

# Modelling Emotional Attachment: an Integrative Framework for Architectures and Scenarios.

Dean Petters

School of Social Sciences  
Birmingham City University  
UK

Email: dean.petters@bcu.ac.uk

Everett Waters

Department of Psychology  
SUNY Stony Brook University  
USA

**Abstract**—Humans possess a strong innate predisposition to emotionally attach to familiar people around them who provide physical or emotional security. Attachment Theory describes and explains diverse phenomena related to this predisposition, including: infants using their carers as secure-bases from which to explore, and havens of safety to return to when tired or anxious, the development of attachment patterns over ontogenetic and phylogenetic development, and emotional responses to separation and loss throughout the lifespan. This paper proposes that one way for computational modelling to integrate these phenomena is to organise them within temporally nested scenarios, with moment to moment phenomena organised within ontogenetic and phylogenetic sequences. A number of existing agent-based models and robotic attachment simulations capture attachment behaviour, but individual simulations created with different tools and modelling approaches typically do not integrate easily with each other. Two ways to better integrate attachment model are proposed. First, a number of simulations are described that have been created with the same agent-based modelling toolkit, so showing that moment to moment secure base behaviour and the development of individual differences in attachment security can be simulated with closely related architectural designs. Secondly, an integrative modelling approach is proposed where the evaluation of, and comparison between attachment models is guided by reference to a shared conceptual framework for architectures provided by the CogAff schema. This approach can integrate a broad range of emotional processes including: the formation of a set of richer internal representations; and loss of control that can occur in emotional episodes.

## I. INTRODUCTION

Attachment Theory describes how our closest relationships develop and function across the life span. Attachment bonds are formed early in infancy and can reform and develop through the life-span. Attachment Theory has biological, cognitive, cross-cultural and psychopathological perspectives [1], [2]. It doesn't just explain moment to moment interactions between infants and carers. Attachment interactions can also be observed in all stages across human lifespan development from infancy to old age, and in adult romantic relationships. It is also a comparative and evolutionary theory. So because Attachment Theory explains phenomena over a range of timescales from moment to moment interactions to ontogenetic and phylogenetic development a complete modelling approach to attachment phenomena needs to be capable of simulating temporally nested scenarios. Sometimes the modeller will just want to explain a few minutes of interaction, but on other occasions modelling attachment development over a lifespan

or over evolutionary trajectories may be desired.

The developmental processes described by Bowlby involve intimate interaction between lower level processes, such as simple reflexive responses, and emerging higher level structures and mechanisms. New resources are created over time. Integrating elements into a system depends on internal factors like biases in infant learning abilities and external factors such as information and structure in the expectable caregiving environment. So an important issue for modelling of attachment phenomena is that the control system as proposed by Bowlby is not just preformed and waiting to be triggered or maturing without experience. Rather it is highly interactive between inner and outer factors, through its construction over ontogenetic development, and operation in moment to moment interactions.

During this process of ontogenetic change there is typically stability in the individual differences in behavioural patterns exhibited by individuals. This continuity is in part explained because often early appearing fragments of instinctive behaviour are integrated into later appearing complete sequences with their normal mature functional consequences. According to Bowlby the attachment control develops from being reliant on simple mechanisms such as simple reflexes through many intermediate forms such as fixed action patterns, chaining of simple actions, goal corrected mechanisms, simple plans, hierarchical plans, and Internal Working Models, to finally being mediated by complex high level representations such as natural language [1].

In addition to explaining ontogeny, Attachment Theory is also a theory of phylogenetic development. It was founded upon the Behaviour System concept from Ethology. In this framework Behaviour Systems control behaviours such as mating, fighting and feeding. Each Behaviour System carries out a species-specific function, and has survived in the genome because its functions contribute over evolution to biological fitness. Therefore Behaviour Systems are related to one meaning of the term 'instinct' [3]. The Behaviour Systems that Bowlby linked to attachment behaviour in human infants are the attachment, fear, sociability and exploration systems [1]. For Bowlby, behaviours resulting from the attachment behaviour system and the fear system have the predictable outcome of maintaining access and proximity to its primary carer and avoiding environmental hazards like predation. Bowlby proposed the concept of the Environment of Evolutionary Adaptedness (EEA) when explaining how

proximity seeking and maintaining systems can be adaptations [1]. In this view infants have evolved adaptations to keep safe by bringing their caregiver closer by, either by signalling or by actually moving towards their caregiver. However, infants also need to learn about the world and caregivers help structure infant learning. So infants possess adaptations that allow a balance to be struck between keeping close to the carer, and so keeping safe and learning in a structured manner, and exploring the broader world away from the carer. One of the evolutionary assumptions in Attachment Theory (from an adaptationist analysis of the EEA) is that if the infant perceives its current environment is risky, the balance between exploration and security should shift towards keeping more secure. Conversely, if the environment is perceived as less risky, the infant will benefit from shifting this balance towards more exploration. Another key assumption of the evolutionary approach in Attachment Theory is that, although infants in contemporary contexts are removed from the characteristics of the EEA, they still possess the innate triggers and learning biases sculpted by evolution to match the EEA rather than modern conditions.

