k \leq h \\ 0 & \text{if } k > h \end{cases}$$

However, this gives a sharp step function that seems a little implausible for an infant's learning. It is unclear how a drop from maximum to minimum reward over a single timeslice might be accomplished in real infants. In addition, the simulation may slice time thickly or thinly. The discount scheme should be able

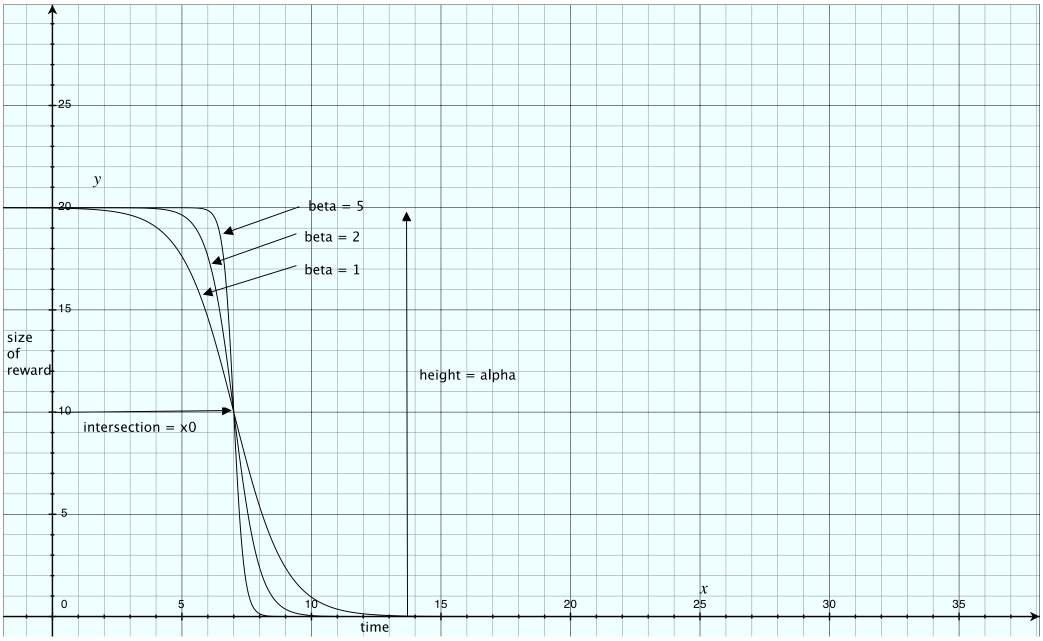

Figure C.2: The novel discount scheme in equation C.6 operating with three different values of β showing how the sharpness of the step function can be varied.

to cope with variation in how thinly time is sliced in the simulation.

A function that combines properties of the geometric and finite horizon discount schemes is

$$g_k^+ = \frac{\alpha}{1 + e^{\beta(k-(k^m)}} \tag{C.6}$$

Figure C.2 shows the shape of the discount scheme that has been derived from equation C.6 and chosen for the positive reward. It combines the positive aspects of the geometric and finite horizon discount schemes. Figure C.2 shows that if the β constant is set very high then this discount scheme tends towards the step function given by the finite horizon discount scheme.

The equations C.7 and C.3 both show that the infant receives a reward and punishment of the same size, here given an amplitude of one. In future work this might be developed so that depending on how the carer responded or didn't respond the amplitude of the reward and punishment signals would change.

$$r_k^+ = 1, \forall k \tag{C.7}$$

C.0.0.3 Putting it all together

Since the rewards and punishments are all set to have a value of one, equation C.2 can be rearranged to become equation C.8:

$$\underline{r_t} = \sum_{k=0}^{\infty} g_k^{+} - g_k^{-} \tag{C.8}$$

Substituting in the details of the positive and negative discount schemes we get the total return

$$\underline{r_t} = \frac{\alpha}{1 + e^{\beta(k-(k^m)}} - \rho k \tag{C.9}$$

The constants α and ρ control the maximum rewards and punishments, respectively, that the infant receives. The constant β sets the gradient of the decrease of the positive reward, a small positive value for β produces a gradual decrease in the size of reward each time step, higher positive values for β produce decreases in reward that approach a step function from maximum to minimum reward over small periods of time. The time step where the steepest decline in reward occurs is set by the terms k^m, which is the minimum possible time in which a carer could respond. Its inclusion means that infants do not expect carers to respond faster than the laws of the virtual world allow. When the time elapsed from bid to reward is close to k^m the positive reinforcement is large. Another way of saying this is, prompt responses give large reinforcement signals and this results in the Safe-range distance being increased. When responses take more time reinforcement signals become negative ‘punishment’ signals. In this context ‘punishment’ doesn’t mean that the actions that were taken are less likely to be taken in future. Quite the opposite occurs. The Safe-range limit is reduced by the value of the punishment signal. Therefore distances that are previously considered by the infant to be safe, are now beyond the Safe-range distance. The carer still has to forage and may still need to go as far afield in the future, so the chances are that after a decrease in Safe-range the carer will be less responsive in future.

This positive feedback mechanism, operating over a long training period, may be what drives the infant-carer pairs into the Secure/Insecure clustering seen in the Strange Situation studies. A carer whose performance is initially intermediate between Secure and Insecure may come to be perceived as at either extreme of caregiving.

Appendix D

Detailed description of infant attachment behaviours

D.1 The 82 measures of behaviour reported in Ainsworth *et al* (1978)

Ainsworth *et al.* (1978) made a large number of different measures of infant and carer behaviour. These are listed in figure D.1.

Mother at home Ist quarter	**Mother at home 4th quarter**	**Infant at home Ist quarter**	**Infant at home 4th quarter**	**Strange Situation at 1 yr**
Ignoring of crying	Ignoring of crying	Freq of crying	Freq of crying	Crying
Duration crying not responded to	Duration crying not responded to	Duration crying	Duration crying	Avoidant behaviour
Mean duration of pick-up episode	Acknowledge on entrance	Positive response to being held	Positive response to being held	Resistant behaviour
Affectionate pick-ups	Affectionate pick-ups	Negative response to being held	Negative response to being held	Exploratory manipulation
Tender, careful holding	Tender, careful holding	Positive response to put-down	Positive response to put down	Exploratory locomotion
Provide unpleasant physical experience	Abrupt/interfering pick-ups	Negative response to put-down	Negative response to put down	Exploratory vision
Aversion to physical contact	Freq of verbal commands	Terminates face-to-face	Initiate pick-up	Maintaining contact
Inept holding	Inept holding	Bouncing	Initiate put-down	Oral behaviour
Contingent pacing	Routine manner	No response	Sink-in	Smiling
Silent, unsmiling	Freq of physical interventions	Vocalising	Active contact	Looking
Routine manner	Maternal rigidity	Smiling	Tentative contact	Vocalising
Quantity/ending of feeding	Lack of emotional expression		Separation distress	Search behaviour
Handling the babies preferences	Sensitivity		Follows when separated	Seeking proximity
Synchronisation of feeding rate to intake	Accessibility		Positive reunion	Distance interaction
Timing of feeding	Cooperation		Mixed reunion	Oral behaviour
Maternal rigidity	Acceptance		Communication	
Lack of emotional expression			Compliance	
			Anger	

Figure D.1: A table showing all the kinds of behaviour that were coded in the original Strange Situation study (Ainsworth *et al.*, 1978), organised into five lists according to when the behaviours were observed and who performed the behaviours.

D.2 Further details on the distinction between the A, B and C attachment groups

Figure D.2 breaks observations down into three sections. The first section describes between-group differences in carer behaviour at home. The groups that the carer were put into to gain this data were gained from analysis of their infant's behaviour in the Strange Situation. From this data will be derived the Independent Variables for the simulation. The next section describes between-group differences in infant behaviour at home. This data is to be treated as dependent variables. The differences in this data should emerge from the infant's interaction in their particular caregiving environments. The last section describes the group differences that are measured in the Strange Situation. It is these differences that are used to categorise infants and their carers into A, B or C categories.

EMPIRICAL BEHAVIOUR all measures observed at home unless 'in SS' (in Strange Situation)	**GROUP DIFFERENCES**
CARER BEHAVIOUR AT HOME THAT DIFFERENTIATES B FROM (A/C) GROUPS	
tender, not inept, affectionate contact	B > A/C
carer response to crying	B > A/C
carer response to reunions	B > A/C
CARER BEHAVIOUR AT HOME THAT DIFFERENTIATES A FROM C GROUP	
aversion to contact, unpleasant contact	B/C < A
abrupt interferring contact	B/C < A
lack of emotional expression and rigidity	B/C > A
INFANT BEHAVIOUR AT HOME THAT DIFFERENTIATES ALL GROUPS	
infant anger/resistance	A > C > B
INFANT BEHAVIOUR AT HOME THAT DIFFERENTIATES B FROM (A/C) GROUPS	
infant crying	B < A/C
positive affect in response to pick up/put down	B > A/C
INFANT BEHAVIOUR AT HOME THAT DIFFERENTIATES A FROM C GROUP	
sinking in during contact	B/C > A
infant movement in response to separations	A/B > C
INFANT BEHAVIOUR IN THE STRANGE SITUATION THAT DIFFERENTIATES ALL GROUPS	
infant crying in SS	A < B < C
infant movement in exploration in SS	A > B > C
INFANT BEHAVIOUR IN THE STRANGE SITUATION THAT DIFFERENTIATES A, AND (B/C) GROUPS	
infant crying at separations in SS	A < B/C
proximity seeking (particularly in reunion) in SS	A < B/C
infant treatment of strangers in SS	A > B/C
infant avoidance at reunions in SS	A > B/C
INFANT BEHAVIOUR IN THE STRANGE SITUATION THAT DIFFERENTIATES C, AND (A/B) GROUPS	
infant anger/resistance at reunion in SS:	C > A/B

Figure D.2: Table showing a selection of empirical measures of home and Strange Situation carer and infant behaviours particularly relevant to the distinctions made in the REUNION scenario. Only statistically significant group differences have been included ($p < 0.05$)

D.3 Four crucial dimensions of coding in the Strange Situation

Four dimensions of behaviour in the reunion episodes of the Strange Situation were found to be crucial in distinguishing the classificatory groups. These were: proximity and contact seeking behaviour, contact maintaining behaviour, avoidance and resistance.

