A New Direction for Attachment Modelling: Simulating Q Set Descriptors

Dean Petters (dean.petters@bcu.ac.uk)
Department of Psychology, 4 Cardigan Street
Birmingham, B4 7BD, UK

Abstract

Modelling infant-carer attachment relationships is an emerging field at the intersection of research in Attachment Theory and computational modelling of emotion. Existing attachment models vary from very abstract models to simulations of specific experimental protocols, such as the Strange Situation Procedure. This paper argues for the benefits in broadening attachment modelling of infants and young children to also include simulating attachment Q set descriptors. The attachment Q set (AQS) is a 90 item list of attachment related behaviors used to assess the balance between attachment and exploratory behavior in home and other naturalistic settings. The AQS descriptors provide a broader and more rounded challenge for attachment modelling than other types of systematic attachment measure because they can be observed in naturalistic contexts and are less dependent on the specific details of laboratory settings. A computational attachment model is presented from which a selection of 8 attachment Q set descriptors will be simulated. Thsee initial descriptors to be simulated are concerned with the time an infant takes to recover from anxiety. A 'route map' for progress towards capturing all 90 Q sort descriptors is discussed.

Keywords: Attachment Theory; Attachment Modelling; Agent-based modelling; Attachment Q sort

Introduction

Attachment Theory describes and explains the nature of emotional bonds which form in close relationships (Cassidy & Shaver,2016). There are a small but growing number of computational attachment models which have been implemented as software and robotic simulation. For recent reviews see (Petters & Waters,2015) and (Petters & Beaudoin,2017). This paper will illustrate how empirical data in the form of attachment Q set (AQS) descriptors is well suited for the purposes of forming test-cases in scenarios and specification of requirements for attachment models. Two key contributions of this paper are that (i) it illustrates different ways that empirical data that can be used for modelling affective phenomena, and in particular it highlights the constraints and biases for simulations in this domain; and (ii) it provides an examples of how an existing simulation has been adapted to model Q set descriptors.

A short introduction to Attachment Theory

In its early theoretical development, an idea which was important in distinguishing Attachment Theory from learning theory is that attachment between an infant and main caregiver is a rich 'love' relationship (Bowlby,1969). This means that whilst attachment relationships can be tracked by observable behaviour patterns, attachment arises from a complex internal information processing architecture, termed by Bowlby the 'attachment control system' (ACS) (Bowlby,1969). The ACS acts to maintain a balance between attachment behaviour and exploration. Cues to danger momentarily move this balance from exploration to attachment. Over longer ontogenetic timescales, the complex organisation of attachment behaviour is sensitive to environmental factors. This means that both normative routines and individual difference patterns of attachment are learnt, with the aid perhaps of some evolutionary biases in infants' learning abilities. Individual differences in attachment are conceptualised as differences in an individual's ability to use their attachment figure as a secure-base. This means that the attachment-exploration balance for any individual reflects its past history of sensitive and effective responses by its caregiver in support of exploration and when the infant is distressed. (Petters,2006a).

Initially in ontogenetic development, the ACS is composed of relatively simple mechanisms, such as reflexes and fixed action patterns. However, later in development the ACS becomes comprised of a diverse range of information processing structures and mechanisms, from simple reflexes to goal corrected mechanisms and processes of planning, deliberation about future consequences of possible actions, and representing aspects of the self and environment in natural language to facilitate these processes and to communicate with others (Bowlby,1969;Petters,2006a). In addition to better capturing the behavioral complexity and underlying processes in play during infant-mother interactions, viewing attachment in control system terms clarifies assessment criteria. In principle, it is much easier to evaluate whether an attachment system is tracking set goals such as maintaining access to the carer or regulating affect than to evaluate the "qualities" of attachment as a dyadic relationship or a social network. (Waters & Deane,1985).