Attachment development and individual differences have been investigated using the 'Strange Situation' [4]. This is a standardised laboratory procedure where all infants are presented with the same controlled and replicable set of experiences. The Strange Situation procedure was created to investigate how infant attachment behaviour was sensitive to context. What was found was that nested within the normative trends that illustrate infant's sensitivity to context are patterns of infant response reflecting the infant's confidence in the caregiver's responsiveness. To capture infant responses to changes in context, like separations and reunions, the Strange Situation procedure consists of 8 three minute episodes which are designed to activate, intensify or relax attachment behaviours of one year olds in a moderate and controlled manner. The infant and carer enter the laboratory setting together, but then undergo a separation, when the carer leaves from the room, before a reunion in a subsequent episode. After the first reunion episode the infant also meets an unfamiliar 'stranger' in the laboratory, before a further separation. In each episode the infant's behaviour is recorded from behind a two-way mirror. In the final episode the mother returns to her infant after the infant has been left alone for three minutes in the unfamiliar setting. The infant's response in the reunion episodes correlate strongly with patterns of maternal behaviour and infant responses intensively observed throughout the previous year. An infant's responses to reunion in the Strange Situation can therefore act as a shorthand for the infant's home relationship with their carer. Individual differences clustered into three distinct patterns, labelled: Secure (type B), Avoidant (type A), and Ambivalent (type C). More recent studies have categorised a fourth type of Disorganised pattern of attachment (type D) that is the least well characterised or understood type and describes a very small proportion of infants in the general population ([2], page 26). Each of the main categories also has subtypes.

The way that behavioural patterns from the Strange Situation are explained exemplifies why three temporal perspectives are needed for a complete explanation of attachment phenomena. Ainsworth speculated that some subtypes in the Strange Situation were due to the infant's short term state, and

arise because immediately prior short-term interactions have an effect on Strange Situation responses [4]. In this analysis, whether infants are tired or have just experienced insensitive responses will lead to different sub-group categorisation. Bigger differences in patterns of attachment are due to long term status arising from the many repeated interactions the infant experienced over the span of ontogenetic development. Lastly, according to Attachment Theory, the basic pattern of secure-base behaviour that all categories have in common results from the evolved structure of the attachment control system and the way it is biased to develop through an individual lifespan. So a control system for secure-base behaviour shared by all infants is an evolutionary adaptation.

So for attachment modelling to provide a fully comprehensive approach to explaining phenomena such as infant responses in the Strange Situation it needs to fulfil requirements from temporally nested scenarios. Moment to moment interactions exist within ontogeny, and single life spans are also situated within evolutionary trajectories.

Agent-Based Modelling (ABM) is a particularly valuable type of modelling for attachment research because of how this type of modelling type represents the passage of time. In ABM agents exist in a virtual environment over multiple time-slices. So dynamic interactions can unfold over time - with the granularity of time able to be made arbitrarily fine. This can therefore aid clear thinking about developmental phenomena that exist over various time-scales and how individuals engage in different dynamic interactions over these various time-scales. ABM can vary the granularity of time to simulate daily interactions over a year or second by second interactions over a few minutes or hours. Populations of agents can also be 'born', differentially reproduce and 'die', simulating multi-generational evolutionary changes. Real infants and carers exist within a continuing stream of time and can affect and sometimes control the nature of the sensations they receive from their external and internal environments. For example, if an infant turns the direction of their head, they can then see a completely different view from before that may activate different desires, goals or intentions. When there are two agents interacting the situation becomes more complex and contingent on the immediately previous context. When an infant or carer acts the other will react and change the situation and context for both of their next 'moves'. ABM can capture these dependencies and over lengthy passages of time agents can engage in extended exchanges or 'dialogs' with their environments where each new step in the dialog is influenced by all the preceding steps ([5], p 28).

Whereas many phenomena of interest in cognitive development are captured in controlled laboratory settings, attachment phenomena, such as secure-base behavioural patterns, can be observed in naturalistic settings. Schlesinger and Parisi [5] explain how the distinction between controlled laboratory settings giving experimental results and naturalistic settings giving ethological observations maps onto the distinction between computational models with online and offline sampling:

*"The key difference between offline and online sampling concerns how time is experienced by the model. During offline sampling, each input pattern constitutes a discrete, independent moment in time.*

[...] In addition, although the model may receive feedback after each output (e.g., a training signal), it does not experience any short-term or long-term consequences of the response. In many respects, the model's experience during offline sampling is analogous to a human subject in a typical psychology experiment: (1) a random sequence of discrete training exemplars is presented, (2) responses are independent (i.e., one response does not affect another), and (3) interference effects between successive exemplars are treated as noise and consequently minimized by repeated training runs."

"If we think of offline sampling as a procedure analogous to a laboratory experiment, then online sampling is akin to an ethological study: we first design a virtual organism and a quasi-natural learning environment (or perhaps, a robotic organism in a real environment) and then leave the organism free to explore and learn through interaction with its environment. [...] Because each behavior is followed by not only a potential training signal (where appropriate), but also the sensory consequences, our organism experiences both the immediate and the long-term effects of its behavior." ([5], p 126)

So ABM is a useful approach to modelling how two or more agents respond to each other in continuous time. Though implemented on digital computers and necessarily confined to discrete time steps, these can be set very finely.