Proximity and contact seeking was scored such that:

> "[Highest scores were given where] *the baby had shown most active initiative in seeking proximity to an adult - approaching without delay and without needing to be invited, and approaching fully to make contact, clambering up on the adult without needing to elicit her cooperation.*[Lowest scores were given where] *there seemed to be no overt effort to gain proximity and no behaviour that seemed to be a clear-cut signal inviting the adult's approach.* [Somewhat higher scores than the lowest] *included behavioural episodes in which a baby made an 'intention movement' toward a person - a slight and incomplete approach"* (page 51 Ainsworth *et al.*, 1978)

Contact maintaining behaviour was scored such that:

> *The highest scores are given to babies who repeatedly resist release and who, as a consequence, succeed in maintaining physical contact throughout most of the episode in question. Resisting release implies intensified clinging when the adult attempts to put the baby down (or merely to shift his position), or turning back immediately to clamber up again when put down. Mere protest, without active effort to maintain contact, is scored lower.* (page 53 Ainsworth *et al.*, 1978)

Avoidant behaviour was scored such that:

> *Highest scores are given to infants who persistently ignore their mothers, continuing to play without acknowledging mother's return despite here effort to invite the baby's approach. Somewhat lower scores are given to infants who mingle greeting responses with moving away, turning away, or looking away.* (page 53 Ainsworth *et al.*, 1978)

Resistant behaviour was scored such that:

> *The highest scores are given to babies who persistently manifest intense angry and/or resistant behaviour to an adult. The resistance is shown by pushing away from, striking out at, or squirming to get down from an adult who has offered contact, or by pushing away, throwing away, or otherwise rejecting toys throuhg which an adult attempts to mediate interaction. The highest scores imply an obvious angry emotional tone, although in lower scores the resistant behaviour may be seemingly without negative affect. Resistant behaviour is not incompatible with proximity seeking or contact maintaining. An angry, resistant infant may nevertheless strongly seek to gain and maintain contact, although such a combination suggests ambivalence* (scoring scheme for resistance in the Strange Situation, page 53 Ainsworth *et al.*, 1978)

D.4 Describing subtle aspects of Avoidant infant behaviours in more detail

Secure and Ambivalent patterns of behaviour do not possess the kind of covert behavioural characteristics that Avoidant patterns possess. Main and Weston (1982) review a number of papers which describe Avoidant behaviour as a general phenomenon in human infants:

> *"Social avoidance* [...] *refers to movements directed away from a prospective social partner - to, for example, gaze aversion, movements away of head or upper body, turning the back, and moving out of contact of immediate proximity. It also refers to failure to communicate acts of the social partner and even to apparent lack of recognition of a familiar partner.* [It] *is surprisingly prevalent in human infants reunited with their attachment figures following long and stressful major separations, and in rejected infants it appears following even very brief separations. Avoidance of the attachment figure presents as something of a problem for attachment theorists, simply because it* **is** *relatively common. Are we to presume that the attachment system breaks down so easily that mildly rejected infants regularly avoid rather than approach their mothers?"* (page 31, Main and Weston, 1982)

Avoidance ranges from clear and obvious examples, such as turning and moving the entire head or body away, to more subtle and mild forms of rejecting behaviour, such as:

> *"a certain expression of the eyes* [that] *is often described as an 'empty' or 'blank' expression, also a 'letting the mental shutters down' - it is a vague, expressionless look, often aimed slightly past the adult's eyes ...This initial response of the child is, if the observer keeps looking at it, followed by partial or complete closing of the eyes. This can be, but is not always, very slow; sometimes so slow that one has the impression that 'the eyelids have a very long way to go'. When closing the eyes is total, the eyelids look completely smoothed out - not even slightly creased, pressed, or 'screwed up' as eyes look when closed on other occasions except when in sleep"* (page 179, Tinbergen and Tinbergen (1973), quoted in Main and Weston, 1982, page 35)

Main and Weston (1982) describe how Avoidant infants come to seem as if they are exploring when their principal underlying goal is assessed as avoidance of their carer:

> *"infants viewing the parent across a distance cannot maintain simple gaze aversion for long. The infant again looks toward and perhaps greets the parent, or turns fully about and moves away, or searches in a rather disorganised way for something to do with his hands. Often he seizes upon a nearby inanimate object. This search for something to do (seize, touch, handle) is striking. A close repeated viewing of films of the apparent 'exploratory' behaviour that occurs immediately on reunion reveals that the inanimate object generally has far from the infant's full attention. The infant may, for example, turn rather frantically toward a table leg and finger it (but with eyes fixed blankly on the wall ahead), or in a clumsily decisive move, he may suddenly drop a toy in to a box but then (staring straight ahead) close the box on his hand. The general impression is that the infant could not succeed in maintaining his avoidance of the parent without the aid of the seized-upon object"* (page 35, Main and Weston, 1982)

> *Researchers have measured physiology in an attempt to understand the competencies of the young infant, and have inferred the presence or absence of particular affective states on the basis of these physiological responses. These studies are usually based upon certain assumptions about the role of physiological systems in the generation of behavioural responses. Such assumptions are the product of a long history of research in psychophysiology with adults and older children.* (page 228, Fox and Card, 1999)[1] .

---

[1] Describing how psychophysiological researchers validate their inferences about the affective

D.5 Cross-cultural attachment studies

Cross-cultural patterns of attachment are important when considering if the original Strange Situation study on which the REUNION scenario is based provides conclusions about an architecture that is common to all human infants, or is specific to just infants brought up in some cultures. A meta-analysis of cross-cultural patterns for 2000 Strange Situation classifications across 8 countries found that the original Baltimore study, and most other studies (from countries including the US, China, West Germany, Great Britain, Netherlands, Sweden and Japan) fitted into a group where about two thirds of infants were assigned to the Secure (B) category, a fifth assigned to the Avoidant (A) category and an eighth to the Ambivalent (C) category (van Ijzendoorn and Kroonenberg, 1988). Studies with statistically outlying distribution patterns included: Israeli and US studies with elevated proportions of Ambivalent infants; a West German study with higher numbers of Avoidant infants; and a Japanese study where the number of Avoidant infants was lower than the international average and the number of Ambivalent infants higher. This data suggests that the three main styles of attachment are not derived from culture specific experiences in infancy, though the proportion of infants possessing each style is affected by cultural factors.

---

state of infants from physiological measurements is beyond the scope of this thesis, and is described in Fox and Card (1999).

Appendix E

Further evolutionary analysis of function

Section 4.2.3 analyses the possible evolutionary function of different styles of infant attachment behaviour. The sections in this appendix extend this analysis.

E.1 Details of different possible adaptive benefits of avoidant infant behaviour

E.1.1 Maintaining proximity by stopping the carer from moving away

Main and Weston (1982) propose a theory that avoidance acts to maintain proximity by stopping the carer moving away. This theory was adapted from an ethological theory originally proposed by Tinbergen and Moynihan (1952, cited in Main and Weston (1982)).

> *"Visual avoidance has been studied extensively in birds (particularly gulls and terns), and also in mammals including primates. It has been observed most often in antagonistic encounters, but these include courtship situations (in which tendencies to flee, to attack, and to approach are competing or alternating), as well as situations involving simple questions of submission and dominance. Visual avoidance is often interpreted as indicative of submission or appeasement to a dominant animal, although occasionally it is interpreted as defiance."*(page 45, Main and Weston, 1982)

Tinbergen and Moynihan (1952, cited in Main and Weston (1982)) hypothesise that the avoidance exhibited in these examples of animal behaviour can paradoxically serve proximity maintenance by performing the 'signal' function of concealment. For example, in situations such as courtship, black-headed gulls engage in 'head flagging' that is believed to minimise the exposure of their conspecific courtship partners to threatening elements of their own facial structure. That is, an animal acts avoidantly so that it removes from view stimuli that are threatening to others. This conception of avoidance is that avoidance is an evolved pattern in some animal species that works to maintain proximity. It is reworked by Main and Weston (1982), who apply it to the case of avoidance in human infancy. According to Main and Weston (1982) infants may avoid their carer's gaze because this kind of contact may drive the carer away. In addition, the evolutionary analysis set out above maintains that mild neglect is likely to be one end of a spectrum of mistreatment by the carer that an infant may trigger with strong signalling at the wrong time.

E.1.2 Maintaining proximity by stopping the infant itself moving away

Chance (1962, cited in Main and Weston (1982)) reinterprets the same data as Tinbergen, and additional data on male-male encounters in rats. His contrasting view is that the same movements have the same ultimate effect of increasing proximity, but the effect is brought about by 'cutting-off' processes in the avoidant animal. Chance hypothesises that it is the sight of another animal that arouses tendencies to flee or attack. Avoidance, in this view, is an external form of inhibition that an animal possesses to minimise its own proximity-decreasing tendencies. The animal doesn't internally inhibit the urge to flee, it externally moves so that its perceptual field does not give rise to the urge to flee.

Main and Weston (1982) adapt Chance's theory and suggest that infants may act avoidantly towards carers who have been rejecting and unresponsive because otherwise the infants would produce angry behaviour that might result in decreased proximity to the carer.

Tinbergen and Moynihan's, and Chance's theories show that avoidance can be an evolved pattern that is: *actually consonant with the working of a proximity-maintaining system.*

E.1.3 Evaluation of analysis of function - perhaps avoidance is not about proximity seeking

Main and Weston (1982) present a theory of the adaptive benefit of avoidant behaviour for infants in reunion episodes that does not centre on gaining safety. According to this theory, the benefit for the infants in acting avoidantly may be in maintaining self-control and behavioural organisation.

It may be that the avoidant behaviour helps shift the infant's attention so that it does not become distracted and lose attention due to its being in an angry emotional state. Therefore the infant maintains flexibility in behaviour. This view of avoidance is somewhat similar to the previously described view that infants possess mechanisms that lead them to act avoidantly because of the benefit in containing their own protest. However, it is different because in this version the benefit to the infant is not in keeping the carer close, it is in keeping control of its own behaviour for purposes other than maintaining safety.

This hypothesis might be tested by independently manipulating how angry and how insecure an infant would be expected to be in reunion episodes.

E.2 Why are forms of holding no longer correlated with forms of maternal behaviour?

Why should forms of holding lose their predictive power? Firstly, in developed countries with higher standards of living, less frequent extreme conditions such as famines, the availability of birth control and abortions, a legal system that strictly penalises infanticide and a culture that disproves of it, infanticide is simply much less frequent. The relatively high prevalence of carers that demonstrate high levels of unpleasant physical contact is not linked to high levels of serious abuse like infanticide because these types of serious abuse are very limited in modern societies. Secondly, it may be that unpleasant forms of physical contact were much lower in frequency in the EEA, so that examples of unpleasant contact were strongly against the norm and possessed greater significance. Hrdy (1999) reviews observations on what occurred when !Kung hunter gatherers became settled. The inter birth time became much smaller, mothers having less time for all infants. It may have been that in the EEA in times of plenty mothers would have experienced less stress. Therefore in good times carergiving contact would be pleasant and infanticide low. Only when conditions were unusually harsh would caregiving contact deteriorate and infanticide increase. This interbirth theory is contradicted by data on the Hadza people, which shows that some less mobile hunter-gatherer peoples do have small interbirth differences (Hrdy, 1999). However, modern hunter gather environments differ from the EEA because all the most productive and in histroric periods has been taken by agriculture and only the most marginal pieces have been left for hunter gatherers. It is not that agriculture guarantees a better or more reliable food source. Agriculturists only displaced hunter gatherers by force of numbers, not because individual agriculturalists had better lifestyles (Diamond, 1997). Before the displacement of hunter gatherers to marginal environments the stresses of hunter-gathering may have been very different to those observed in contemporary anthropological studies.

To summarise, it is likely that the EEA included periods when parental care was not only limited, but also that infant attempts to signal louder may have provoked either greater levels of negligence, abandonment or abuse. What this thesis is suggesting is that it is plausible that human infants possess mechanisms that provide the ability to sense when crying louder would be detrimental to the achievement of their aims.

Avoidance in human infancy has been interpreted as a behavioural strategy that is similar to displacement activities observed in non-human ethological studies of behaviour (Hinde, 1983). An example of displacement activity from animal behaviour might be found when an animal is faced with a con-specific with which it might fight with or flee from, but instead starts to groom itself (Bowlby, 1969). Displacement activities occur in animals when two strongly activated behaviours ‘cancel each other out’, and a seemingly inappropriate behaviour becomes active. Main and Weston (1982) propose that Avoidant behaviours that seem like displacement behaviours might increase infant safety in several ways:

- Keeping the infant safe by stopping the carer from rejecting the infant and moving away
- Keeping the infant safe by stopping the infant itself from moving away from the carer
- Keeping the infant safe by optimising its behavioural flexibility

These Main and Weston (1982) emphasise that these affects are not mutually exclusive and can be combined[1]. A key issue that this chapter is attempting to address is that if avoidance is a form of displacement behaviour in infants, is it implemented with mechanisms more similar to other animal species or through distinctly human mechanisms that perhaps include deliberation?

