Implementing a Theory of Attachment: A Simulation of the Strange Situation with Autonomous Agents

Dean Petters
University of Birmingham
School of Computer Science
Birmingham B15 2TT, United Kingdom
d.d.petters@cs.bham.ac.uk

Abstract

This paper presents a number of autonomous agent simulations as theories of infant attachment. In particular, the simulations model the behaviour of one year old human infants observed in Strange Situation studies. A key question this study is attempting to answer is, for a normal population of infants, why do responses tend to cluster into three groups of similar response? In addition this work explores what architectural mechanisms underlie the ability to switch goals, learn different predispositions to hold goals and learn about the effectiveness of different actions. Simulations are described that provide explanations of these phenomena. Computational experiments were carried out to investigate the relative effects of two learning mechanisms. Firstly, infants learned about their attachment relationships solely from carer responses in episodes when they experience anxiety. Secondly, infants learned solely from episodes where they experience warmth or rejection in a non-anxious environment. Results show that learning from episodes where anxiety was experienced causes greater clustering of responses. This result explains the clustering phenomena at a high level of explanation that is independent of any low level physiological implementation.

1 Introduction

This work is an attempt to understand what is going on inside a one year old human infant's mind/brain as it engages in attachment behaviours, in particular the behaviours observed in the Strange Situation Experiment (Ainsworth *et al* 1978). 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[1].

The Strange Situation Experiment is not strictly an experiment, rather it is a standardised laboratory procedure that presents infants with a controlled and replicable set of experiences. The eight episodes of three minute duration 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. The final episode involves the mother returning to her one year old infant after the infant has been left alone for three minutes in the unfamiliar setting. Ainsworth *et al* (1978) found that an infant's response to such a reunion can act as a shorthand for the infant's home 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.

Various theories have been put forward to explain these correlational phenomena (Goldberg 2000, chapters 4 and 5). Some emphasise evidence that the correlations are due to innate temperamental traits possessed by the infant, others emphasise evidence that infants are adapting to the caregiving environments provided by their mothers. In a meta-analysis of thirteen studies Goldsmith and Alansky (1987) found some support for both views. However, their analysis showed that infants' Strange Situation behaviour more closely matches their mothers' previous behaviour than their own. This suggests that the major direction of causation is from patterns in caregiving style to patterns in infant response. This work is reporting the creation of a set of agent-based simulations that act as explanatory theories for the attachment behaviours seen in the Strange Situation. All the simulations reported in this paper 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 flexible enough for straightforward inclusion of these mechanisms in future work.

A number of architectures have been implemented that vary in the complexity of the structures, mechanisms and representations that they possess. The simplest architecture, termed the **goal-switching** (**GS**) architecture, 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 the short time period of the Strange Situation. Two implementations have been created of 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-warmth** (**GLW**) architectures. Computational experiments have been carried out to compare the performance of these architectures, and the results of these experiments make an important contribution to the understanding of how individual difference categories in attach-

[1] Infant and adult autonomous agents and other objects have been implemented in a two dimensional virtual world using the simagent toolkit http://www.cs.bham.ac.uk/research/poplog/packages/simagent.html

ment behaviour may form.

A further reactive architecture has been implemented that incorporates the **GL** 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 principle design element that sets the **RAL** architecture apart from **GL** architectures is the manner in which the generated goals interact, before final selection. None of the **GLA**, **GLW** or **RAL** architectures are intended to represent sense-decide-act models of cognition. The time-slicing of these simulations is meant to approximate continuous and concurrent processing. 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 reactive goal generators and the same repertoire of possible actions as the **RAL** architecture. The **HAR** architecture differs from the **RAL** architecture because instead of a reactive, preprepared arbitration of the competing goals, this architecture allows on-line simple reasoning about what actions to take. This reasoning occurs in a deliberative subsystem in which processing is initiated from the reactive level and actions are carried out at the reactive level. However within the deliberative subsystem processing does conform to a sense-decide-act design.