The agents in agent-based models can incorporate within their control architectures the kinds of structures and mechanisms that Bowlby set out in his description of the attachment control system. ABM can also support a focus on understanding cognitive processes of organism-environment interaction which are situated across time and used to model the developmental role of sensory self-selection and perception-action linkages [5]. Sensory self-selection may be relevant to the avoidant pattern of behaviour, where insecure infants may avert their gaze from their caregivers as a form of emotion regulation. Perception-action linkages may be involved in close coupled interactions between infants and carers. For example, when infants 'sink-in' when being held or engage in reciprocal vocalisations.

## II. LITERATURE REVIEW OF COMPUTATIONAL ATTACHMENT MODELS

Only a few years after Bowlby's [1] first formulation of the Attachment Control System, Bischof [6] presented an attachment model that could run as a simulation on a computer. The simulation represented carers and infants as agents within a virtual environment (though they are not termed agents in Bischof's paper). The infant agent's social motivation is formalised as a cybernetic control system that is based upon the behaviour systems that Bowlby postulated. So a key advantage of Bischof's model is it instantiates attachment control within a broad motivational system. This attachment control system includes computational elements that produce attachment security and exploration behaviour implemented as feedback control circuits. This means that the infant modelled in this simulation either possesses attachment or exploration

goals at any one time and switches between moving towards the simulated mother and exploring the environment.

More recently, Hiolle and Canamero and co-workers [7], [8] constructed a robot with a reactive perception-action architecture implemented as a neural network that allowed a robot to imprint the perceptual attributes of a person standing in front of it. After a short period of imprinting the robot could then follow this person (as long as they did not move too fast!). A notable achievement of this robotics simulation is that the robot learns to move towards an attachment figure by 'imprinting' the perceptual 'image' of the attachment figure and then acting to maintain this 'image' in its perceptual system. So this system might also learn to signal instead of move if this brought the carer back and hence conserved the imprinted perceptual image. Stepping aside from the engineering aims of this research and instead considering it as a model of human attachment there is clearly a mismatch between what this demonstrates and human performance because human infants only learn to move when they already possess quite sophisticated abilities in bringing their carers to them. This robot knows nothing about the world except its attachment figure's imprinted image and the actions needed to maintain this image in perception. So it is more like a precocial species than altricial human infants. Hiolle and Canamero and co-workers [9] also constructed robots that explored their surroundings but which curtailed this exploration and returned to the proximity of their carer when they experienced too much 'arousal' resulting from too much stimulation from novel perceptual input. So these simulations demonstrated systems learning about objects but not learning about carer effectiveness in how to support effective exploration.

Multi-generational evolutionary simulations have been constructed which explore the evolutionary pressures linked to care-giving and feeding [10], [11] and lifetime learning and protection [12]. Parisi and co-workers created simulations of artificial neural networks in virtual 2D environments that forage for simulated food. They showed that 'child' neural networks will evolve to follow 'adult' networks because the 'adult' networks provide food which the 'child' networks cannot gather themselves. As the simulation progresses through many generations the networks evolved a tendency to form small social aggregations of kin-related individuals that moved together towards food. However, this coordination into 'families' was not driven by agents in the adult phase of their development. Rather, it was the predisposition to maintain proximity to adult networks that evolved to be expressed in the infant developmental phase that drove this emergence of foraging 'families'.

In similar research, Bullinaria [12] used evolving neural networks in an artificial life exploration of how life-time learning, periods of parental protection from predation, and age at first reproduction, interact and affect life-history evolution. He found that higher learning levels occurred when periods of parental protection were relatively long and procreation by the protected 'child' neural networks was precluded. If procreation is allowed earlier in a lifecycle and can occur during the protection period then lower levels of learning occur - because the evolutionary pressure to learn too quickly (to achieve at an early age the fitness required to succeed in the competition to procreate) results in poorer final learned performance. So

neural networks evolved predispositions to spend time procreating rather than learning! As Bullinaria ([12], p 406) notes, in this particular research, as with most artificial life simulations, simplifications are required to render them computationally feasible. There are several obvious simplifications that limit their relevance to Attachment Theory: parental protection in the simulation does not influence how learning occurs and what is learned; either protection periods end when procreation starts, or both can occur at full strength at the same time (so no gradual decrease of protection), no grand-parenting effects in childcare; no individual differences in protection strategies are possible; protection and other parental investment strategies cannot be contingent on environmental conditions.

### III. MODELLING SECURE-BASE BEHAVIOUR IN SIMPLE AND ADVANCED SCENARIOS WITH RELATED ARCHITECTURES

#### A. A basic secure base scenario - an infant exploring a play area

A clear example of secure-base behaviour is provided by a naturalistic study of toddlers between ages 1 and 2 which investigated how infants balance the opportunity for exploration and the security provided by their carers [13]. In London parks, the infant and carer pairs were covertly observed, in observation periods averaging 15 minutes duration. This study found that most infants moved away from their carers to explore, but kept within a caregiver's line of sight, and periodically 'checked-in' by gaining the attention of their carers or by moving back to closer proximity. This main scenario describes infants switching between approach to the caregiver and exploratory behaviour in naturalistic environments - so showing cycles between being nearer to the carer and then exploring further afield. It can be augmented with a number of related mini-scenarios ([2], p. 51-78), including: infant signalling to their carer agents and relying upon their carer agents to provide energy and other physiological needs; infants seeking close contact when fatigued; that close contact with caregiver has a restorative effect on infant play and exploration; infants being reassured by communication over distance through signalling as well as close physical contact; and coy behaviour as a mixed response to fear and sociableness towards strangers.