---

[1] Further details of the observations and theories which support these mechanisms are presented in appendix E.1

Appendix F

Further details of further work

F.1 Alternative information processing forms of anger

Section 4.4.5 noted that the representation of anger in the infant was limited and required further elaboration. An infant's angry protest with their carer might be interpreted as:

- **Anger as an attitude**. Adapting a definition of anger from Ortony *et al.* (1988), we might define infant anger as a cognitive state that arises when the infant attributes the results of the carers actions as detrimental to its prospects, in terms of the infant's current goals. According to this definition, for an infant to be angry with its carer the infant would be disapproving of an action of the carer that was assessed as blameworthy. At home this might be not responding sensitively and in the Strange Situation it might be leaving the infant alone. Anger would only result if the infant were displeased about the related event because it was undesirable, and because it blamed the carer for bringing this event about.

- **Anger as the activation of a rectifying motive**. Adapting a definition of anger from Sloman (1982), a definition of anger would mean that for an infant to be angry with their carer would involve the infant believing that the carer was responsible for actions which violated one of the infant's motives. In addition, for the infant to be angry, rather than regretful, the infant's belief must interact with a motive generator to produce a new motive that involves violating one of its carer's motives. This motive need not be acted upon, but it would need to be fairly intense desire to do something to the carer to be defined as anger.

- **Anger as the violation of an expectation**. Adapting a definition of anger from Rolls (1998), we might define infant anger as a state that arises when the infant expects to receive a response or security and doesn't. In this definition of anger, an organism can be angry without any knowledge of other agents.

- **Anger as innate reflex** The above definition of Roll's involves anger resulting from circumstances where an outcome important to a subject is different from that subjects expectation. A novel definition of an anger like state can be applied to reflexive systems that involve no change in expectation. Some reflexive reactions could be categorised as being examples of implicit anger because the evolutionary origins of those reactions are similar to the origins of non-innate responses, and the actual functions for the organism are the same.

The representation of anger in the architecture presented in this thesis might be adapted according to these information processing definitions of anger. An initial prototype anger mechanism has been implemented which is based upon the fourth interpretation of anger above. This mechanism possesses an innate measure of when a response is late, and when any given bid for attention has not been responded to by this time the infant acts by increasing the negative valency of it's signal. Adding protest to a infant agent's signal in this way brings carer's responses quicker if the protest takes the signal over the carer's action-threshold. This mechanism has not been systematically tested and validated due to time constraints. However, even these preliminary results suggest a hypothesis that angry protest acts to make unresponsive carers more like responsive ones. In empirical studies real carers very rarely left infants to cry until extinction of crying. Therefore, in reality, nearly all episodes of crying or protest, however long in duration, are terminated by some form of reward. This is duplicated in the simulation. Secure infants protest less because less intense signals already get a prompt response and increasing the intensity of these signals does not increase responsiveness. An extension of this mechanism is to allow the infants to learning, by associative mechanisms to modulate their emotional state (Arbib and Fellous, 2004). Then infant agents who repeatedly get their anger triggered will do so more in the future. Infant agents who get their anger triggered less do not learn this habit and consequently become less angry in future. The affect that quick responses deliver less demanding infants might be counter to the intuitions of many parents and is another prediction to be empirically tested.

F.2 Minsky's description of a form of culturally supported attachment learning

Section 4.4.6 suggested that cultural learning was a possible extension for the architectures described in this thesis. Minsky (1987) describes a form of culturally derived attachment learning:

> *"Suppose a child were playing in a certain way, and a stranger appeared and began to scold and criticise. The child would become frightened and disturbed and try to escape. But if, in the same situation, the child's parent arrived and proceeded to scold and criticise, the result would be different. Instead if being frightened the child would feel guilty and ashamed, and instead of trying to escape, the child would try to change what it was doing, in attempts to seek reassurance and approval*
>
> *I suspect that these two scenes engage different learning mechanisms. In the encounter with the forbidding visitor, the child might learn* **"I should not try to achieve my present goal in this kind of situation."** *But when scolded by someone to whom the child is "attached," the child might learn* **"I ought not to want to achieve that goal at all!"** *In the first case, it is a matter of learning which goal to pursue in which circumstance; in the second instance, it is more a question of what goals one should have. "* (page 175, Minsky, 1987)

F.3 A partially implemented evolutionary algorithm for the simulation

Section 4.4.6 suggested that the simulations described in this thesis might be adapted to form multi-generational evolutionary computation experiments.

The algorithm in figure F.1 has been partially implemented. It is intended that in evolutionary experiments the aspects of the infant architecture that are to be evolved will be represented in terms of a number of numerical parameters. So the generate procedure in figure F.1 would produce a population of infants expressed as vector of real numbers. The evaluate procedure would involve each individual in the population being run in a simulation intended to match what is known about the EEA. Therefore this environment will include a predator (which can be viewed as an abstraction of all types of threat) and things that need to be learnt. The vector

```
define evolve
    generate()
    repeat 10000 times
        evaluate()
        select()
        crossover()
        mutate()
    endrepeat
enddefine
```

Figure F.1: An evolutionary algorithm that might be used to validate the initial parameters within infant architectures

of real numbers for each infant would determine the value of critical parameters within the infants architecture. These might include: the initial values of the safe-range parameter; the intrinsic attraction of exploration and socialisation; the intrinsic level of object and agent wariness; and the parameters that control the rate at which goals and other internal states are increased and decreased. The fitness function for each infant in its selection phase would be a function that balanced learning with safety. The select procedure would involve keeping some proportion of the most successful vectors, that is keeping all the architectures that performed most successfully in the evaluation stage. Because the infants are represented as vectors there are a number of crossover methods available and mutation can be carried out according to fast evolutionary strategies (Yao and Liu, 1997).

Appendix G

Further details of scenarios and evaluation

G.1 Commentary on figures from Park Scenario simulation

Commentary on figure G.1(b)

The same image has been is shown magnified in subfigure G.3(a). Subfigure G.3(b) provides a complimentary table that gives information gained from the output window. This shows that at 19 steps after the initialisation of the simulation, although the BABY agent is moving away from the CARER agent its distance to the CARER (240) is still less than the

This level is arrived at as the activation is a function of two variables, the distance to the carer and the time that was spent not receiving gratification from exploration.

Commentary on cycle 1

The BABY is nearby the position that it started the simulation in, just by the side of the CARER. It is well within the Safe-range distance and therefore it gains a zero rating for its security goal. It can sense the presence of the object b4 as the closest object (it is in fact the only 'toy' object visible to the BABY at this time).

Commentary on cycle 9

The BABY is continuing towards object b4. The security goal hasn't risen, despite the BABY moving further away, but the exploration goal strength has risen, because this is a function of the distance to the object. Nearby objects that are unfamiliar provide a stronger goal than objects that are farther away.

Commentary on cycle 19

The baby is continuing towards object b4 and has learnt its first feature about object b4, there are 10 to learn in total. That objects have fixed numbers of features is an unrealistic assumption. This learning occurred because the BABY got closer than the limiting distance. The exploration subsystem has two modes, getting gratified and learning or being out of focus or too far or another goal being active.

Commentary on cycle 24

CARER is now just inside the edge of the Safe- range but the secure goal strength remains at zero. This is an unrealistic assumption. We wouldn't expect a sharp boundary such as this. Number of feature in memory increasing as is the explore goal strength

Commentary on cycle 28

The BABY is now across the Safe-range but the security goal is still zero. Memory and goal exploration still increasing

Commentary on cycle 30

Security goal has now started to rise. The rate of increase of the security goal is a major variable in the patterns of behaviour that the BABY shows, if the rate was high then the security goal would quickly overtake the exploration goal.

Commentary on cycle 35

BABY has now learnt all of the object b4's features and the goal activation resulting from this object drops to zero. Since security goal activation has been steadily growing this becomes active.

Commentary on cycle 46

As the BABY returns within the Safe-range the security goal diminishes to zero. The exploration subsystem then perceives object b5 as the next nearest object that is novel (ie hasn't been learnt about) so this becomes the target of exploration and the BABY moves towards it.

Commentary on cycle 56

Before the BABY reaches the object b5, and even before it learns about any of its features, the BABY is already over the Safe-range limit. Because b5 is further away the exploration goal is lower and hence a smaller security goal is needed to become active and start the BABY moving back towards the carer.

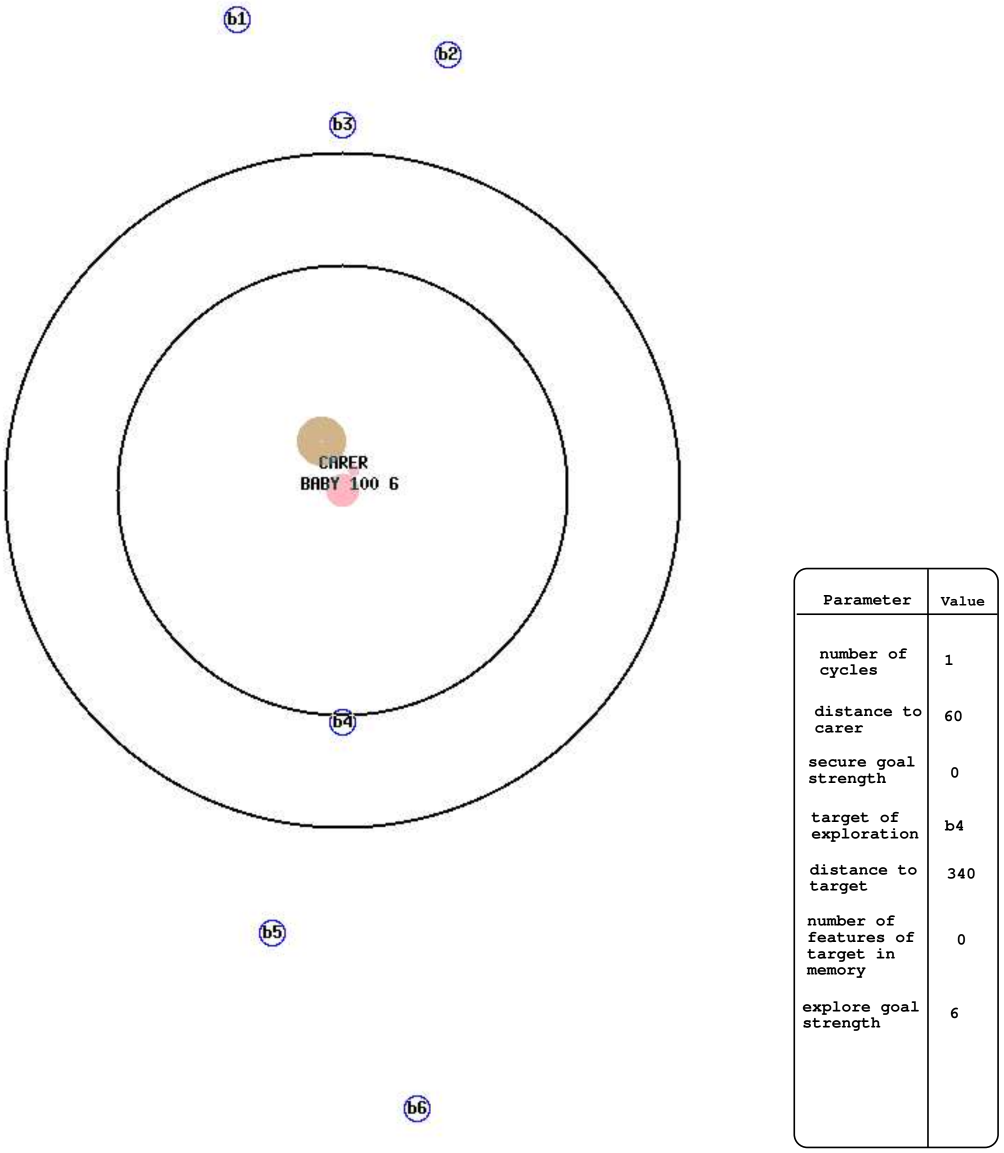

Parameter	Value
number of cycles	1
distance to carer	60
secure goal strength	0
target of exploration	b4
distance to target	340
number of features of target in memory	0
explore goal strength	6

(a) A screenshot of the 2D virtual environment

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.1: This shows the position of the simulation after 1 time cycle has elapsed. The BABY agent started the simulation close to the CARER. Since it is well within its Safe-range the BABY only has one goal with above zero activation. This is the explore goal and its strength of activation is set at 6.