2 The design process and scenario formation

Adapting psychological data and theory so that it can be represented in an agent based simulation involves several processes. These include the analysing and ascribing of function to behaviours, abstracting them, and then presenting the behaviours in a form that can act as a specification of requirements for the purposes of simulation design and evaluation. Petters 2004 gives an overview of how these processes have been applied to the problem of simulating infant attachment behaviour. Petters 2005 gives a more detailed review of the attachment literature and gives further details of how the scenario for the simulations described in this paper were drawn up.

The key data incorporated in the simulation are derived from the study of Ainsworth *et al* (1978). This study found that when the many measures of infant response in the Strange Situation were evaluated using a Multiple Discriminant Function Analysis the infants were found clustered into three major categories of attachment style[2], labelled: Avoidant (type A), Secure (type B) and Ambivalent/resistant (type C).

Avoidant infants respond to their mothers on reunion in the Strange Situation by: not seeking contact or avoiding their mother'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 that measured the physiological correlates of stress for infants undergoing the Strange Situation have been carried out and show that the 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 mothers 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. The avoidant pattern of behaviour is considered to have the function, for the infant, of minimising the chance of receiving actual or psychological harm.

Secure infants respond to their mothers on reunion in the Strange Situation by approaching her 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 mothers 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. Secure infants are considered to initially possess the goal of increasing their safety but then change relatively quickly in reunion episodes to hold the goal of exploration.

Ambivalent/resistant infants respond to their mothers on reunion in the Strange Situation by: not being comforted and being overly passive or showing anger towards their mothers. 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; 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 infants were more angry, and they cried more. Ambivalent/resistant infants are considered to possess the goal of increasing their safety during the three minutes of the final reunion episode, and using increased protest to gain this safety.

3 An architecture for switching between goals

All infants that take part in the Strange Situation exhibit a range of behaviours. What differentiates the infants is not that some show fear and cry and others do not, but rather the particular frequency and intensity of particular behaviours in the specific episodes. The **goal-switching** (**GS**) architecture, shown in figure 1, can produce patterns of switching behaviour where all available goals are active at some time. The **GS** architecture has three major divisions: a perceptual subsystem comprising goal activator modules, a central selection and arbitration subsystem, and an action subsystem. (see figure 1). There are four goal activator modules. Each can be viewed as possessing an implicit goal or function, and each provides proposals for action with a variable activation level. They include: the fear module (maintaining safety from unfamiliar objects); the anxiety module (maintaining safety from remoteness of the carer); the exploration module (learning about objects); and the socialisation module (learning about agents). 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. The second is to mod-

[2]Recent studies have categorised a fourth type of disorganised attachment style that forms a very small proportion of infants in the general population and which is not currently represented in the simulation

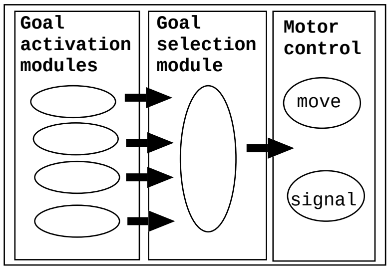

Figure 1: The **GS** architecture: A Reactive Design for choosing the top level goal of an infant. The goal activation modules that are represented are exploration, socialisation, anxiety and fear.

ulate this activation according to internal factors. The third task is to make selections for targeting. There are also two action subsystems, moving and signalling. 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. What this means is that 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.

The exploration and socialisation goal activators have a similar structure. Both sense the environment and pass increased activation when a target of exploration or socialisation is close by and in the focus of their sensory apparatus. However activation for exploration and socialisation rises over time even if no objects or agents are within sensory range. This is because the internal state of 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 these goal activators is that infants possess a kind of hunger for learning about objects and agents. The longer infants don't experience novel object or agent interactions the higher the activation for this kind of experience becomes. This state of affairs is hand-coded in the simulations described here, however in reality the same behaviour may emerge from perceptual systems that amplify signals from novel sources.