#### B. A reactive architectural model of basic secure base behaviour

A simulation that possesses a richer internal organisation than those described in the literature review in section II (but which is still over-simple) is the reactive agent architecture illustrated in Figure 1. This fulfils the requirements set out by the basic secure-base scenario (and related mini-scenarios) ([2], p. 51-78). In this work infant, caregiver and stranger agents were created and evaluated in a 2D virtual environment where they interacted with each other and objects representing food, furniture and toys. The infant's actions can include changing its physical location and signalling with an affective tone from positive (like smiling) through neutral to negative (like crying) to very negative (like 'hard crying'). Figure 1 illustrates the four Behaviour Systems the architecture possesses: attachment-security, exploration, wariness, and socialisation, and it is the interaction of these behaviour systems that produces behaviour




Fig. 1. A reactive attachment architecture. Safety, exploration, socialisation and physical need 'behaviour systems' gain activation levels as a result of internal and external factors and compete to control actions for movement and signalling in a 'winner-take-all' selection mechanism

patterns that match the requirements set out by the scenarios. In the infant control architecture the highest activated behaviour system controls the actions that are carried out by the agent. Switching behaviour patterns occur because when goals are satisfied activation for this goal is reset and variables that activate safety, exploration, social interaction, and physical contact, rise and fall. The activation levels vary because inputs to component behaviour systems vary according to the current context - for example, placing a novel toy right in front of an infant agent leads to higher activation of the exploration behaviour system irrespective of other activation levels. If several behaviours are activated only the one with the highest activation level generates a goal.

How should this architecture be evaluated? It produces abstract patterns of secure-base behaviour, but then so does Bischof's 1975 simulation, described in section II. It might be evaluated by how closely it satisfies the modelling scenarios, or the mechanisms it uses to do so, or because it provides a good foundation for broader and deeper scenarios [2], [14]. Clearly it is not satisfactory to have a model with an implausibly sophisticated architecture - this can simulate everything but not the way an infant would. But also, overly simple architectures can be bad models if the researcher wants to use them as the basis for explaining how behaviour patterns and competencies develop. For example, attachment-focused simulations are not the first computational models to simulate these distinctive behavioural patterns of cyclical recession from a target and then return. Much older examples of this phenomenon are provided by the pioneering cybernetics research on photo-tropic robots created between 1912 [15] and the 1950s [16]. These photo-tropic (light seeking) robots move around attraction points provided by light sources in the same manner that agent-based and robotic attachment simulations involve infant agents or robots moving around an attraction point of their carer-agents location. Such basic 'tropic' robots demonstrate that abstract patterns like secure-base behaviour can be modelled with relatively simple systems that do not serve as a good foundation for integrating existing and future information processing structures and mechanisms set out in Attachment Theory. So one way to evaluate the architecture illustrated in figure 1 is to show that this moderately complex architecture can be made more complex still by forming the reactive component of a hybrid architecture that would form the basis of a better attachment model for older children,

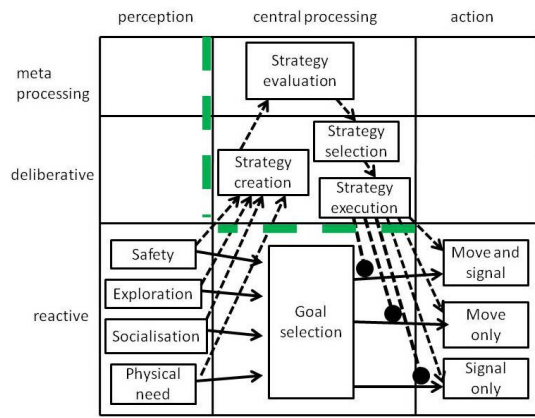


Fig. 2. A hybrid attachment architecture with reactive, deliberative and meta-management subsystems. The resource constrained deliberative subsystem takes input from the reactive-subsystem, carries out ‘look ahead’ reasoning, and can inhibit the reactive subsystem and execute alternative actions. The green dashed line represents the fact that in the human attachment system deliberative and meta-management processes require attention and so are resource bound, which limits the number that can be concurrently active.

adolescents and adult.