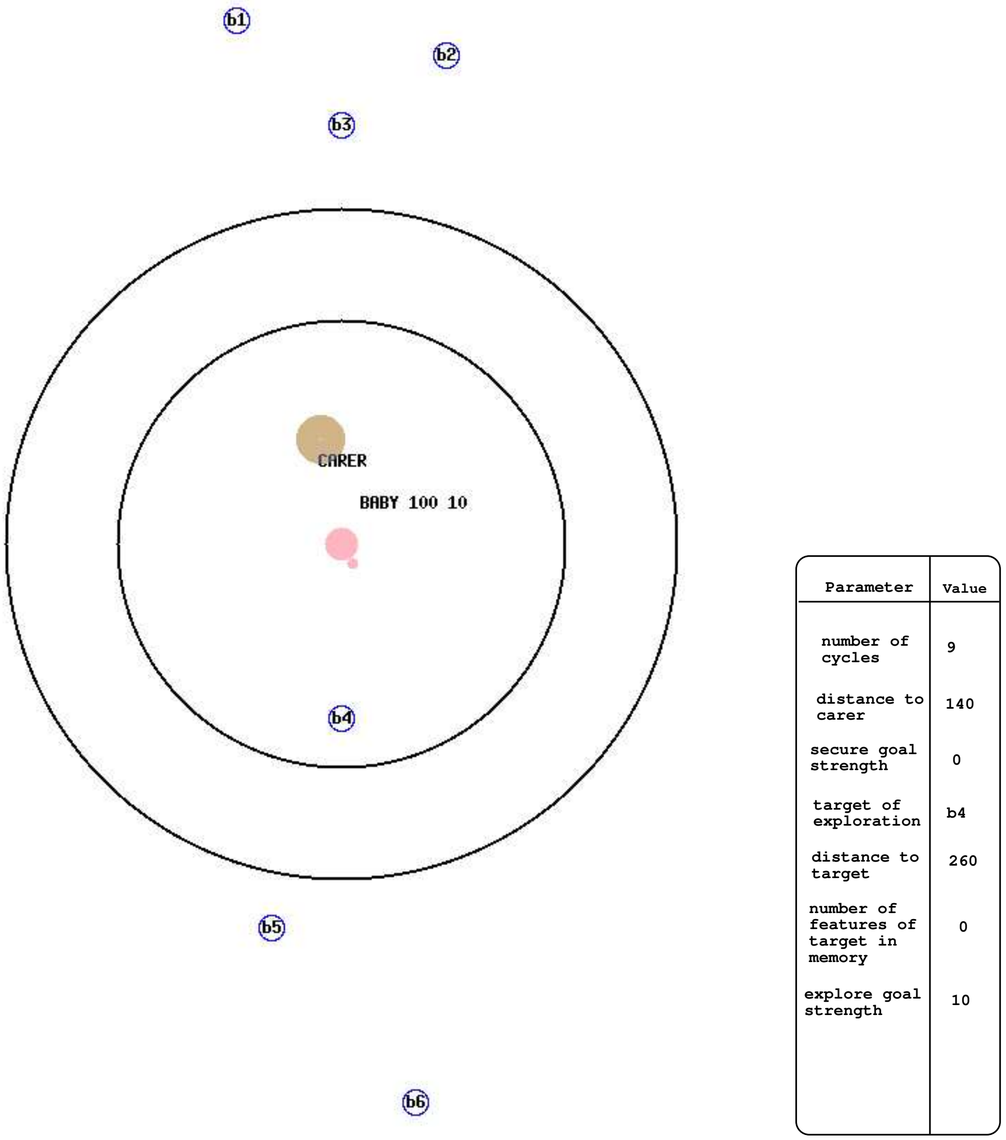

Parameter	Value
number of cycles	9
distance to carer	140
secure goal strength	0
target of exploration	b4
distance to target	260
number of features of target in memory	0
explore goal strength	10

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.2: This shows the position of the simulation after 9 time cycles have elapsed. The BABY is now further from the CARER but still within the Safe-range therefore the security goal is still set at zero activation. Since the BABY agent is closer to its target the exploration goal activation strength is no set to 10.

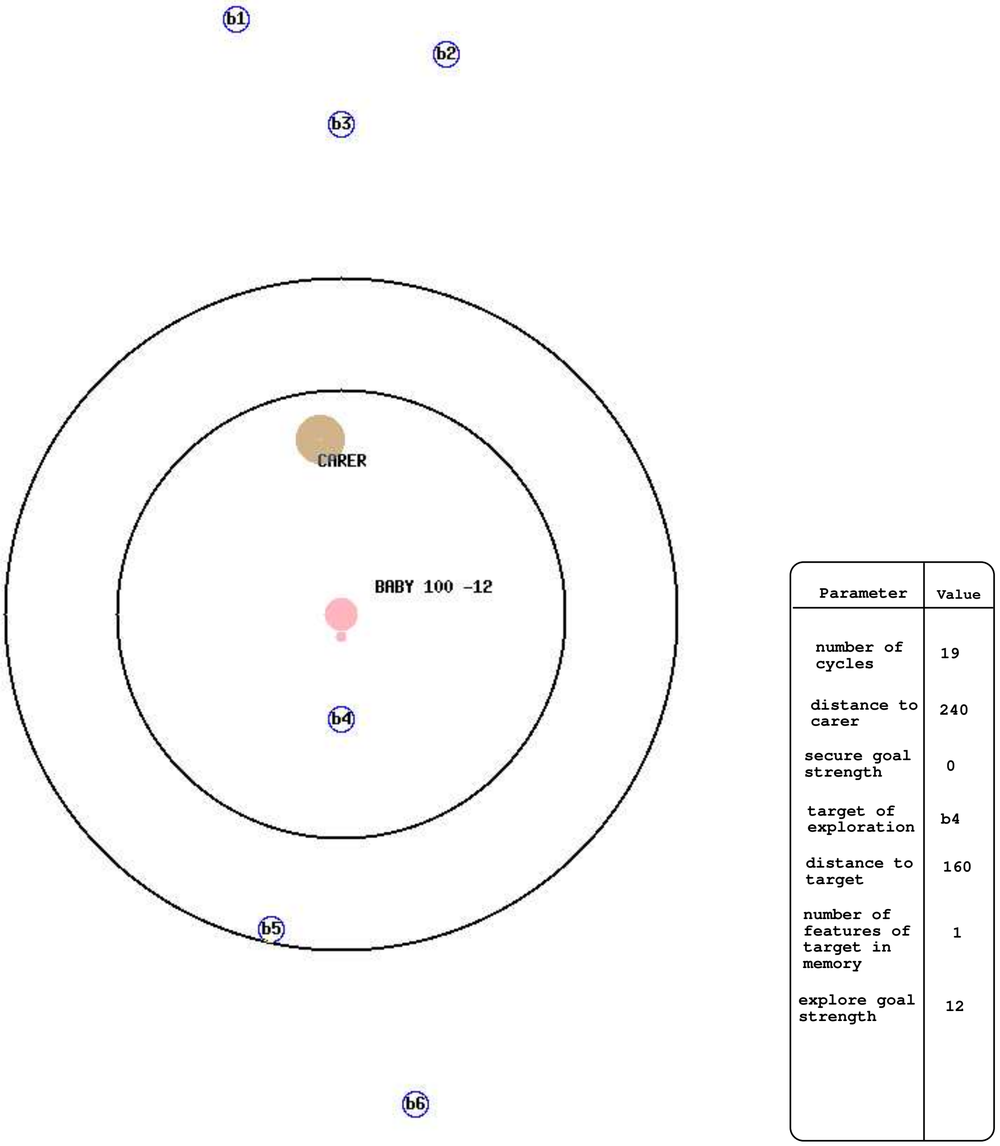

Parameter	Value
number of cycles	19
distance to carer	240
secure goal strength	0
target of exploration	b4
distance to target	160
number of features of target in memory	1
explore goal strength	12

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.3: The BABY at 19 cycles

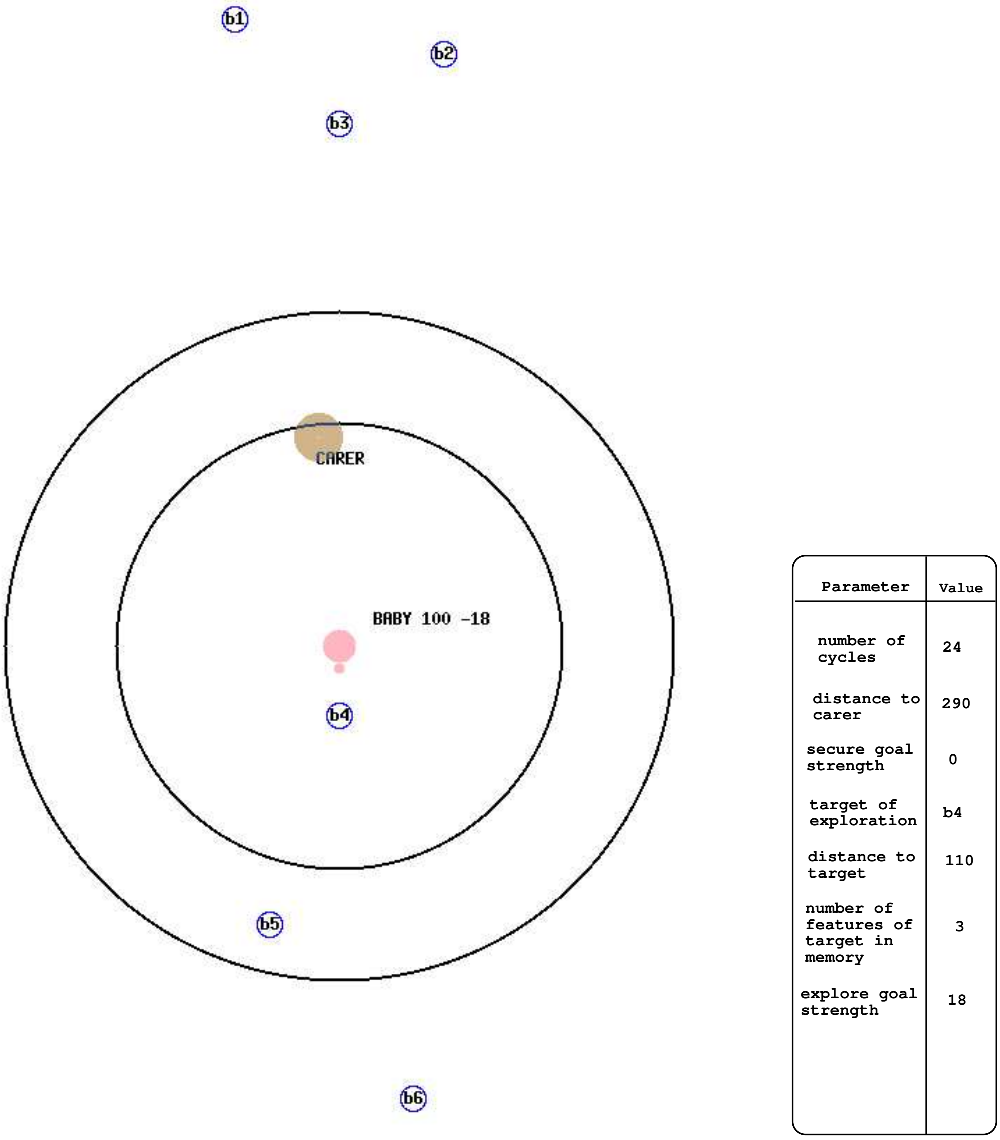

Parameter	Value
number of cycles	24
distance to carer	290
secure goal strength	0
target of exploration	b4
distance to target	110
number of features of target in memory	3
explore goal strength	18