Whilst Bowlby set out the details of the ACS, Ainsworth and co-workers initiated the 'individual difference phase' of attachment research by developing the Strange Situation Procedure (SSP) (Ainsworth, Blehar, Waters, & Wall,1978). The SSP is conducted in a 4m x 4m room with chairs for two adults, toys for the infant to explore, and one-way glass for observation and video recording. The assessment is divided into 8 three minute episodes. At two critical points, the carer leaves the infant in the room for three minutes (once with a responsive but unfamiliar adult and once all alone). In early research, it was thought that response to these separation episodes (esp. crying) would be the best predictor of prior experience with the carer and of later adjustment. However, smooth adaptive responses to reunion (as opposed to anger or avoidance) soon proved to be much more revealing of home environment. The context changes that occur in the transitions between the eight episodes, and the infant's responses to these transitions provide a valuable data-set for contemporary researchers interested in designing attachment behaviour simulations (Petters,2006a,2006b).

Normative behaviour patterns across episodes highlight the infant's sensitivity to context that would be difficult to explain in terms of traits or operant control and justify the use of a control systems approach (e.g., more play, different kinds of signalling, less proximity seeking when carer is present) . Although the SSP assesses rather complex behavior, it does so in a restricted context and time frame. Therefore, it has been important to validate SSP based assessments against observations in more naturalistic settings and over longer time intervals. Ainsworth undertook this using detailed ethological observations, For each dyad, infant and maternal behaviour observed in home for up to 16 hours toward the end of the infant's first year (Ainsworth et al.,1978). The creation of the SSP triggered the development of a huge number of diverse measurement tools in attachment research, ranging from trait questionnaire measures similar to those used in personality research to the AQS methodology, which can be compared with the ethograms used in ethological research. More recently, Waters and Deane developed a more economical method for observing and quantifying infant-mother interactions. Their AQS descriptors cover the full range of attachment and exploratory behaviors that Ainworth recorded. However, rather than generating narrative records of the observations, the items are scored and compared to a template that describes skillful, well-organized use of the carer as a secure base (Waters & Deane,1985).

Different ways to model attachment

Attachment phenomena have been modelled in a very abstract fashion using Artificial Neural Nets (ANNs) (Fraley,2007;Edalat & Mancinelli,2013). In these attachment models the ANN can be viewed as an extremely abstract representation of an individual. The 'experiences' and 'behaviour' of the individuals in these simulations are also extremely abstract, being constituted of data that are an independent sequence of discrete training exemplar and response pairs. The main result (finding) of these simulations matches the high level of abstraction that these models have been created at. This is that in these artificial neural network simulations early prototypes are not over-written, and so show greater continuity, when new relationship experiences are inconsistent. But consistent presentation of new prototypes does result in gradual change (Fraley,2007;Edalat & Mancinelli,2013).

Agent-based models have also simulated the SSP (Petters & Waters,2015) and infant secure-base behaviour. These models are less abstract then the models based on neural nets. The main result (finding) from these simulations is that within a design space for attachment architectures, some attachment architectures show system properties like sensitivity to initial conditions (c.f. the butterfly effect) and saddle points in developmental trajectories (Petters,2006a,2006b). So where the neural networks learn item by item in 'batch jobs', and provide a result in terms of how many new learning experiences it takes to undo existing learning, the agent-based models exist in online dynamically changing virtual environments and provide results consonant with this type of dynamic simulation. In the agent based modelling case, inputs to an agent at any given time are contingent on what occurred the moment before. This means that these simulations help explain findings in terms of repeated contingent interactions that result in positive feedback driving the system away from its starting conditions towards extreme levels of 'secure' or 'insecure' interactions.

In summary, whereas the ANN results describe change in an internal representation acted upon by an independent sequence of 'offline' discrete training exemplars, agent-based modelling (ABM) results follow the changing trajectory for an agent in a broader system as that agent is acted upon and in turn influences the broader system in 'online' fashion. These findings illustrate a key principle in the art and science of cognitive modelling is the importance in finding the right level of abstraction for a simulation. This paper is concerned with discussing the benefits and drawbacks for attachment modelling in taking various approaches to deciding upon an abstraction level for computational attachment models. The paper introduces the AQS as a source of empirical constraints and requirements specifications not used before in attachment modelling. It will illustrate how modelling AQS data will provide some specific benefits over simulating other sources of information in the form of trait measures, frequency and time sampling data, and the SSP.