### C. A hybrid architectural model of basic secure base behaviour

Figure 2 illustrates a hybrid architecture which situates reactive subsystems alongside a deliberative planning subsystem (that allow ‘look-ahead’ reasoning) and a simple meta-management subsystems (where cognitive meta-processes operate on other cognitive processes) [2, p. 103-151]. In this hybrid architecture attentive processes occur within the bounds of a resource limited variable attention filter. These resource bound serial deliberative processes arise from non-attentive reactive or perceptual processes which operate in parallel. Reactive motive generativators therefore are triggered and activated by any possibly relevant internal and external events. In the attachment domain there will be possible threats but also possible exploratory and social opportunities. When these conditions are met a motivator is constructed which may ‘surface’ above the attentional filter and be operated upon by processes in the deliberative or meta-management levels. Amongst the deliberative attachment processes generated by motivators are the creation, selection, and execution of action plans. Evaluation processes occur in the meta-management layer. This hybrid architecture can produce the same external behaviour patterns as the reactive architecture which it extends. What it adds is a reconceptualisation of the attachment control system in light of insights from dual process theories in psychology and artificial intelligence [17].

Both architectures in figures 1 and 2 possess reactive learning mechanisms so that toy objects and ‘strangers’ become familiar. But neither architecture possesses learning mechanisms than change the long term relative tendencies for particular behaviours to become activated. That is, neither architecture changes its predisposition to explore or seek security according to evidence from the level of responsiveness or sensitivity of the caregiving it receives. This deficiency is remedied by the architecture described in the next section III-E.

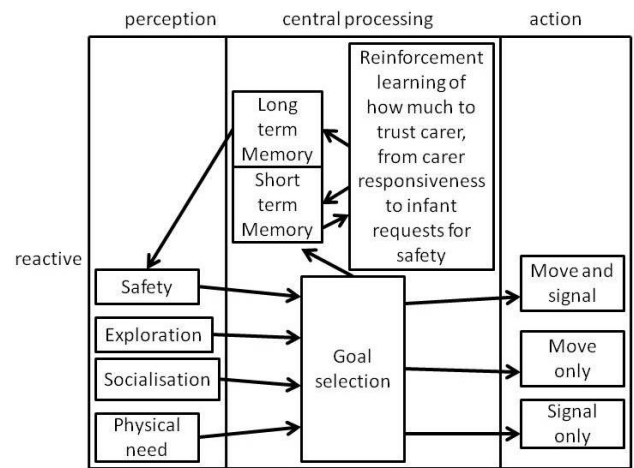


Fig. 3. A reactive attachment architecture that includes mechanisms for learning a trust level from repeated carer responses to infant requests for proximity and communication. A ‘winner-take-all’ selection mechanism means that only one goal is selected, and when the goal of safety is activated the short term memory encodes when the request for safety is made, and so how long the carer agent takes to respond to this request. A reinforcement learning subsystem then uses this information and updates trust levels in long term memory.

### D. An advanced secure-base scenario - simulating the development of individual differences in infant-carer attachment dyads

The basic secure base scenario described moment to moment interactions between carer and infant agents and the infant agent’s balancing of exploration and security maintenance over these short timescales. Scenarios which describe longer term ontogenetic development of individual difference attachment patterns need to include all these shorter duration scenarios and also describe how infant agents adapt over ontogeny to adopt either secure or insecure interaction patterns with caregiving agents; and for infant and carer agents to reproduce the different home and laboratory behaviour patterns observed in the year long home experiences and short test experiences of the Strange Situation Experiment [4]. This includes mini-scenarios that set out requirements for how separation and reunion behaviours should result from learning about past experiences over ontogeny.

### E. A reactive architectural model of the development of individual differences in infant-carer attachment dyads

The infant agent architecture created to fulfil the requirements of the advanced secure-base scenarios includes reactive level learning mechanisms that allow the infant agents to infer a measure of implicit trust from the responsiveness and sensitivity of their carer agents from the results of very many episodes where the infant agent signals for a response from the carer agent [2, p. 79-102][18], [19]. So there are repeated instances where the carer agent goes beyond the infant’s Safe-range limit, is called back by the infant agent, and then responds promptly or otherwise. Figure 3 illustrates how a reinforcement learning subsystem is connected to a more basic non-learning architecture, and allows infant agent trust in the carer agent to adapt to the carer agent’s responsiveness. In a formal description of the reinforcement learning mechanism,

the reward that the infant agent picks up at time  $t$  is defined as  $r_t$  [20]. This is a measure of immediate reward. Long term return ( $\underline{r}_t$ ) is what the infant agent is attempting to maximise by its behaviour. In reinforcement learning systems, return is often defined as a discounted sum of the future rewards [20]. 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:

$$\underline{r}_t = \sum_{k=0}^{\infty} g_k^- r_{t+k}^- + g_k^+ r_{t+k}^+ \quad (1)$$

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. 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 in all contexts. 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:

$$g_k^- = \rho, \forall k \quad (2)$$

So the negative discount function treats each time slice that the carer does not respond is a small punishment of the same amplitude. So for any given timeslice, if the infant agent has previously signalled, it will experience the same ‘disappointment’. The positive reward function is described by:

$$g_k^+ = \frac{\alpha}{1 + e^{\beta(k-(k^m))}} \quad (3)$$

This function combines desirable properties of geometric and finite horizon discount schemes. It overcomes the problem of steep initial discounting in geometric discounting, which is psychologically implausible in the attachment domain, whilst also having a smoother step function than a traditional finite horizon model. The constant  $\alpha$  controls the maximum rewards that the infant receives. The constant  $\beta$  sets the gradient of the decrease of the positive reward, a small positive value for  $\beta$  produces a gradual decrease in the size of reward each time step, higher positive values for  $\beta$  produce decreases in reward that approach a step function from maximum to minimum reward over small periods of time. So if the  $\beta$  constant is set very high then this discount scheme tends towards the sharp step function given by a simple finite horizon discount scheme. 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. Since the rewards and punishments are all set to have a value of one, equation 1 can be rearranged:

$$\underline{r}_t = \sum_{k=0}^{\infty} g_k^+ - g_k^- \quad (4)$$

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

$$\underline{r}_t = \frac{\alpha}{1 + e^{\beta(k-(k^m))}} - \rho k \quad (5)$$

This reinforcement learning function therefore means that 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.