The internal workings of the anxiety module operate differently from the exploration and socialisation goal activators. If the carer remains very close to the infant then the anxiety module passes no activation. Activation only starts to be passed when the agent senses that the carer is beyond a Safe-range distance. At this point the activation level of the anxiety module may be below 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 cause behaviour to switch. If a carer remains beyond the safe range distance for a long period then the activation of the anxiety module will win-out over other goals and the infant will act to decrease the proximity to the carer. This can be accomplished by moving towards the carer and signalling to the carer so that the carer responds by moving back within the Safe-range distance.

The **GS** architecture can be viewed as a simplified version of a Contention Scheduling mechanism (Cooper and Shallice 2000). In both architectures different goals are generated and compete in a winner takes all contest to see which goal becomes active and directs the external behaviour of the infant. Both architectures link perception to action, but the architecture for Contention Scheduling does so via a hierarchically organised network of action schemas, which represent goals and multiple levels of subgoals. The architecture for infant attachment is flat, with one level of goals activating a single level of atomic actions. These actions are also more abstract than those found in the model of Contention Scheduling. This is because the implementation of contention scheduling is modelling tasks, such as making a drink, in greater detail for a relatively short period of time, whereas the architectures described here are intended to capture patterns of behaviour that form over an entire year. Future work may involve augmenting the two infant actions of moving and signalling with the inclusion of less abstract and more numerous atomic actions. These may need to be respresented at multiple levels of abstraction, and hence require hierarchical organisation. The main drawback of the **GS** architecture as a theory of infant attachment behaviour is that it doesn't possess any mechanisms that allow the infant to learn the effects of choosing different goals. In reality we find that infants learn about the responsiveness of their carers and change their behavioural predispositions accordingly.

4 An architecture for choosing goals

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. An infant architecture can perform optimally when it possesses an accurate measure of its security. This means it can maximise its learning opportunities whilst keeping within an acceptable level of risk. In the simulations described here the goals of safety and learning are balanced. Each type of learning goal is continuously activated and only extinguished when it has gained control and been 'exercised'. Goals related to safety are only activated when some form of threat is sensed. We can view the Safe-range distance as representing the infant's confidence in the responsiveness of the carer. An infant that has learnt that its carer provides a lot of security can allow the carer to move further before it attempts to regain proximity. We can view the infant as leaving its security in the hands of its carer. An infant that has learnt its carer is unreliable needs to act itself to maintain its 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.

In computational experiments that model the development of different attachment styles all infants start with Safe-range parameters set at the same level. How then do infants learn about the effectiveness of their carers? This section describes two very similar architectures that are based upon the **GS** architecture but that adapt to the pattern of care that they receive by changing the value of their Safe-range parameter.

Infants possessing the **goal-learning-from-anxiety** (**GLA**) architecture only learn about the secure-base nature of their carer from their carer's response when they experience in-security. This information is used to change the Safe-range parameter, increasing it when carers respond effectively and decreasing it when carers are insensitive to the infant's bids for response. Infants possessing the **goal-learning-from-warmth** (**GLW**) architecture learn information about their carer's effectiveness from the carer's warmth of response, in episodes of socialisation not anxiety. In the case of the **GLW** the infant is learning how to respond in situations where they perceive danger, however this learning is occurring in a different type of situation, where they don't perceive threat. The **GLA** and the **GLW** mechanisms have differing advantages and disadvantages, and in reality the infant may possess both types of mechanism. The key trade-off is between the closeness of fit of the episodes that the infant is trying to learn from and the number of those incidents that are available for learning. In other words, how does evolution provide infants with mechanisms to learn about their carer's likely response in times of danger, such as potential abandonment. Infants can learn about their carers likely response from a small number of episodes that closely match episodes of potential danger. Or they can assess the much more numerous responses in episodes that match less well.