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.4: The BABY at 24 cycles

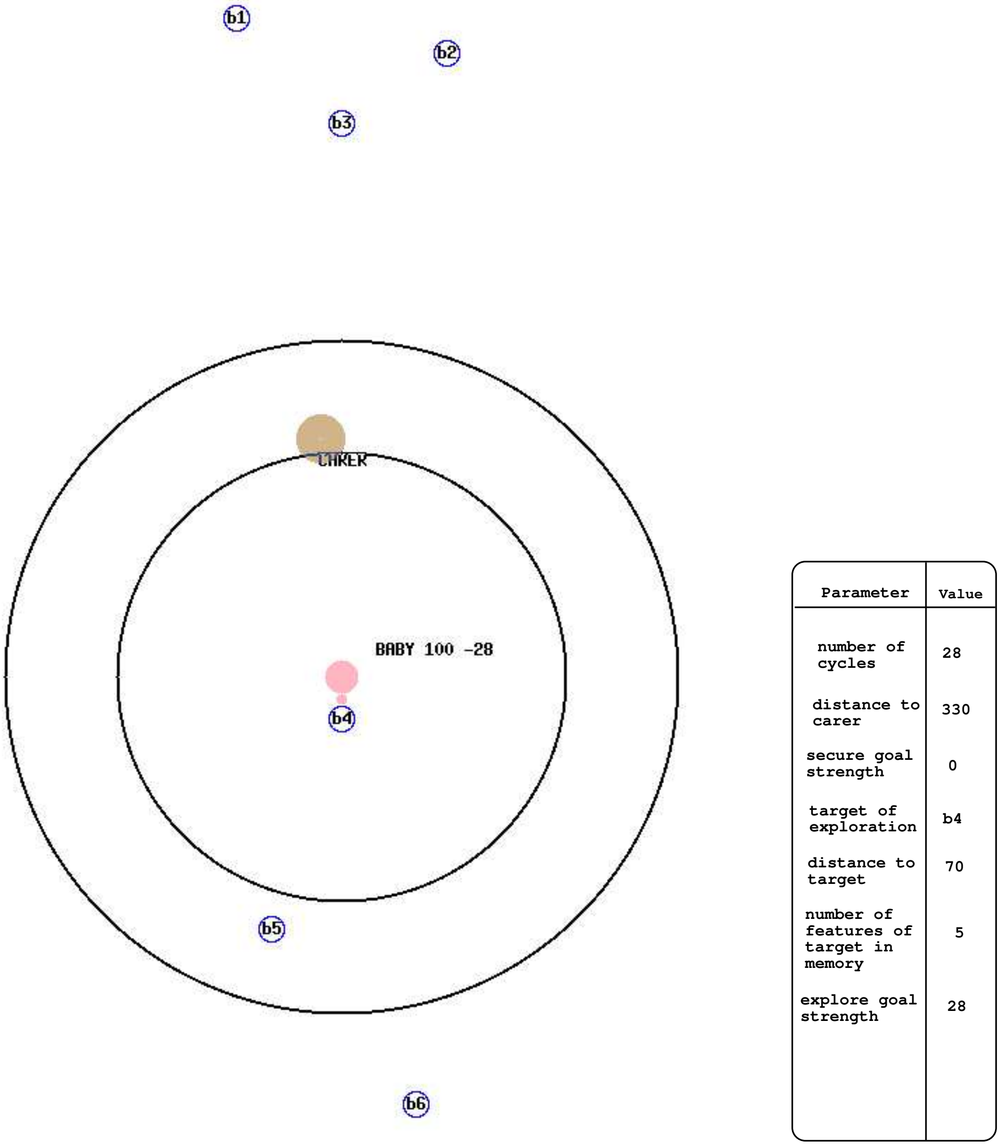

Parameter	Value
number of cycles	28
distance to carer	330
secure goal strength	0
target of exploration	b4
distance to target	70
number of features of target in memory	5
explore goal strength	28

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.5: The BABY at 28 cycles

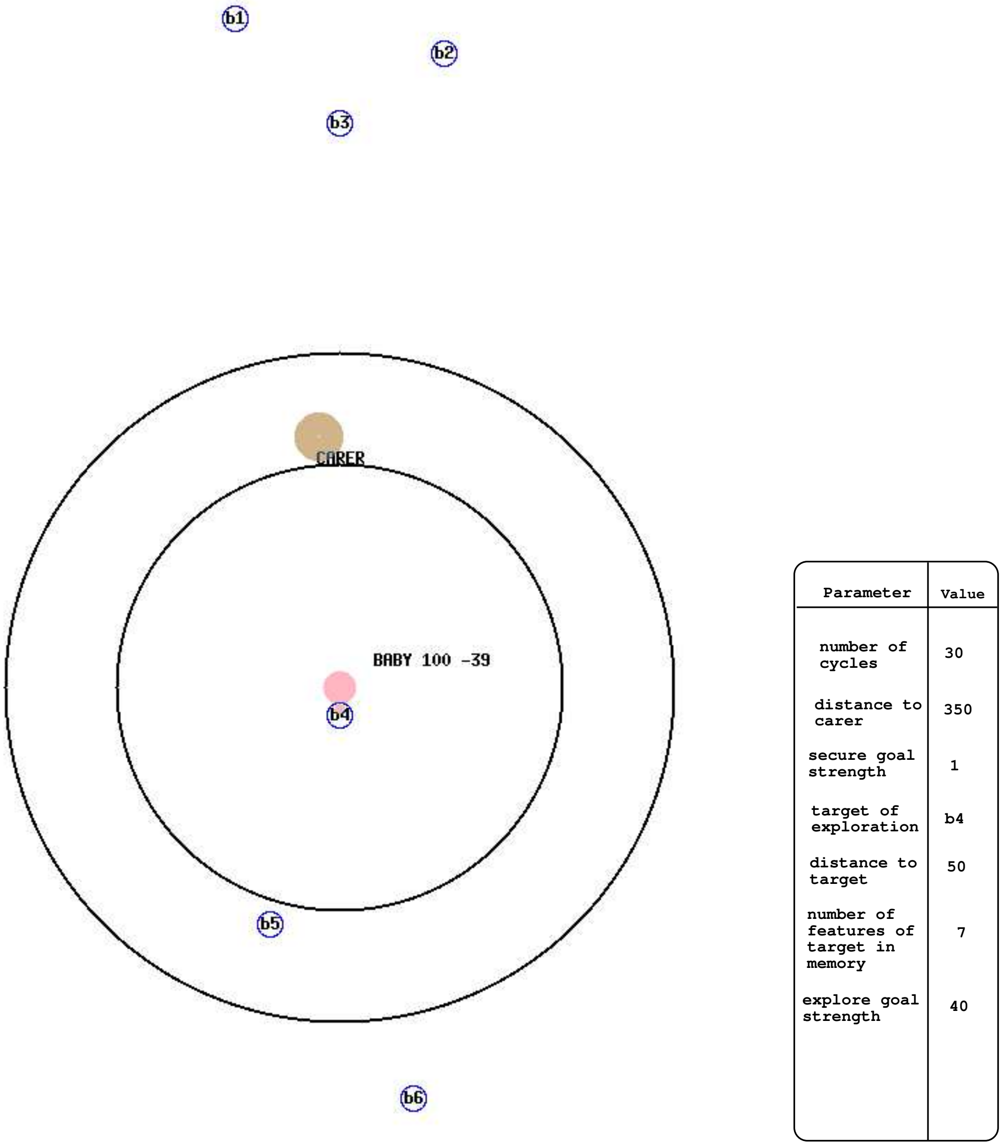

Parameter	Value
number of cycles	30
distance to carer	350
secure goal strength	1
target of exploration	b4
distance to target	50
number of features of target in memory	7
explore goal strength	40

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.6: The BABY at 30 cycles

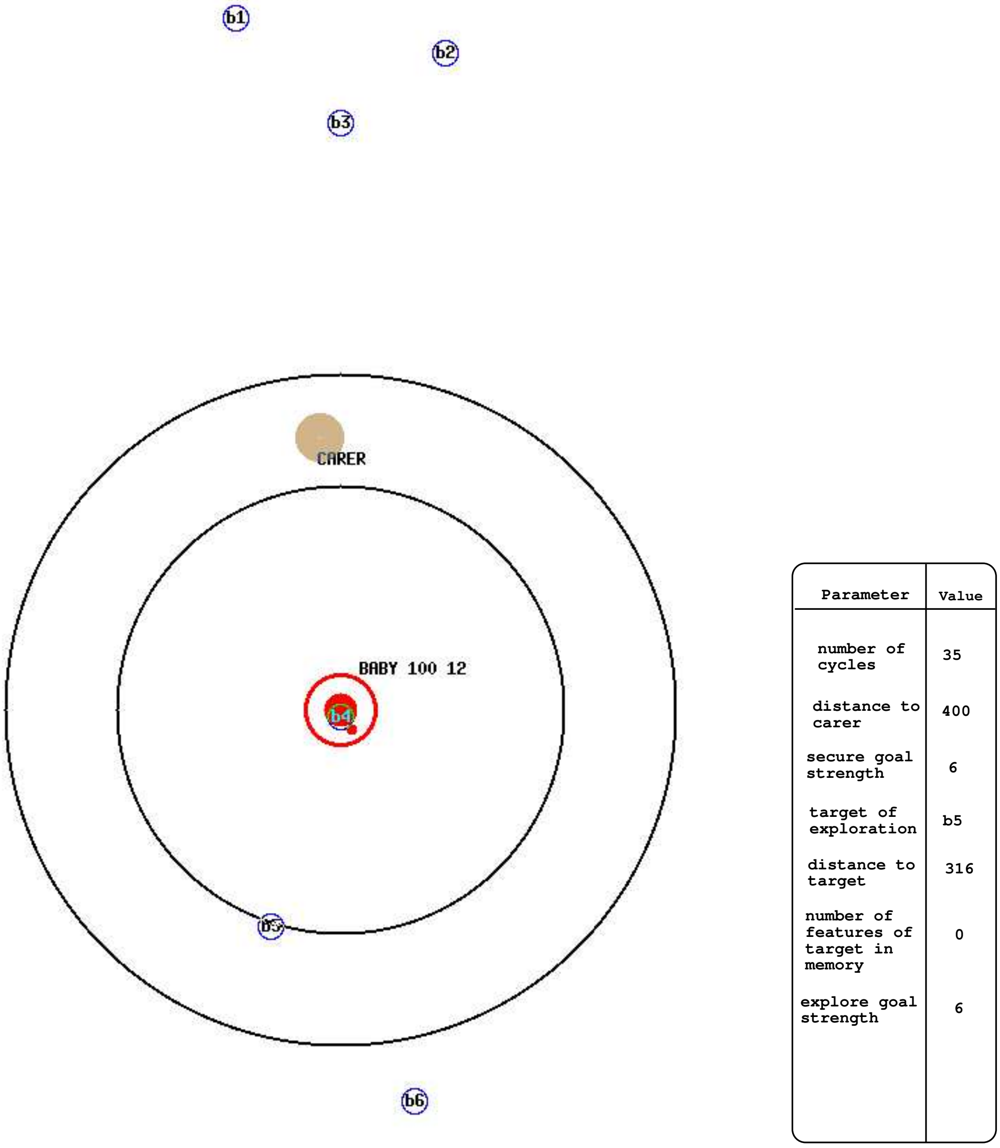

Parameter	Value
number of cycles	35
distance to carer	400
secure goal strength	6
target of exploration	b5
distance to target	316
number of features of target in memory	0
explore goal strength	6