Figure 4 illustrates the results of a computational experiment where ten identical infant agents (all initialised with the same level of trust) paired with carer agents with identical responsiveness and sensitivity. The figure shows that over time these averagely trusting infants with their averagely responsive carers diverge into two classes of infant-carer dyads (one group of five dyads with high trust and one group with low trust) because of positive feedback that operates on small differences in the initial conditions in random location of agents and toy and food objects. 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, and so makes novel contribution to the taxonomicity debate in attachment categorisation [2, p. 97]. A carer whose performance is initially intermediate between Secure and Insecure may come to be perceived as at either extreme of caregiving. Figure 4 shows an asymmetry between the shape of the positive and negative updates to the Safe-range limit because two different novel discount rules for positive and negative immediate rewards are implemented.

#### F. Future work on integrating scenarios and architectures

Future work includes integrating the hybrid and reinforcement learning architectures into a single architecture and varying parameters and functions within this architecture to investigate how the overall dynamic and emergent properties relies upon particular parameters and reinforcement learning discount functions. The ontogenetic developmental trajectories reported here will also be nested with phylogenetic simulations, to assess how different parameters and functions impact on evolutionary fitness in environments similar to the likely characteristics of the human EEA.

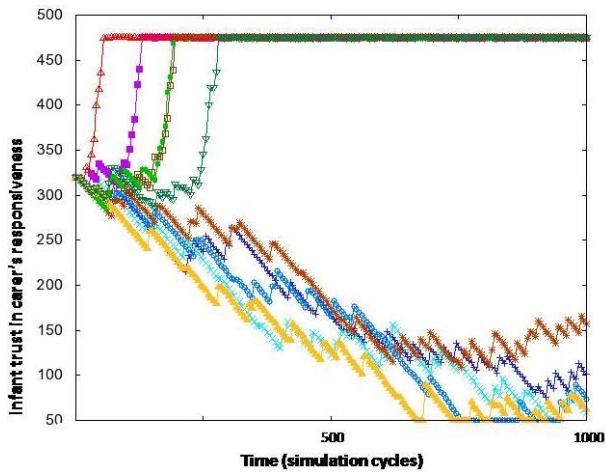


Fig. 4. Evidence that positive feedback can drive taxon formation in infant security classification. All ten infant and carer agents were initialised with identical internal parameters. Small random differences in locations were acted upon by positive feedback so that 5 infant-carer dyads developed secure relationships, and 5 infant-carer dyads developed insecure relationships.

#### IV. TAKING STOCK AND BRINGING THINGS TOGETHER: SITUATING ATTACHMENT MODELLING COMPONENTS WITHIN THE COGAFF SCHEMA

So far this paper has argued that a comprehensive understanding of attachment phenomena needs to include explanations that range from moment to moment interactions to longer term ontogenetic and phylogenetic development. Being broad minded enough to include such disparate perspectives can aid integration, as it allows seemingly dissimilar phenomena to be interrelated. This paper has also argued that secure-base behaviour is a particularly important basic phenomenon within Attachment Theory that can be simulated by a variety of agent-based models. With multiple models of the same phenomena, some questions arise: how should these diverse simulations be compared with each other? Do particular modelling approaches complement each other or compete with each other? So a further contribution of this paper is to argue that situating attachment modelling components of different models within the conceptual framework provided by the CogAff schema aid such comparisons.

The CogAff schema (figure 5) is a systematic architectural framework which is a useful conceptual tool because it makes some high level distinctions that are helpful in situating attachment models within a shared integrative framework [2], [21], [22]. The CogAff schema organises information processing by overlaying three layers (reactive; deliberative; and meta-management) and three columns (perception, central processing and action). The reactive layer of control and cognition was produced earlier in phylogenetic development and emerges earlier in ontogenetic development, but remains active throughout the lifespan. The architectures in figures ??, 2 and 3, which were designed to simulate secure-base behaviour, have both been presented in the ‘CogAff’ format. Attachment related emotions, such as being startled, terrified or delighted, arise as global interrupts in the lower reactive layer of cognitive architectures. Interruptions to deliberative processing in the middle layer gives rise to attachment emo-