The nature of empirical data constrains the nature of the simulation

The importance of structural fidelity

Gaining structural fidelity is an important objective when constructing psychological measurement tools, such as personality scales and related questionnaires (Simms & Watson,2010). This is because any behavioural measure should provide data congruent with the type of construct it is designed to assess. There are two aspects to structural fidelity (Simms & Watson,2010). The first is a structural component of construct validity which requires that structural relations between the chosen test items in the measurement tool parallel structural relations for other manifestations of the construct in question, which did not get chosen to be test items. So this is a requirement that test items are representative of the possible manifestations of the construct in terms of their structural relations (Simms & Watson,2010). This aspect of test item choice is clearly relevant to the computational modeller. To produce a model based on the underlying phenomenon rather than arbitrary aspects of observed data a modeller should not abstract and simulate test items that systematically differ from other manifestations of the construct they intend to model. The second aspect of structural fidelity regards the assumptions underlying the chosen test set matching the theoretical model underlying the construct (Simms & Watson,2010). Loevinger (1957 cited in (Clark & Watson,1995)) was the first researcher to highlight these issues, and contrasted scales and tools which were based on a *"deeper knowledge of psychological theory"* (Loevinger 1957, p. 641, cited in (Clark & Watson,1995)) with tools based on an atheoretic "answer-based" technology (Clark & Watson,1995). Clearly, for the computational modeller this issue is critical. When possible, computational modellers should draw upon empirical data that align with appropriate underlying theory.

Limitations of trait rating, frequency counts and time sampling behavioural measures for attachment modelling

Trait measures are flexible and economical, take context into account, and demonstrate coherence over time (Waters & Deane,1985). However, they are not suited to assessing non-quantitative data and they score low in structural fidelity because attachment is not a trait. Waters and Deane note: *"trait language should only be used to summarise behaviour - never as a substitute; never as an explanation"* ((Waters & Deane,1985), p.44). Waters and Deane suggest that trait rating are coercive and conservative in forcing researchers to view constructs in terms of pre-existing scales, and working against introduction of new scales or measures during the process of measurement (Waters & Deane,1985). It is also difficult to disentangle affect and cognition in trait rating data (Waters & Deane,1985). Observational data in the form of frequency counts and time sampling retain much more behavioural detail than trait rating methods. They also have good structural fidelity. However, their expense and difficulty to use mean that often only small numbers of behavioural categories are assessed in a single study. In addition, particularly interesting behaviours may occur at very low frequencies (Waters & Deane,1985). This means that they are a very good way of getting very detailed data on behaviours of specific interest if those behaviours occur relatively frequently. For practical reasons of resources and time, what these methods is not so good for is gaining a comprehensive overview of an entire behaviour domain that possessed many different salient behaviour types (Waters & Deane,1985).

Limitations of the SSP as a source for attachment modelling

The SSP (Ainsworth et al.,1978) involves a set of scoring protocols that includes behaviour coding, frequency and percentage measures. All these measures were developed to provide insight in to the underlying ACS. This is done partly by including separation and reunion episodes which are mildly to moderately stressful as a way of activating the ACS in a controlled manner. Because the procedure is designed specifically to uncover the state of an infant's ACS, the SSP affords very high structural fidelity. However, the behaviours produced in the SSP do not correlate directly to behaviours in naturalistic settings like the home environment. For example, crying rate in the SSP does not predict the rate of these behaviours at home. Rather, the behaviours produced in the SSP are used to infer the state of the ACS, and an ACS in this state will produce different behaviour patterns depending on context. Other limitations of the SSP for psychological research include the narrow age range it can be used (21-18 months), strong carry over effects (infants recognise the context if it is repeated soon after), the expense and difficulty of administration, and it does not capture developmental change well (Waters & Deane,1985). For the attachment modeller, it is also too narrow in measures used and number of contexts it describes.