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.7: The BABY at 35 cycles

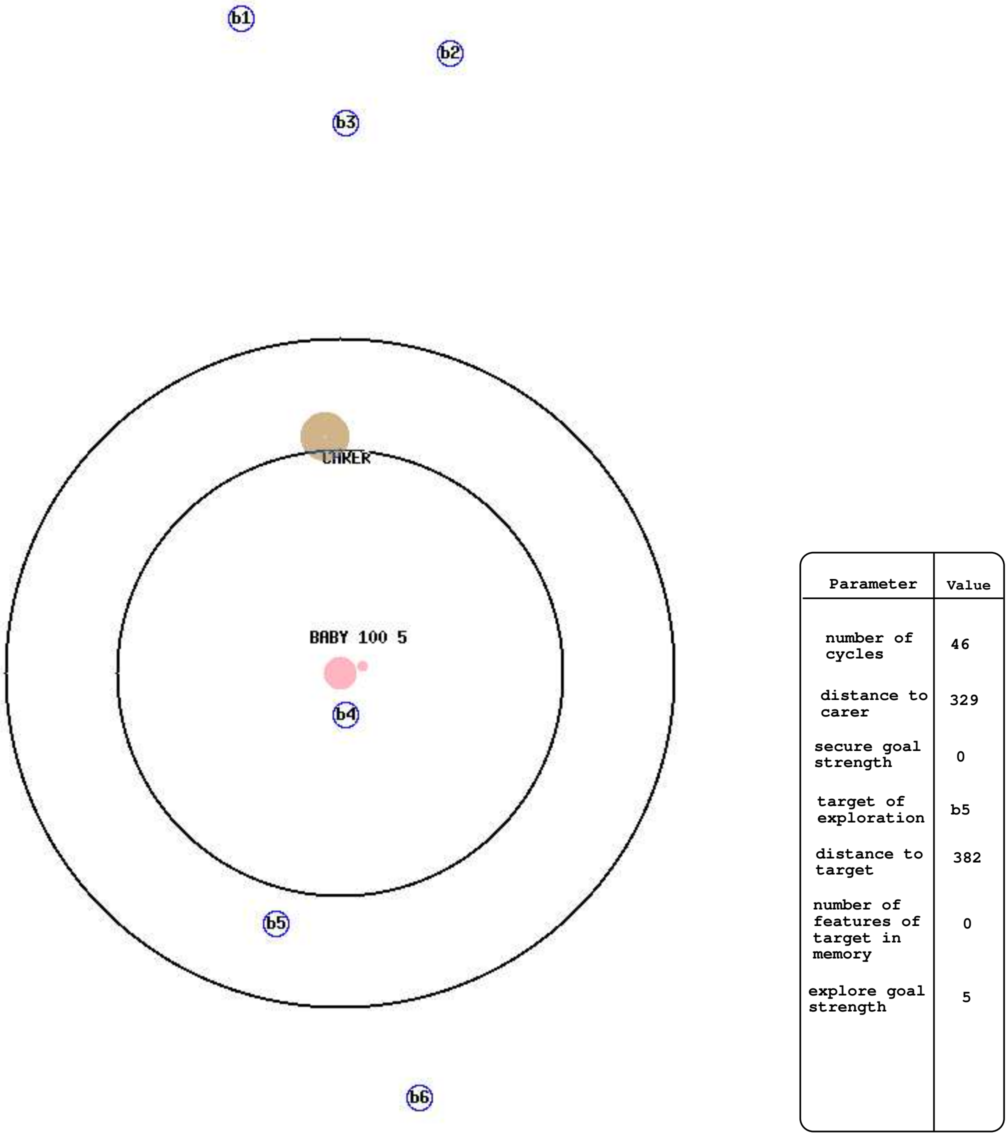

Parameter	Value
number of cycles	46
distance to carer	329
secure goal strength	0
target of exploration	b5
distance to target	382
number of features of target in memory	0
explore goal strength	5

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.8: The BABY at 46 cycles

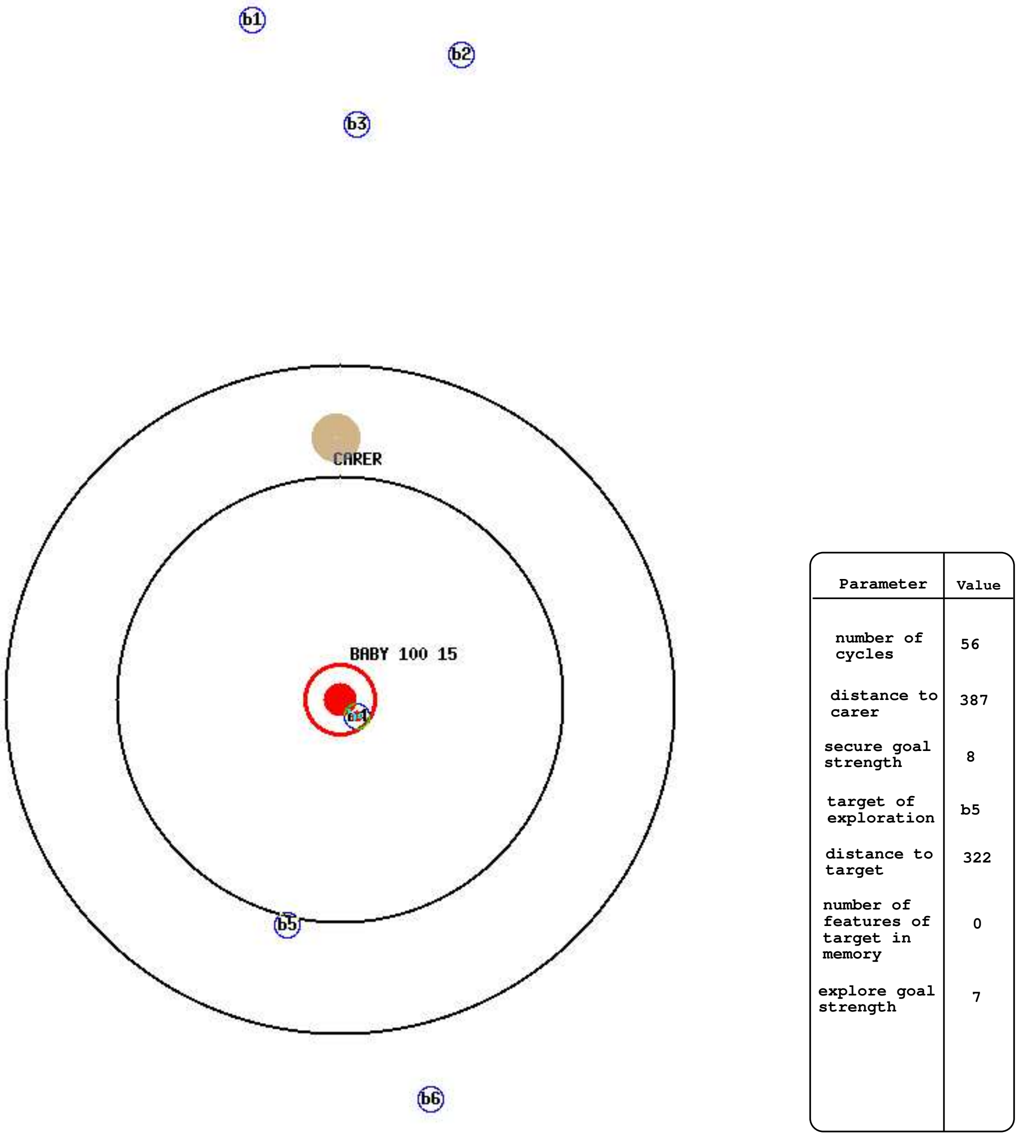

Parameter	Value
number of cycles	56
distance to carer	387
secure goal strength	8
target of exploration	b5
distance to target	322
number of features of target in memory	0
explore goal strength	7

(a) A screenshot

(b) Parameters that describe the BABY agent, including measures of its inner state

Figure G.9: The BABY at 56 cycles.

Bibliography

M. Ainsworth, M. Blehar, E. Waters, and S. Wall. *Patterns of Attachment: a psychological study of the strange situation*. Erlbaum, Hillsdale, NJ, 1978.

M. Ainsworth. *Infancy in Uganda: Infant Care and the Growth of Love.* Johns Hopkins Press, Baltimore:, 1967.

J.R. Anderson and C. Lebiere. *The atomic components of thought*. Lawrence Erlbaum Associates, Mahwah, NJ, 1999.

J.W. Anderson. Attachment behaviour out of doors. In N. Blurton-Jones, editor, *Ethological Studies of Child Behaviour*, pages 199–215. Cambridge University Press, Cambridge, UK, 1972.

J.R. Anderson. *Learning and memory*. Wiley, New York, 1995.

M.A. Arbib and J-M. Fellous. Emotions: from brain to robot. *Trends in cognitive sciences*, 8:554–61, December 2004.

M.A. Arbib. Schemas for the temporal organisation of behaviour. *Human neurobiology*, 4:63–72, 1985.

M. A. Arbib. From Rana Computatrix to Homo Loquens: A computational neuroethology of language evolution. In R. I. Damper et al., editors, *WGW'02 Biologically-inspired robotics: The legacy of W. Grey Walter*, pages 12–31, Bristol, 2002. Hewlett Packard Research Labs.

R. Baillargeon, M. Graber, J. Devos, and J. Black. Why do young infants fail to search for hidden objects? *Cognition*, 36:255–284, 1990.

J. Bates, A. B. Loyall, and W. S. Reilly. Broad agents. In *AAAI spring symposium on integrated intelligent architectures*. American Association for Artificial Intelligence, 1991. (Repr. in SIGART BULLETIN, 2(4), Aug. 1991, pp. 38–40).

L.P. Beaudoin. *Goal processing in autonomous agents.* PhD thesis, School of Computer Science, The University of Birmingham, 1994. (Available at http://www.cs.bham.ac.uk/ research/ cogaff/).

J. Belsky. Modern evolutionary theory and patterns of attachment. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 141–162. Guilford Press, London, 1999.

D. K. Belyaev, I. Plyusnina, and L. Trut. Domestication of the silver fox (vulpes fulvus desm): changes in physiological boundaries of the sensitive period of primary socialisation. *Applied Animal Behaviour Science*, 13:359–70, 1985.

J. Bowlby. *Attachment and loss: volume 1 attachment.* Basic books, New York, 1969. (Second edition 1982).

J. Bowlby. *A secure base.* Basic books, New York, 1988.

R.J. Brachman and H.J. Levesque, editors. *Readings in knowledge representation.* Morgan Kaufmann, Los Altos, California, 1985.

V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology.* The MIT Press, Cambridge, Mass, 1984.

C. Breazeal and B. Scassellati. Infant-like social interactions between a robot and a human caretaker. *Adaptive Behavior*, 8:49–74, 2000.

I. Bretherton and K.A. Munholland. Internal working models revisited. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 89–111. Guilford Press, London, 1999.

G.W. Bronson and W.B. Pankey. On the distinction between fear and wariness. *Child Development*, 48(4):1167–1183, 1977.

S. Bullock. An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models. In *Fourth European conference in Artificial Life*, pages 456–462. MIT Press, Cambridge. MA, 1997.

G. Butterworth. Joint visual attention in infancy. In Theories of Infant Development*, eds. G. Bremner & A. Slater*, pages 317–354. Blackwell Press, Oxford, 2004.

M. Carpenter, K. Nagell, and M. Tomasello. Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, 63, 1998.

J. Cassidy and L.J. Berlin. The insecure/ambivalent pattern of attachment: Theory and research. *Child Development*, 65(4):971–91, 1994.

J. Cassidy. The nature of the child's ties. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 3–20. Guilford Press, London, 1999.

J.S. Chisholm. *Death, hope and sex.* CUP, Cambridge, 1999.