The **GLA** and the **GLW** architectures both provide explanations for infant behaviours in the reunion episode of the Strange Situation. In computational experiments involving both types of infant architecture Secure infants get back to attentive exploration earlier than other infants because their Anxiety subsystem is less activated. In an unfamiliar environment the Exploration subsystem will always posses at least moderate activation and when the activation of the Anxiety subsystem drops below this level the behaviour switches to exploration. A figurative way of saying this is that, although they have just undergone a distressing separation, Secure infants feel safer in the reunion episodes than the other infants do. Exploration is activated because less threat is perceived. This is because Secure infants have learnt, either from episodes of anxiety or socialisation in their previous home experiences, that their carers provide reliably effective responses. Insecure infants have not previously received consistently effective responses in episodes of anxiety or socialisation. These infants therefore have higher levels of activation in their anxiety module in the reunion episode of the Strange Situation. Sroufe and Waters 1977) 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, alter the level of security perceived by the infant. We can view the Safe-range parameter as an abstract representation of a much richer measure of security that exists in reality. This does not change the core of our theory, whereby incidences of high or low carer responsiveness change the criteria by which the level of future security is assessed.

So how do the performance of the **GLA** and the **GLW** architectures differ? Figure 2 shows the results of ten experiments, where identical carers were matched with infants which possessed the **GLA** architecture. All the carers' responsiveness levels were set to be intermediate between being secure and insecure. In these experiments, five infants ended up secure and five ended up insecure, the difference due to small random variations that occured early in each simulated

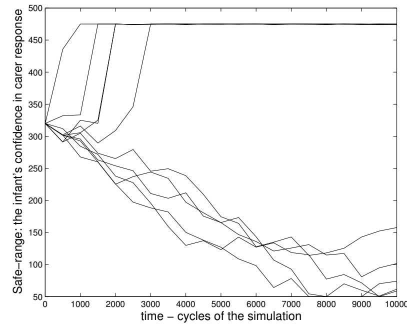

Figure 2: Infant Safe-range changing over time, with infants only learning from episodes of anxiety. Each infant had an identical initial state, differences resulting from random elements early in the 10000 cycles (Learning updates only occur on a small proportion of cycles).

run. We can explain this bifurcation by considering that after each episode the infant updates its Safe-range distance. If the infant gets a late response, 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 within the sphere 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 will become increasingly small. This means the infant is becoming chronically untrusting of the carer's performance. The reverse obviously holds true for carers that carry out a series of prompt returns. 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 being at either extreme of caregiving. Any successful response makes another successful response more likely, any unsuccessful response increases the probability that the carer will have further to travel to get within the bounds of Safe-range distance.

Another way of saying this is, *prompt* responses give large reinforcement signals and this results in the Safe-range distance being increased, whereas *delayed* responses produce negatively reinforcing 'punishment' signals. In the **GLA** architecture, '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 is more likely to go further than the new Safe-range distance.

The same experiments were repeated with infants possessing the **GLW** architecture. Figure 3 shows the results. As with the previous results, all the carers' responsiveness levels were set to be intermediate between being secure and insecure. In these experiments there is no positive feedback and

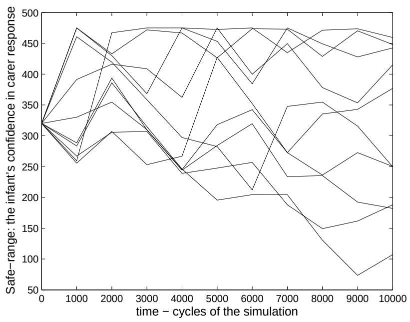

Figure 3: Infant Safe-range changing over time, with infants learning about carer response only from interactions with no anxiety (Learning updates only occur on a small proportion of cycles).