Limitations of modelling in a 'method-bound' research domain

A further limitation for attachment modelling related to attachment measures arising from the large variety of attachment measures currently available. It might be imagined that having numerous attachment measures to choose from would help the attachment modeller. However, the current situation has given rise to what Fonagy terms the 'method-bound' nature of Attachment Theory:

> *"Attachment theory* [...] *has been in some ways method-bound over the past 15 years. Its scope was determined less by what fell within the domain defined by relationship phenomena involving a caretaking-dependent dyad and more by the range of groups and behaviors to which the preferred mode of observation, the strange situation, the adult attachment interview, and so forth, could be productively applied."*((Fonagy,1999), p. 5)

There is therefore drawback in attempting to model behaviour in a domain which is 'method-bound'. If the methods leave gaps in empirical data coverage, the gaps will not get modelled. So a researcher interested in behaviour will need to consider carefully how to get a representative sample of behaviour in this kind of domain. The next section describes the AQS, which overcomes the limitations for attachment modelling of behaviours described as traits, frequency counts, time samples and SSP patterns. It also allows strong structural fidelity between observable behaviours and internal processes and provides more comprehensive coverage of attachment any other single measure.

Overview of Attachment Q Set Behavioural Descriptors

Q sort methodology can be applied to research in any given area of behavioural science (Waters & Deane,1985). First, it involves developing a set of descriptive items. These should ideally be extensive enough to be an overview of the entire behavioural domain of interest. For example, Waters and Deane spent two years developing a 100 item Q set for infant attachment. They reviewed relevant literature; developed a list of relevant constructs (security, dependency, detachment, self-efficacy, aspects of object orientation, communication skills, predominant mood, response to physical comforting, fearfulness, anger and trust); rated infants and toddlers on these variables and then specified the behaviour that led to or was congruent with these ratings (Waters & Deane,1985). This is important because it provides an emphasis on simulating ordinary as opposed to traumatic experiences when attempting to model the development of information processing architectures for attachment. When 'ordinary' architecture development can be modelled, trauma modelling can follow.

Each AQS item refers to a particular behaviour patterns in a specific context. As Waters and Deane note, because the AQS *"covers a broad range of secure-base and exploratory behavior, affective response, social-referencing and other aspects of social cognition* [...] *it can be construed as an overview of the entire domain of attachment relevant behavior, as currently understood within an ethological/control systems perspective"* ((Waters & Deane,1985), p. 7. This means that the AQS captures a more comprehensive description of attachment relevant behaviour that other behavioural measures that might be used in computational modelling.

What is of particular interest for attachment modelling is how the AQS descriptors were initially constructed. Waters and Deane describe four stages in the initial development of the attachment Q-set (Waters & Deane,1985). These four stages involve procedures for developing sets of items which empirical psychologists use when observing behaviour, processing and analysing behaviour, and ultimately producing a classification for the individual observed. However, computational modellers can use these descriptions to construct representative behavioural scenarios from which to direct model design and simulation implementation, and guide model evaluation and validation.

The first stage of Q-set production is of most interest to computational modellers because it involves procedures for developing sets of items. Developing a Q set requires careful examination of extensive observational data. Even when initial descriptor sets are produced they need to be trialled to weed out highly correlated descriptor pairs (Waters & Deane,1985). It also requires close attention to detail, focusing on distinctions and ambiguities that may not be apparent in measurement tools at a higher abstraction level (Waters & Deane,1985). One of the major advantages in this methodology for empirical psychologists is that observers new to the domain will evaluate the same context as the experts who designed and calibrated the AQS (Waters & Deane,1985). This is precisely the property of a measurement tool that computational modellers require: providing broad and comprehensive coverage but also focused on behaviour of interest, filtering out irrelevant behaviours from analysis, and a level of clarity in actions and context that a novice can understand (and learn from). Another major advantage of the AQS methodology is it gives an helpfully strong focus on the role of context, and effectively defines behaviours as *"acts plus context"* as context is integral to each Q set item (Vaughn, Waters and Teti,forthcoming).