A. Clark. *Being there: putting brain, body, and world together again.* MIT Press,, Cambridge, Mass, 1997.

R. Cooper and T. Shallice. Contention scheduling and the control of routine activities. *Cognitive Neuropsychology*, 17(4):297–338, 2000.

R. Cooper. *Modelling High-Level Cognitive Processes.* Lawrence Erlbaum Associates, New Jersey, 2002.

K. Craik. *The Nature of Explanation.* Cambridge University Press, London, New York, 1943.

P.M. Crittenden. Attachment and psychopathology. In Attachment theory: social, developmental and clinical perspectives*, eds. S. Goldberg R. Muir & J. Kerr*, pages 367–406. The Analytic Press, Hillsdale, NJ, 1995.

S.B. Crockenberg. Infant irritability, mother responsiveness, and social support influences on the security of infant-mother attachment. *Child Development*, 52:857–69, 1981.

H. Cruse. The evolution of cognition - a hypothesis. *Cognitive Science*, 27:135–155, 2003.

M. Daly and M.I. Wilson. *Homicide.* Aldine, New York, 1988.

L. Darden and N. Maull. Interfield theories. *Philosophy of Science*, 44:43–64, 1977. 1.

M.R.W. Dawson. *Minds and Machines: Connectionism and Psychological Modelling.* Blackwell Publishing, Oxford, 2004.

D. Dennett. Cognitive science as reverse engineering: Several meanings of 'top-down' and'bottom-up',. In Logic, Methodology and Philosophy of Science IX*, eds. D. Prawitz, B. Skyrms & D. Westersthl*, pages 679–689. Elsevier Science, BV, Amsterdam, North-Holland, 1994.

D.C. Dennett. *Darwin's Dangerous Idea: Evolution and the Meanings of Life.* Penguin Press, London and New York, 1995.

D.C. Dennett. *Kinds of minds: towards an understanding of consciousness*. Weidenfeld and Nicholson, London, 1996.

A. Diamond, L. Cruttenden, and D. Neiderman. A not b with multiple wells: 1. why are multiple wells sometimes easier than two wells? 2. memory or memory + inhibition? *Developmental Psychology*, 30(4):192–205, 1994.

A. Diamond. Development of the ability to use recall to guide action, as indicated by infants' performance on a notb task. *Child Development*, 56(4):868–883, 1985.

A. Diamond. Differences between adult and infant cognition: Is the crucial variable presence or absence of language? In Thought without language*, ed. L. Weiskrantz*, pages 335–370. Oxford Science Publications, Oxford, 1988.

A. Diamond. Frontal lobe involvement in cognitive changes during the first year of life. In Brain Maturation and cognitive development: Comparative and cross cultural perspectives*, eds. K. R. Gibson & A.C. Petersen*, pages 127–180. Aldine de Gruyter, Hawthorne, NY, 1991.

A. Diamond. Neuropsychological insights into the meaning of object concept development. In The epigenisis of minds: Essays on biology and cognition*, eds. S. Carey & R. Gelman*, pages 67–110. Erlbaum, Hillsdale, NJ, 1991.

J. Diamond. *Guns, Germs and Steel*. Jonathan Cape, 1997.

A. Diamond. Close interrelation of motor development and cognitive development and of the cerebellum and prefrontal cortex. *Child Development*, 71(1):44–56, 2000. (special issue: New directions for child development in the 21st century).

A. Diamond. Normal development of prefrontal cortex from birth to young adulthood: Cognitive functions, anatomy, and biochemistry. In Principles of Frontal Lobe Function*, eds. D.T. Stuss & R.T. Knight*, pages 466–503. Oxford University Press, London, UK, 2002.

Sandra Pinto Clara do Carmo Gadhano. *Reinforcement Learning in Autonomous Robots: An Empirical Investigation of the Role of Emotions*. PhD thesis, The University of Edinburgh, 1999.

J. Doran and M. Palmer. The eos project: integrating two models of palaeolithic social change. In Artificial Societies*, eds. N. Gilbert & R. Conte*, pages 103–125. Guilford Press, London, 1995.

K. Durkin. *Developmental social psychology*. Blackwell, 1995.

J. Elman, E.A. Bates, M.H. Johnson, A. Karmiloff-Smith, D. Parisi, and K. Plunkett. *Rethinking innateness : a connectionist perspective on development*. MIT Press, London, 1996.

J.H. Flavell, P.H. Miller, and S.A. Miller. *Cognitive Development*. Pearson Education, Upper Saddle River NJ, 2002.

N. A. Fox and J.A. Card. Psychophysiological measures in the study of attachment. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 226–248. Guilford Press, London, 1999.

R.C. Fraley and S.J. Spieker. Are infant attachment patterns continuously or categorically distributed? a taxometric analysis of strange situation behaviour. *Developmental Psychology*, 39:387–404, 2003.

M.J. Frank. When and when not to use your subthalamic nucleus: Lessons from a computational model of the basal ganglia. In J.J. Bryson, T.J. Prescott, and A.K. Seth, editors, *Proceedings of Modelling Natural Action Selection*, pages 53–60. AISB Press, School of Science and Technology, University of Sussex, Brighton, 2005.

B. Girard, V. Cuzin, A. Guillot, K.N. Gurney, and T.J. Prescott. Comparing a brain-inspired robot action selection mechanism with winner-takes-all. In *Proceedings of the seventh international conference on simulation of adaptive behavior on From animals to animats*, pages 75–84. MIT Press Cambridge, MA, USA, 2002.

J. Gleick, editor. *Chaos: Making a New Science*. Vintage, London, 1996.

S. Goldberg. *Attachment and Development*. Arnold, London, 2000.

R. Grush. The emulation theory of representation: motor control, imagery, and perception. *The Behavioral and Brain Sciences*, 27(3):377–396, 2004. (To appear).

K.N. Gurney, T.J. Prescott, and P. Redgrave. A computational model of action selection in the basal ganglia I. A new functional anatomy. *Biological Cybernetics*, 84:401–410, 2001.

K.N. Gurney, T.J. Prescott, and P. Redgrave. A computational model of action selection in the basal ganglia II. Analysis and simulation of behaviour . *Biological Cybernetics*, 84:411–423, 2001.

D. Haig. Genetics conflicts in human pregnancy. *Quarterly Review of Biology*, 86:495–531, 1993.

‘H. Harlow. The Nature of Love. *American Psychologist*, 13:573–685, 1958.

J. Haugeland. What is mind design? In J. Haugeland, editor, *Mind Design II*. MIT Press, Cambridge, MA, 1997.

M.D. Hauser and E.S. Spelke. Evolutionary and developmental foundations of human knowledge: A case study of mathematics. In M. Gazzaniga, editor, *The Cognitive Neurosciences*. MIT Press, Cambridge, MA, 2004.

L. Hertsgaard, M. Gunnar, M. F. Erickson, and M. Nachmias. Adrenocortical responses to the strange situation in infants with disorganised/disoriented attachment relationships. *Child Development*, 66:1100–1106, 1995. 4.

E. Hesse. The adult attachment interview, historical and current perspectives. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 395–433. Guilford Press, London, 1999.

R.A. Hinde. Ethology and child development. In *Hand book of child psychology, eds. P.H. Mussen*, pages 27–93. J. Wiley and Sons, New York, 1983.

M.A. Hofer. Hidden regulators in attachment and loss. In Attachment theory: social, developmental and clinical perspectives*, eds. S. Goldberg R. Muir & J. Kerr*, pages 203–230. The Analytic Press, Hillsdale, NJ, 1995.

J. Holmes. *John Bowlby and Attachment Theory*. Routledge, 1993. (revised edition).

I.R. Hoos. Methodology, methods and models. In F.E. Emery, editor, *Systems thinking*, pages 42–62. Penguin, Brighton, 1981.

J.C. Houk, D. Fraser, A. Fishbach, S.A. Roy, L.K. Simo, C. Bastianen, D. Fansler-Wald, L.E. Miller, P.J. Reber, and M. Botivinick. Action selection in subcortical loops through the basal ganglia. In J.J. Bryson, T.J. Prescott, and A.K. Seth, editors, *Proceedings of Modelling Natural Action Selection*, pages 84–92. AISB Press, School of Science and Technology, University of Sussex, Brighton, 2005.

S.B. Hrdy. *Mother Nature: A History of Mothers, Infants, and Natural Selection*. Pantheon Books, London, 1999.

J. Kagan. *Three Seductive Ideas*. Harvard University Press, Cambridge: MA, 2000.

E.R. Kandel, J.H Schwartz, and T.M Jessell. *Principles of Neural Science*. Mcgraw Hill, London, 2000. Fourth Edition.

A. Karmiloff-Smith. *Beyond modularity : a developmental perspective on cognitive science*. MIT Press, 1994.

A. Kerr and P.D. Zelazo. Development of 'hot' executive function: The children's gambling task. *Brain and Cognition*, ?(?):?, 2004. In Press, Corrected Proof from Science Direct.

R.R. Kobak, H.E. Cole, R. Ferentz-Gillies, W. S. Fleming, and W. Gamble. Attachment and emotion regulation during mother-teen problem solving: A control theory analysis. *Child Development*, 64:231–245, 1993.

G.W. Kraemer. A psychobiological theory of attachment. *Behavioural and Brain Sciences*, 15:493–541, 1992.

I. Lakatos. *Criticism and the Growth of Knowledge*. Cambridge University Press, New York, 1970.

K.N. Laland and G.R. Brown. *Sense and Nonsense*. Oxford University Press, Oxford, 2002.

J.E. LeDoux. *The Emotional Brain*. Simon & Schuster, New York, 1996.

M. Likhachev and R.C. Arkin. Robotic comfort zones. In *Proceedings of SPIE: Sensor Fusion and Decentralized Control in Robotic Systems*, pages 27–41, 2000.

B. Logan. Biorealistic simulation of baboon foraging using agent-based modelling. In J.J. Bryson, T.J. Prescott, and A.K. Seth, editors, *Proceedings of Modelling Natural Action Selection*, pages 127–134. AISB Press, School of Science and Technology, University of Sussex, Brighton, 2005.

M. Macdonald. Warmth as a developmental construct: an evolutionary analysis. *Child Development*, 63:753–773, 1992. 2.

M. Main and D.R. Weston. Avoidance of the attachment figure in infancy. In The place of attachment in human behavior*, eds. M. Parkes & J. Stevenson-Hinde*, pages 31–59. Basic Books, New York, 1982.

M. Main, N. Kaplan, and J. Cassidy. Security in infancy, childhood and adulthood: A move to the level if representation. In Growing points of attachment theory and research*, eds. I. Bretherton and E. Waters*, pages 66–107. Monographs of the society for research in child development, 50(1-2 serial no 209), London, 1985.

M. Main. Cross cultural studies of attachment organisation; recent studies, changing methodologies, and the concept of conditional strategies. *Human Development*, 33:48–61, 1990.

M. Main. Metacognitive knowledge, metacognitive monitoring, and singular (coherent) vs multiple (incoherent) model of attachment: findings and directions. In Attachment across the lifecycle*, eds. C.M. Parkes & J. Stevenson-Hinde & P. Marris*, pages 127–159. Routledge, London, 1991.

R.S. Marvin and P.A. Britner. Normative development: The ontogeny of attachment. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 44–67. Guilford Press, London, 1999.

P. Meehl. Clarifications about taxometric method. *Applied and Preventative Psychology*, 8:165–174, 1999.

E. Meins. *Security of attachment and the social development of cognition*. Psychology Press, Hove, 1997.

M. L. Minsky. *The Society of Mind*. William Heinemann Ltd., London, 1987.

F. Montes-Gonzalez, T.J. Prescott, K. Gurney, M. Humphries, and P. Redgrave. An embedded model of action selection mechanisms in the vertebrate brain., 2001. Technical Report, Department of Psychology, University of Sheffield, Western Bank, Sheffield S10 2TP, UK.