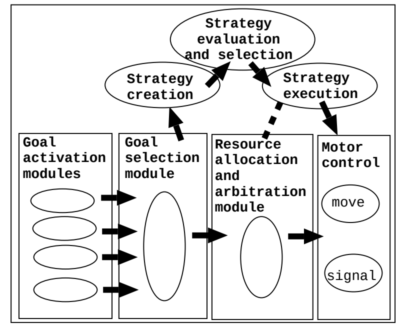

Figure 4: The Hybrid Design **HAR** . Deliberative mechanisms provide a secondary route to action, activated as a result of interrupts to reactive selection and arbitration and imposing different actions. The dashed line represents an inhibitory interaction initiated by the Strategy execution module.

the attachment status of the infants varies from secure to insecure without clustering. This is because the learning of the Safe-range distance is independent of the infant's behaviour in episodes of anxiety. In reality infants are likely to possess a balance of learning mechanisms of both the **GLA** and **GLW** types.

5 An architecture for choosing actions

The behaviour we are attempting to explain in this section is the Avoidant behaviour of type A infants. The previously described **GLA** and **GLW** architectures adequately support the behaviour of Secure type B infants and Ambivalent type C infants. When these infants execute goals the choice of action is straightforward. Secure infants either possess the goal of security and move and signal to their carer or they possess the goal of exploration and move to a target object. Ambivalent infants predominantly possess the goal of security and move to and signal to their carer. The puzzle is that Avoidant infants hold the goal of security but limit how close they move to their carers and do not signal to their carers. This work describes two architectures, the **reactive-action-learning** (**RAL**) and the **hybrid-action-reasoning** (**HAR**), which can both support the patterns of home and Strange Situation behaviours observed in all three attachment styles.

The **RAL** is based upon the **GLA** architecture, but has been augmented so that it can support a phenomenon that ethologists have termed 'displacement activity' (Bowlby 1982, Main and Weston 1982). 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, or flee from, but instead starts to groom itself (Bowlby 1982). Displacement activities occur when two strongly activated behaviours 'cancel each other out', and a seemingly inappropriate behaviour becomes active. 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. To support displacement phenomena the **RAL** architecture possesses additional components. These are an Avoid-pain goal activation module, and a more sophisticated resource allocation and arbitration module. This allows a greater repertoire of possible actions. A **RAL** infant with highly activated goals of safety and pain avoidance can signal and move to the carer, just move to the carer and not signal, or even not move to the carer but engage in some other action with a low activation such as exploration. This last option means that the infants can engage in displacement activity. 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 differ in their reunion behaviour from Insecure Ambivalent infants because at close distances to their carers the goal of avoiding close contact inhibits active expression of secure goal behaviour, leading to behaviours linked to the exploration goal being activated as displacement behaviours. Two criticisms of the **RAL** architecture are: firstly, that it introduces a hand-coded 'custom' mechanism that is local rather than general; and secondly, it proposes that actions are chosen in a single route through the architecture, and research in cognitive development provides evidence of dual routes to action in nine month old infants (Willatts 1999).

Crittenden (1995) explains Avoidant infant behaviour in reunion episodes without recourse to the idea of displacement activities. 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"*(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 "*(Crittenden 1995). A hybrid architecture called **hybrid-action-reasoning** (**HAR**) has been created which develops these ideas (see figure 4). 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. In the **HAR** architecture represented in figure 4 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 continuous 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. The major criticism of the **HAR** architecture is that it introduces a whole range of advanced structures and processes in a single agglomerate package that better fits the observed capabilities of eighteen rather than nine month old infants (Willatts 1999).

The **RAL** and **HAR** architectures currently both show clustering of responses in separation episodes. However this is due to discontinuities which are implementational rather than core aspects of both architectures. Carer levels of aversion to contact remain fixed and both infant architectures force intermediate levels of infant contact aversion to be externally manifested as either avoidance or proximity seeking. Ongoing work involves introduction of mechanisms that will allow positive feedback between carer and infant levels of aversion, so that clustering emerges through dynamic interaction in a theoretically deeper manner.