The second stage of a Q-sort methodology involves then assigning scores to descriptors when assessing individual study participants, depending how well the participants matches the behaviour. Then the third stage of a Q-sort methodology involves data reduction and analysis and there are a wide variety of procedures for doing this (Waters & Deane,1985). The Q-set methodology allows an infant or child's behaviour to be observed and measured so that it gives a set of scores which can be correlated against a hypothetical 'most secure baby' Q-sort. So a very secure Q set would give a high correlation (around r = 0.6 of a theoretical maximum of 1). Very insecure infants give correlations around r = 0, because insecure behaviour does not involve doing the exact opposite of secure behaviour. Details of how the Attachment Q sort procedure is actually used in empirical research by psychologists to assess infants is beyond the scope of this attachment modelling paper but described in more detail by Waters and Deane (Waters & Deane,1985), with the full set of Q sort items listed by Waters (Waters,1987).

Unlike other measurement tools, the AQS provides an abstract generalised template for computational attachment modellers. As Vaughn, Water and Teti note, it is similar to an 'ethogram', because it is *"rooted in observation and attempts to catalogue the full suite of behaviors associated with a particular behavioral system"*.(Vaughn, Waters and Teti,forthcoming, p.14).

Modelling results

Modelling of Q set descriptors has been undertaken using an existing agent-based model of the SSP as a point of departure (Petters,2006a,2006b). Figure 1 shows a hybrid infant architecture with reactive components and a simple deliberative subsystem. This architecture simulates the SSP by 'experiencing' the pattern of caregiving in a home 'training' stage and then producing typical SSP behavioural patterns in a 'test' stage. It has been used as the basis for implementing AQS descriptors by being augmented with further perceptual, memory and action mechanisms.

The most recent version of the AQS has 90 behaviour descriptors (Waters,1987). Waters and Deane present these items as a single list. The first task that has been undertaken in this current research is to analyse these descriptors to assess the best order to place them in a 'route map' for eventually capturing all Q sort descriptors in a single implemented simulation. So for this current modelling effort, the AQS descriptor list has been analysed into three main sets of descriptors: those that could be modelled by the existing agent-based architecture with manageable extensions to that architecture (20 items); those that were well beyond the capabilities of the existing implemented agent architecture and would require a significantly more sophisticated architecture to be simulated (35 items); and 'filler' items not linked to attachment phenomena and which were added to the Q set for pragmatic reasons to make the AQS sorting procedure run more smoothly (35 items) (Waters & Deane,1985). There were two main reasons that items were assessed as being significantly beyond the capabilities of the existing simulation: that the descriptor

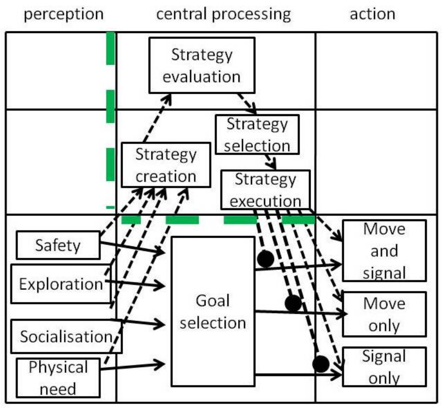

Figure 1: A hybrid attachment architecture with reactive, deliberative and meta-management subsystems. This archiecture has been extended to store in infant agent memory the details of episodes when it interacts with carer agents, with the responsiveness and sensitivity of the interaction stored along with other details of context

required a more sophisticated perception and understanding of others than is currently implemented in existing attachment simulations (for example, AQS descriptor 42: *'Child recognizes when mother is upset. Becomes quiet or upset himself. Tries to comfort her. Asks what is wrong, etc.'*); and that the descriptor requires a more complicated model of the simulated world than is currently implemented (for example, AQS descriptor 53: *'Child puts his arms around mother or puts his hand on her shoulder when she picks him up'*. Filler items included AQS descriptors such as number 89: *'Childs facial expressions are strong and clear when he is playing with something'*.