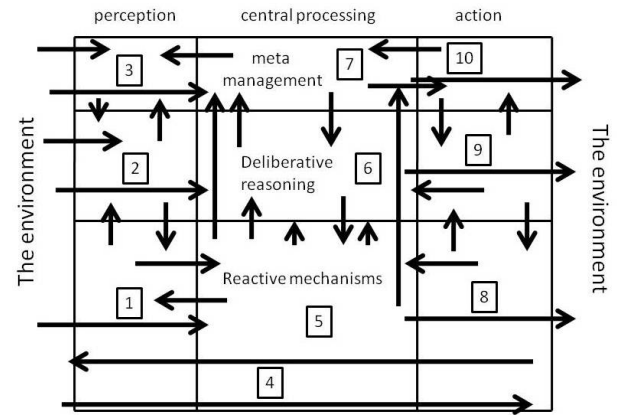


Fig. 5. A representation of the CogAff schema with additional numbered references to attachment specific processes. There are three kinds of perceptual process in the attachment domain: [1] is perception of physical distances, gaze direction and tone of voice; [2] is perception of linguistic statements such as attachment related narratives; and [3] represents the perception of a whole attachment model from another person, perhaps a mother, adult partner or therapist. Process [4] represents direct interaction between perception and action without central processing, such as when infants ‘sink in’ (which is not very naturally represented in the 2D CogAff diagram!). There are three kinds of emotional interrupt distinguished: [5] represents alarms from changes in the reactive layer, such as fear or wariness; [6] represents alarms from predicting the future, such as apprehension a carer may leave again after returning from a previous separation; and [7] represents tertiary emotions which are not alarms but perturbations to normal meta-management, deliberative and reactive processes, perhaps triggered by repeated references to an attachment figure being activated for deliberation. In symmetry with the perception column, there are also three levels of action [8, 9, 10], which deal with comparable information transfers, like movements, linguistic actions, and sharing of attachment models. Processes which transfer information in the three layers differ in their representational requirements: in the reactive layer information processes may lack compositional semantics, in the deliberative layer this is required, and in metamanagement layer information is represented with meta-semantics.

tions like being anxious, apprehensive or relieved. Modelling emotions like infatuated love and grief, that involve a loss of control, can be accomplished by first showing how information processing architectures maintain control in the highest meta-management layer in typical unemotional circumstances. Then emotional states where control is lost can be modelled as perturbances to meta-management processes which result in disruptions to intentional control [23], [21], [24]. So such perturbant emotions involve uncontrolled thoughts, negative or positive, being activated for processing with the attention requiring deliberative layer, thus ‘pushing’ other thoughts out. Caregiving patterns that have been described as ‘enmeshed’, ‘preoccupied’ or ‘avoidant’ are likely to be linked to limitations in meta-processing and self-reflection [2]. The perceptual, central processing and action columns also help compare components of different computational attachment models. Reactive perceptual input is sense data about the physical world. Imprinting mechanisms found in non-human species like Geese may be best located as reactive perceptual mechanisms. Deliberative perceptual inputs and action outputs are structured, often in the form of natural language or other compositional representations that can be automatically parsed. So they do not require attention to be perceived but are preprocessed for later deliberative processing. Attachment story narratives and representations of counterfactual situations

are an example of deliberative level perceptions. When entire attachment models or ‘world views’ are shared between individuals these will be meta-level perceptions [25]. For example, Bowlby describes how caregivers support infants by manipulating the environment and providing information directly through language use so that “*instead of each one of us having to build his environmental and organismic models entirely for himself, he can draw on models built by others*” ([1], p 82). Systems for recording attachment experiences and influencing future attachment behaviour can be situated in reactive, deliberative or meta levels depending on whether their processing requires attention and acts over compositional semantics or involves meta-semantics.

## V. CONCLUSION

Attachment Theory is formulated with a behavioural component that is composed of a great variety of attachment behavioural descriptions and a cognitive component that is composed of diverse information processing structures and mechanisms. This means that attachment phenomena are an attractive domain for modelling - which can aim to update Bowlby’s ideas as well as provide extra detail and new directions. From the following quote, from 1969, we can imagine that Bowlby would have approved of this happening:

*“Before the days of computers a major difficulty was to imagine by what possible means the kinds of detailed instructions required for the execution of instinctive behaviour could be drawn up and stored, and then made available for use at the required time and place. Now the means whereby they come to be stored and made available are no longer wholly beyond the powers of imagination, even though the processes used appear to be far more intricate and ingenious than any that man has yet learned to employ”* ([1], p. 70)

ABM is a natural way to model attachment interactions. Many existing attachment models are used to simulate secure-base behaviour. Although this is a good foundation for an attachment simulation, just simulating secure-base behaviour leaves a model just producing the same kinds of behavioural patterns as very early photo-tropic robots. This paper argues that the increasing number of simulations and modelling approaches that attempt to go beyond the tropic approach can be compared and contrasted by being situated within the architectural framework provided by the CogAff schema.