6 Conclusion

This paper has described a number of architectures that are intended to act as explanatory theories of infant attachment behaviour. They provide a greater degree of precision and detail than existing linguistically termed theories (Bowlby 1982, Goldberg 2000). These theories have been implemented using autonomous agent techniques that are based upon simplified models of complete systems that endure over time. The infant and carer agents in the simulation can therefore respond to each other's behaviours in a dynamic and adaptive manner. This work differs from many developmentally oriented agent-based simulations because of its concentration on central processing and impoverished perceptual capacity (Schlesinger 2001).

None of the architectures presented here is a finalised theory of attachment behaviour and the Strange Situation. The **GS**, **GLA** and **GLW** architectures fail when they are evaluated against the full set of requirements in the scenario. For example, they don't provide an explanation for how Avoidant infants choose actions. The **RAL** and **HAR** architectures don't fit what is known about infant cognition from other sources (Willatts 1999). In addition, failings of the scenario itself include the fact that there is no way an infant in this simulation could ever gain an adult's proximity by smiling or behaving well. This counts as a limitation of the ontology of the simulation; the simulation simply doesn't support this kind of behaviour. This is, therefore, a fundamental limitation of the simulation as it now stands. Future work involves widening and deepening of the scenario and development of an architecture that will be an intermediate between the **RAL** and the **HAR** architectures. Despite the many limitations, this work has demonstrated a means by which theories that were previously expressed only in words and diagrams can be transformed into working models. These models have been tested to see if they have the claimed properties, showing also that different models produce demonstrably different consequences.

6.1 Acknowledgment

Thanks to my PhD supervisor, Aaron Sloman, for much help and encouragement.

7 References

Ainsworth, M. *et al* (1978). *Patterns of Attachment: a psychological study of the strange situation.* Hillsdale,. NJ: Erlbaum.

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

Cassidy, J. & Berlin, L.J. (1994). The insecure/ambivalent pattern of attachment. *Child Development*, *65(4)*, 971–991.

Cooper, R. & Shallice, T. (2000). Contention scheduling and the control of routine activities. *Cognitive neuropsychology*, *17(4)*, 297–338.

Crittenden, P.M. (1995) Attachment and Psychopathology. In S. Goldberg R. Muir & J. Ker(Eds.), *Attachment theory: social, developmental and clinical perspectives* Hillsdale, NJ: The Analytic Press

Goldberg, S. (2000). *Attachment and development*. London: Arnold.

Goldsmith, H.H. & Alansky, J. (1987). Maternal and infant temperamental predictors of attachment: a meta-analytic review", *Journal of clinical and consulting psychology*, *55(3)*, 805–16.

Hertsgaard, L., Gunnar, M., Erickson, M.F., &, Nachmias, M. (1995). Adrenocortical Responses to the Strange Situation in Infants with Disorganised/Disoriented Attachment Relationships. *Child Development*, *66(4)*, 1100–1106.

Main, M., & Weston, D.R. (1982). Avoidance of the attachment figure in infancy. In M. Parkes & J. Stevenson-Hinde (Eds.), *The place of attachment in human behaviour*. New York: Basic Books

Petters, D. (2004). Simulating infant-carer relationship dynamics. *Proc AAAI spring symposium 2004: Architectures for modelling emotion - Cross-Disciplinary Foundations,* number SS-04-02 in AAAI Technical reports, (pp. 114–122). Menlo Park, CA: AAAI Press.

Petters, D. (2005). Building agents to understand infant attachment behaviour. *Proceedings of the workshop: Modelling natural action selection* (pp. 158–165). University of Sussex, Brighton: AISB Press.

Schlesinger, M. (2001). The agent-based approach: A new direction for computational models of development. *Developmental Review*, *21*, 121–146.

Sroufe, L.A. & Waters, E. (1977). Attachment as an organisational construct. *Child Developmental*, *48(4)*, 1184–99.

Willatts, P. (1999). Development of Means-end Behaviour in Young Infants: Pulling a Support to Retrieve a Distant Object. *Developmental Psychology*, *35(3)*, 651–667.