The 20 items in the set of AQS descriptors which were assessed as being able to implemented with an extension of the existing agent-based model have been categorised into 7 subsets focusing upon: affective communication (2 items); predisposition to cry or be demanding (2 items); the interplay of exploration, anxiety and relief (2 items); aspects of physical need (3 items); how sensation and perception operate in the attachment domain (2 items); time to become anxious (2 items); and time to recover from anxiety (8 items). Figure 2 presents the eight AQS items concerned with the subset of descriptors concerned with 'time to recover from anxiety'. The initial modelling in this AQS simulation project has concentrated on capturing these eight descriptors. This has been done by implemented extra percptual, memory and action mechanisms to support simulation of infant expectations about the immediate future likely responses of the carer agent.

The existing agent-based model of the SSP (in figure 1) already simulates individual differences in the behavioural patterns that result when infant agents return to the proximity of their carer agent after a separation (Petters,2006a,2006b). This occurs because the existing agent-based model possesses 'behaviours' for attachment proximity, exploration, social need, and physical need. These all operate independently and in parallel in proposing new active action goals for the agent. The action selection mechanism is a 'winner-take-all' mechanism which selects the candidate goal with the highest activation. The 'behaviour' subsystem for attachment anxiety goal is activated when the distance between the infant agent and carer agent is beyond a parameter termed the 'safe-range'. This safe-range parameter is learned from the results of all previous episodes when the infant agent has attachment anxiety as its active 'behaviour' goal. If the infant agent has experienced a history of prompt and sensitive responses from its carer agent it will have a large 'safe-range'. This means that the carer agent can move further away before the infant agent's attachment anxiety 'behaviour' goal starts to become activated. If the infant agent has experienced a history of tardy and insensitive responses to its requests for proximity and attention then it will have a small 'safe-range'. This not only means that attachment anxiety will be experienced more often, but that anxiety will take longer to drop back to a normal value when reunions occur. However, the 'safe-range' parameter is a very economical record of previous interactions because the results of the quality of interaction in all different contexts are collapsed into a single numerical value. What the newly implemented 'AQS' extensions to the existing simulations involve is the recording of much more context for each individual episode where attachment anxiety becomes the active goal and the infant agent records carer agent responsiveness and sensitivity. In the 'AQS' extension architecture when the infant agent experiences an episode of active attachment anxiety and signals and moves to reduce its anxiety level the context at initiation and conclusion of the goal is recorded. This context includes external measures, such as the agents and objects present in sensory data, and also internal context, such as relative activations for inactive goals, such as physical and social need. This means that when a new episode of anxiety is experienced this more detailed and specific 'episodic memory' is available to influence responses in a 'recovering from anxiety' time period. This mechanism will therefore support simulating the expectations of carer response apparent in the AQS 'time to recover from anxiety' descriptor subset. The production of the AQS simulation extension is a work-in-progress with the aim of ultimately capturing all 90 AQS behavioural descriptors. Mechanisms for encoding very simple episodic memories for anxiety episodes have been implemented. Detailed mini-simulations of AQS descriptor items 2, 13 and 33 from the 'time to recover from anxiety' subset have been completed with progress ongoing for the AQS descriptor items 34, 43, 70, 71, and 78.

Conclusion

Evaluation and validation of attachment models is less well defined than the quantitative evaluation and validation which can occur with some cognitive models that involve simulating quantitative data like reaction times or accuracy measures.