## REFERENCES

- [1] J. Bowlby, *Attachment and loss: volume 1 attachment*. New York: Basic books, 1969.
- [2] D. Petters, “Designing agents to understand infants,” Ph.D. dissertation, School of Computer Science, The University of Birmingham, 2006, (Available online at <http://www.cs.bham.ac.uk/research/cogaff/>).
- [3] R. Hinde, “Ethology and child development,” in *Hand book of child psychology*, eds. P.H. Mussen. New York: J. Wiley and Sons, 1983, pp. 27–93.
- [4] M. Ainsworth, M. Blehar, E. Waters, and S. Wall, *Patterns of attachment: a psychological study of the strange situation*. Hillsdale, NJ: Erlbaum, 1978.
- [5] M. Schlesinger and D. Parisi, “The agent-based approach: A new direction for computational models of development,” *Developmental Review*, vol. 21, pp. 121–146, 2001.
- [6] N. Bischof, “A systems approach toward the functional connections of attachment and fear,” *Child Development*, vol. 48, no. 4, pp. 1167–1183, 1977.
- [7] L. Canamero, A. Blanchard, and J. Nadel, “Attachment bonds for human-like robots,” *International Journal of Humanoid Robotics*, vol. 3, pp. 301–320, 2006, 3.
- [8] A. Hiolle and L. Canamero, “Developing sensorimotor associations through attachment bonds,” in *Proc. 7th International Conference on Epigenetic Robotics (EpiRob 2007)*, eds. C. Prince, C. Balkenius, L. Berthouze, H. Kozima, M. Littman. Lund: Lund University Cognitive Studies, 2007, pp. 45–52.
- [9] A. Hiolle, L. Canamero, M. Davila-Ross, and K. Bard, “Eliciting caregiving behavior in dyadic human-robot attachment-like interactions,” *ACM Trans. Interact. Intell. Syst.*, vol. 2, p. 3, 2012.
- [10] D. Parisi, F. Cecconi, and A. Cerini, “Kin-directed altruism and attachment behaviour in an evolving population of neural networks,” in *Artificial societies. The computer simulation of social life*, N. Gilbert and R. Conte, Eds. London: UCL Press, 1995, pp. 238–251.
- [11] D. Parisi and S. Nolfi, “Sociality in embodied neural agents,” in *Cognition and multi-agent interaction: from cognitive modelling to social simulation*, R. Sun, Ed. Cambridge: Cambridge University Press, 2006, pp. 328–354.
- [12] J. Bullinaria, “Lifetime learning as a factor in life history evolution,” *Artificial Life*, vol. 15, pp. 389–409, 2009.
- [13] J. Anderson, “Attachment behaviour out of doors,” in *Ethological Studies of Child Behaviour*, N. Blurton-Jones, Ed. Cambridge, UK: Cambridge University Press, 1972, pp. 199–215.
- [14] D. Petters, “Simulating infant-carer relationship dynamics,” in *Proc AAAI Spring Symposium 2004: Architectures for Modeling Emotion - Cross-Disciplinary Foundations*, ser. AAAI Technical reports, no. SS-04-02, Menlo Park, CA, 2004, pp. 114–122.
- [15] B. Miessner, *Radiodynamics: The wireless control of torpedoes and other mechanisms*. New York: Van Nostrand, 1916.
- [16] W. Walter, *The Living Brain*. London: Duckworth, 1953.
- [17] D. Petters and E. Waters, “A.I., Attachment Theory, and simulating secure base behaviour: Dr. Bowlby meet the Reverend Bayes,” in *Proceedings of the International Symposium on 'AI-Inspired Biology', AISB Convention 2010*. University of Sussex, Brighton: AISB Press, 2010, pp. 51–58.
- [18] D. Petters, “Building agents to understand infant attachment behaviour,” in *Proceedings of Modelling Natural Action Selection*, J. Bryson, T. Prescott, and A. Seth, Eds. School of Science and Technology, University of Sussex, Brighton: AISB Press, 2005, pp. 158–165.
- [19] —, “Implementing a theory of attachment: A simulation of the strange situation with autonomous agents,” in *Proceedings of the Seventh International Conference on Cognitive Modelling*. Trieste: Edizioni Golaridiche, 2006, pp. 226–231.
- [20] R. Sutton and A. Barto, *Reinforcement learning: an introduction*. The MIT Press, 1998.
- [21] D. Petters, “Losing control within the HCogaff architecture,” in *From Animals to Robots and Back: reflections on hard problems in the study of cognition*, eds. J. Wyatt & D. Petters & D. Hogg. London: Springer, 2014, pp. 31–50.
- [22] A. Sloman, “Architectural requirements for human-like agents both natural and artificial. (what sorts of machines can love?),” in *Human Cognition And Social Agent Technology*, ser. Advances in Consciousness Research, K. Dautenhahn, Ed. Amsterdam: John Benjamins, 2000, pp. 163–195.
- [23] I. Wright, A. Sloman, and L. Beaudoin, “Towards a design-based analysis of emotional episodes,” *Philosophy Psychiatry and Psychology*, vol. 3, no. 2, pp. 101–126, 1996.
- [24] D. Petters, E. Waters, and A. Sloman, “Modelling Machines which can Love: From Bowlby’s Attachment Control System to Requirements for Romantic Robots,” *Emotion Researcher*, vol. 26, no. 2, pp. 5–7, 2011.
- [25] D. Petters and E. Waters, “Epistemic Actions in Attachment Relationships and the Origin of the Socially Extended Mind,” in *Proceedings of 'Re-conceptualizing Mental Illness: The View From Enactivist Philosophy and Cognitive Science', AISB Convention 2013*. University of Sussex, Brighton: AISB Press, 2013, pp. 17–23.