Y. Munakata, D. Bauer, T. Stackhouse, L. Landgraf, and J. Huddleston. Rich interpretation vs. deflationary accounts in cognitive development: the case of means-end skills in 7-month-old infants. *Cognition*, 83:B43–B53, 2002.

Y. Munakata. Infant perseveration and implications for object permanence theories: A pdp model of the a-not-b task. *Developmental Science*, 1(2):161–184, 1998.

D. Nau, M. Ghallab, and P. Traverso. *Automated Planning: Theory and Practice*. Morgan Kaufmann, New York, 2004.

A. Newell. Reflections on the structure of an interdiscipline. In The Study of Information: Interdisciplinary Perspectives*, eds. F. Machlup and U. Mansfield*, pages 99–109. Wiley, New York, 1983.

A. Newell. *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA, 1990.

A. Ortony, G.L. Clore, and A. Collins. *The Cognitive Structure of the Emotions*. Cambridge University Press, New York, 1988.

B. Osthaus, S.E.G. Lea, and A.M. Slater. Dogs (canis lupus familiaris) fail to show understanding of means-end connections in a string-pulling task. *Animal Cognition*, 1:37–47, 2005. 1.

J.M. Pearce. *Animal Learning and Cognition*. Psychology Press, Hove, 1997.

D. Petters. Simulating infant-carer relationship dynamics. In *Proc AAAI Spring Symposium 2004: Architectures for Modeling Emotion - Cross-Disciplinary Foundations*, number SS-04-02 in AAAI Technical reports, pages 114–122, Menlo Park, CA, 2004.

J. Piaget. *The origin of intelligence in the child*. Routledge and Kegan Paul, London, England, 1953.

J. Piaget. *The construction of reality in the child*. Routledge and Kegan Paul, London, England, 1955.

S. Pinker. *How the Mind Works*. Penguin Books, London, 1998.

H. J. Polan and M.A. Hofer. Psychobiological origins of infant attachment and separation responses. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 162–180. Guilford Press, London, 1999.

K.R. Popper. *Objective Knowledge, An evolutionary approach*. OUP, Oxford, 1972.

K.R. Popper. *Myth of the Framework*. Routledge, New York, 1997.

T.J. Prescott, P. Redgrave, and K. Gurney. Layered control architectures in robots and vertebrates. *Adaptive Behaviour*, 1999.

T.J. Prescott, K. Gurney, and P. Redgrave. The basal ganglia. In M. Arbib, editor, *The Handbook of Brain Theory and Neural Networks*, pages 147–151. MIT Press, Cambridge MA, 2002.

P. Redgrave, T.J. Prescott, and K. Gurney. The basal ganglia: a vertebrate solution to the selection problem. *Neuroscience*, 89:1009–1023, 1999. 4.

M. Ridley. *The origins of virtue*. Penguin, London, 1996.

P. Rochat. *The Infant's World*. Harvard University Press, Cambridge, MA, 2004.

E.T. Rolls. *The Brain and Emotion*. Oxford University Press, Oxford, 1998.

R.I. Rumiati, S. Zanini, L. Vorano, and T. Shallice. A form of ideational apraxia as a delective deficit of contention scheduling. *Cognitive Neuropsychology*, 18:617–642, 2001. 7.

J. Russell. *Agency: Its role in mental development*. Erlbaum, Hove, 1996.

L. R. Santos, A. Rosati, C. Sproul, B. Spaulding, and M.D. Hauser. Means-means-end tool choice in cotton-top tamarins (saguinus oedipus): finding the limits on primates knowledge of tools. *Animal Cognition*, 8:236–246, 2005. 3.

R.C. Schank and R. Abelson. *Scripts, plans and understanding*. Lawrence Erlbaum Associates, Hove, UK, 1977.

R.C. Schank. *Dynamic memory: a theory of reminding and learning in computers and people*. Cambridge University Press, Cambridge, UK, 1982.

M. Schlesinger. The agent-based approach: A new direction for computational models of development. *Developmental Review*, 21:121–146, 2001.

A.N. Schore. *Affect regulation and the origin of the self : the neurobiology of emotional Development.* Erlbaum, 1994.

M. Shanahan. A cognitive architecture that combines internal simulation with a global workspace. *Consciousness and Cognition*, 15:433–449, 2006.

H. A. Simon. *The Sciences of the Artificial.* MIT Press, Cambridge, Mass, 1969. (Third edition 1996).

J.A. Simpson. Attachment theory in modern evolutionary perspective. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 115–141. Guilford Press, London, 1999.

A. Sloman and R. L. Chrisley. More things than are dreamt of in your biology: Information-processing in biologically-inspired robots. *Cognitive Systems Research*, 6(2):145–174, June 2005.

A. Sloman. Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence. In *Proc 2nd IJCAI*, London, 1971. Reprinted in *Artificial Intelligence*, pp 209-225, 1971, and in J.M. Nicholas, ed. *Images, Perception, and Knowledge*. Dordrecht-Holland: Reidel. 1977.

A. Sloman. Afterthoughts on analogical representation. In R. Schank and B. Nash-Webber, editors, *Theoretical Issues in Natural Language Processing (TINLAP)*, pages 431–439, MIT, 1975. Reprinted in Brachman and Levesque (1985).

A. Sloman. Towards a grammar of emotions. *New Universities Quarterly*, 36(3):230–238, 1982.

A. Sloman. What sort of control system is able to have a personality?, 1995. (Presented at Workshop on Designing personalities for synthetic actors, Vienna, June 1995).

A. Sloman. Interacting trajectories in design space and niche space: A philosopher speculates about evolution. In M.Schoenauer et al., editor, *Parallel Problem Solving from Nature – PPSN VI*, Lecture Notes in Computer Science, No 1917, pages 3–16, Berlin, 2000. Springer-Verlag.

A. Sloman. Beyond Shallow Models of Emotion. *Cognitive Processing: International Quarterly of Cognitive Science*, 2(1):177–198, 2001.

A. Sloman. How many separately evolved emotional beasties live within us? In R. Trappl, P. Petta, and S. Payr, editors, *Emotions in Humans and Artifacts*, pages 29–96. MIT Press, Cambridge, MA, 2002.

M.S. Smart and R.C. Smart. *Children: Development and Relationships*. Macmillan, New York, 1967.

E.A. Smith. Inuit foraging groups: some simple models incorporating conflicts of interest, relatedness, and central place sharing. *Ethology and Sociobiology*, 6:27–47, 1985. 1.

J. Soltis. The signal functions of infant crying. *Behavioural and Brain Sciences*, 27:443–490, 2004. 4.

G. Spangler and K.E. Grossman. Biobehavioural organisation in securely and insecurely attached infants. *Child Development*, 64:1439–1450, 1993. 5.

‘E.S. Spelke. Core knowledge. *American Psychologist*, 55:1233–1243, 2000.

L.A. Sroufe and E. Waters. Attachment as an organisational construct. *Child Development*, 48(4):1184–99, 1977.

L.A. Sroufe and E. Waters. Heart rate as a convergent measure in clinical and developmental research. *Merryl-Palmer Quartely*, 23:1–14, 1977.

L.A. Sroufe, B. Egeland, and E.A. Carlson. One social world: The integrated development of parent-child and peer relationships. In Relationships as developmental context: The 30th Minnesota symposium on child psychology., *eds. W. A. Collins & B. Laursen*, pages 241–262. Erlbaum, Hillsdale, NJ, 1999.

S. J. Suomi. Attachment in rhesus monkeys. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 181–197. Guilford Press, London, 1999.

R.S. Sutton and A.G. Barto. *Reinforcement learning: an introduction*. The MIT Press, 1998.

D. Symons. Adaptiveness and adaptation. *Ethology and Sociobiology*, 11:427–444, 1990. 3.

R.A. Thompson. Sensitivity and security: New questions to ponder. *Child Development*, 68(4):595–597, 1997.

N. Tinbergen. On aims and methods of ethology. *Z Tierpsychol t*, 20:410–433, 1963.

M. Tomasello. Social cognition before the revolution. In Early Social Cognition*, eds. P. Rochat*, pages 301–314. Lawrence Erlbaum, Mahwah, N.J.; London, 1999.

M. Tomasello. *The cultural origins of human cognition*. Harvard University Press,, Cambridge, Mass.; London, 1999.

J. Tooby and L. Cosmides. The past explains the present: emotional adaptations and the structure of ancestral environments. *Ethology and Sociobiology*, 11:375–424, 1990. 3.

R. Trivers. Parent-offspring conflict. *American Zoologist*, 14:249–264, 1974. 3.

T. Tyrrell. *Computational mechanisms for action selection*. PhD thesis, University of Edinburgh, Edinburgh, UK., 1993.

M.H. van Ijzendoorn and M.J. Bakermans-Kranenburg. Maternal sensitivity and infant temperament. In Theories of Infant Development*, eds. G. Bremner & A. Slater*, pages 233–257. Blackwell Press, Oxford, 2004.

M.H. van Ijzendoorn and P.M. Kroonenberg. Cross-cultural patterns of attachment: A meta-analysis of the strange situation. *Child Development*, 59(1):147–156, 1988.

M.H. van Ijzendoorn and A. Sagi. Cross-cultural patterns of attachment: Universal and contextual dimensions. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 713–734. Guilford Press, London, 1999.

M.H. van IJzendoorn, C.M.J.L. Vereijken, M.J. Bakermans-Kranenburg, and J. M. Riksen-Walraven. Assessing attachment security with the attachment q sort: Meta-analytic evidence for the validity of the observer aqs. *Child Development*, 75:1188–1213, 2004. 4.

E. Waters and T.P. Beauchaine. Are there really patterns of attachment? theoretical and empirical perspectives. *Developmental Psychology*, 39(3):417–422, 2003.

E. Waters and K.E. Deane. Defining and assessing individual differences in attachment relationships: Q-methodology and the organisation of behaviour in infancy and early childhood. *Monographs of the Society for Research in Child Development*, 50(1-2):41–65, 1985.

B. Webb. Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24:1033–1050, 2001.

N.S. Weinfield, L.A. Sroufe, B. Egeland, and E.A. Carlson. The nature of individual differences in infant-caregiver attachment. In Handbook of Attachment*, eds. J. Cassidy & P.R. Shaver*, pages 68–88. Guilford Press, London, 1999.

P. Willatts. Development of problem solving in infancy. In Infant development*, eds. A. Slater & J.G. Bremner*, pages 143–182. Lawrence Erlbaum, London, 1989.

P. Willatts. Beyond the 'couch potato' infant. In Infant development: recent advances*, eds. J.G. Bremner, A. Slater & G. Butterworth*, pages 109–135. Psychology Press, Hove, 1997.

P. Willatts. Development of means-end behaviour in young infants: Pulling a support to retrieve a distant object. *Developmental Psychology*, 35(3):651–667, 1999.

M.S. De Wolff and M.H. van IJzendoorn. Sensitivity and attachment: A meta-analysis on parental antecedents of infant attachment. *Child Development*, 68:571–591, 1997. 2.

M. Wooldridge. *An introduction to multiagent systems*. Wiley, chichester, 2002.

I.P. Wright. *Emotional agents*. PhD thesis, School of Computer Science, The University of Birmingham, 1997. (Available online at http://www.cs.bham.ac.uk/research/cogaff/).

J. Wyatt. Reinforcement learning: a brief overview. In IPerspectives on Adaptivity and Learning*, eds. I.O. Stamatescu*, pages 243–264. Springer, New York, 2002.

Xin Yao and Yong Liu. Fast evolution strategies. In Peter J. Angeline, Robert G. Reynolds, John R. McDonnell, and Russ Eberhart, editors, *Evolutionary Programming VI*, pages 151–161, Berlin, 1997. Springer.