AQS number	Descriptor
2	When child returns to mother after playing, he is sometimes fussy for no clear reason
13	When the child is upset by mother's leaving, he continues to cry or even gets angry after she is gone.
33	Child sometimes signals mother (or gives the impression) that he wants to be put down, and then fusses or wants to be picked right back up.
34	When child is upset about mother leaving him, he sits right where he is and cries. Doesn't go after her.
43	Child stays closer to mother or returns to her more often than the simple task of keeping track of her requires.
70	Child quickly greets his mother with a big smile when she enters the room. (Shows her a toy, gestures, or says "Hi, Mommy").
71	If held in mother's arms, child stops crying and quickly recovers after being frightened or upset.
78	When something upsets the child, he stays where he is and cries.

Figure 2: A set of eight AQS descriptors related to descriptions of the time an infant takes to recover from anxiety have been grouped together to act as a starting point for the AQS modelling project. (AQS descriptor numbers relate to the ordering given in (Waters,1987)).

This paper has demonstrated the benefit of using AQS descriptors in attachment modelling because of their structural fidelity, comprehensive coverage as attachment 'ethograms', and ready incorporation in modelling scenarios. In comparison to the AQS item pool, past modelling research has focused on a narrower, and arguably less theoretically interesting range of behaviors and processes. Thus, this paper has examined the AQS item pool with an eye toward identifying content that could be incorporated into existing models and architectures. It has also highlighted content that seems too complex to be easily incorporated and suggested some of the problems that would have to be solved before doing so. New mechanisms have been described to simulate how infants retreat to the caregiver when distressed, and establish and maintain contact until comfortable enough to resume exploration. The next stage is to simulate how infants explore away from the caregiver, evaluate and maintain caregiver access and availability, and seek information or assistance while exploring or manipulating objects or locations.

References

Ainsworth, M., Blehar, M., Waters, E., & Wall, S. (1978). *Patterns of attachment: A psychological study of the strange situation*. Hillsdale, NJ: Erlbaum.

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

Cassidy, J., & Shaver, P. (2016). Handbook of Attachment, 3rd Ed. London: Guilford Press.

Clark, L., & Watson, D. (1995). Constructing validity: Basic issues in objective scale development. *Psychological Assessment*, *7*, 309–319.

Edalat, A., & Mancinelli, F. (2013). Strong attractors of Hopfield neural networks to model attachment types and behavioural patterns. IEEE. Available from http://dx.doi.org/10.1109/IJCNN.2013.6706924

Fonagy, P. (1999). Points of contact and divergence between psychoanalytic and attachment theories: Is psychoanalytic theory truly different. *Psychoanalytic Inquiry*, *19:4*, 448-480.

Fraley, R. (2007). A connectionist approach to the organization and continuity of working models of attachment. *Personality and Social Psychology Review*, *6*, 1157–80.

Petters, D. (2006a). *Designing agents to understand infants*. Unpublished doctoral dissertation, School of Computer Science, The University of Birmingham. ((Available online at http://www.cs.bham.ac.uk/research/cogaff/))

Petters, D. (2006b). Implementing a theory of attachment: A simulation of the strange situation with autonomous agents. In *Proceedings of the seventh international conference on cognitive modelling* (pp. 226–231). Trieste: Edizioni Golardiche.

Petters, D., & Beaudoin, L. (2017). Attachment Modelling: From Observations to scenarios to designs. In P. Erdi, B. Bhattacharya, & A. Cochran (Eds.), *Computational neurology and psychiatry: Volume 6 of springer series in bio-/neuroinformatics.* (pp. 227–271).

Petters, D., & Waters, E. (2015). Modelling Emotional Attachment: An Integrative Framework for Architectures and Scenarios. In *Proceedings of IJCNN.* IEEE.

Simms, L., & Watson, D. (2010). The construct validation approach to personality scale construction. In Handbook of Research Methods in Personality Psychology*, eds. r.w. robins, r.c. fraley & r.f. krueger* (pp. 240–258). London: Guilford Press.

Waters, E. (1987). Attachment q-set (version 3) [Computer software manual]. Available from http://www.psychology.sunysb.edu/attachment/measures/content/aqs_items.pdf (accessed February 2, 2017)

Waters, E., & Deane, K. (1985). 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.