

A Parallel Distributed Processing Model of Accessibility of Attachment Knowledge

Roxanne L. Thrush (rthrush@andrew.cmu.edu)

Department of Psychology & Center for the Neural Basis of Cognition, Carnegie Mellon University, 5000 Forbes Avenue
Pittsburgh, PA 15217 USA

David C. Plaut (plaut@cmu.edu)

Department of Psychology & Center for the Neural Basis of Cognition, Carnegie Mellon University, 5000 Forbes Avenue
Pittsburgh, PA 15217 USA

Abstract

Attachment theory is a prominent social psychological framework for understanding patterns of thinking about relationships. According to this theory, individuals gradually develop mental models of relationships, and these models influence subsequent information processing. Different experiences in relationships are posited to lead to differences in accessibility of particular mental models. The current work proposes that experience leads instead to knowledge about general tendencies in relationships, and that this knowledge is sufficient to explain outcomes thought to be associated with differences in accessibility of particular mental models. Specifically, we advance a connectionist account of the acquisition of attachment knowledge.

Keywords: connectionist modeling; PDP; attachment theory

Background

Attachment theory describes the emergence of patterns of interpersonal cognition, behavior, and emotional tendencies, particularly during times of stress or adversity (e.g., Bowlby, 1969/1982). According to this theory, individuals learn what to expect from other people, given their particular history of interacting with their caregivers. Individuals whose caregivers have generally been available during times of need and who have provided sensitive care are expected to develop generalized beliefs that others are typically the types of people who are reliable, dependable, and sensitively responsive. Other individuals may experience caregivers as unavailable, insensitive, unresponsive, or critical during times of need. These individuals are then expected to develop generalized beliefs about other people that correspond to these experiences. These generalized beliefs, referred to as “working models of attachment,” are thought to be internal representations of relationships. They are often referred to as experience-dependent prototypes or schema, and are thought to guide future information processing outcomes (e.g., perception of and memory for stressful and/or relationship events, expectations of others’ responses) and emotional responses to distressing situations (e.g., Collins & Read, 1994).

Considerable empirical evidence suggests that individuals do process relationship information in an attachment schema-consistent manner: attachment biases are observed in selective attention, perception, memory (recall and reconstruction), and attributions during relationship and

stressful events (e.g., reviewed in Collins, Ford, Guichard, Ford, & Feeney, 2004).

Attachment researchers have recently posited that different experiences in relationships could lead not just to differences in *content* of relationship knowledge, but to differences in the accessibility of more positive or more negative working models of attachment (Collins et al., 2004). In other words, individuals may not differ in the type of information they possess regarding relationships, but the ease with which positive or negative beliefs or expectations could differ, given individual differences in attachment experiences. Consistent with this theoretical proposition, Mikulincer and colleagues (e.g., Mikulincer & Shaver, 2001; Mikulincer et al., 2001; Mikulincer et al., 2003) demonstrated that priming individuals with positive working models of attachment (e.g., by asking participants to recall positive personal attachment experiences, by having participants watch pictorial representations or read stories of supportive others, or by subliminally presenting participants with words characteristic of positive attachment experiences) resulted in a variety of favorable cognitive and behavioral outcomes.

The proposition that individuals differ in the ease with which particular types of mental models are accessible is limited in the sense that it is not accompanied by a mechanistic account of what it means to “access” a mental model, and how the process of accessing knowledge influences online processing. The goal of the current work is to provide a connectionist simulation of the acquisition of attachment knowledge. Specifically, we posit that a system which learns statistical similarities of its experiences in relationships, but does not actually store representations of relationships, will produce information processing biases that are consistent with empirical observations. We further suggest that a connectionist formulation of attachment learning can contribute novel perspectives on the two unresolved emergent issues in attachment described above.

Contributions of Connectionist Modeling to Attachment Theory

Connectionist models simulate cognition via the collective operation of relatively simple, biologically-inspired processing units. (e.g., Rummelhart & McClelland, 1985; McLeod, Plunkett, & Rolls, 1998). Each processing

unit has a scalar activation value, analogous to the firing rate of a neuron, and positively- or negatively-weighted connections to and from other units, analogous to the collection of synapses among neurons. The activation of a particular unit is determined by the integration of excitatory or inhibitory inputs from other processing units to which it is connected. Information is represented as a pattern of activation across processing units. Processing occurs when input to the system (as a pattern of activation across input units) is eventually transformed into a response (a pattern of activation across output units). Over the course of training, the network adjusts connections among processing units so that appropriate responses to input are produced. This adjustment in connection weights occurs as a function of the degree to which the response produced by the network to a particular input differs from the specified target response. As a result of cumulative training, the network develops a system of connection weights which facilitates processing of a collection of patterns similar to those with which the network has experience. Thus, connectionist networks learn to extract statistical similarities among and across their experiences.

Connectionist modeling is especially relevant in the current context because it can be used to derive specific predictions regarding what it means for knowledge to be “accessible.” Structure and function are intimately related in connectionist models: that is, the functioning of working models of attachment, including accessibility effects, is determined by the structure of the system that processes attachment information. Accessibility of working models of attachment can be thought of as the likelihood that a particular representation will emerge from particular input, and this likelihood is determined by the set of connection weights in a network.

To ask whether individuals differ in the extent to which a particular working model of attachment is accessible translates to asking whether a particular property of an attachment experience (e.g., “in distress”) is more strongly connected to a particular type of responsiveness (e.g., sensitive responsiveness, as coded by features such as “approving look,” “supportive touch,” etc.) for one individual than for another. A connectionist perspective on this question might be that the association of different possible features of attachment experiences depends on unique individual history: the extent to which the experience of being in distress is related to the experience of receiving sensitive care depends on how consistently these events were experienced simultaneously. It follows that the accessibility of a sensitive working model of attachment would increase in response to increasing the number of features of an input pattern that are associated with sensitive responsiveness. For individuals who have generally experienced sensitively responsive caregivers, little or no additional information would likely be needed in order to facilitate the emergence of a representation of a relationship partner as sensitively responsive. Individuals with less positive attachment histories, however, may require

additional information (e.g., specific positive behaviors of a relationship partner) in order for a sensitive representation of a relationship partner to emerge.

Simulation

Network Architecture The task of the network was to learn general tendencies among and across attachment experiences. The network was trained with approximations of attachment experiences, and allowed to discover relationships among features of the environment (i.e. whether the situation is attachment-relevant), features of the self (i.e. one’s present emotional state as well as one’s future emotional state), and features of others (i.e. who is present and the type of response that can be expected). An autoencoder network, in which the task of the network is to learn to reproduce input patterns over an output layer, was, therefore, chosen in order to model the process of learning to represent attachment information.

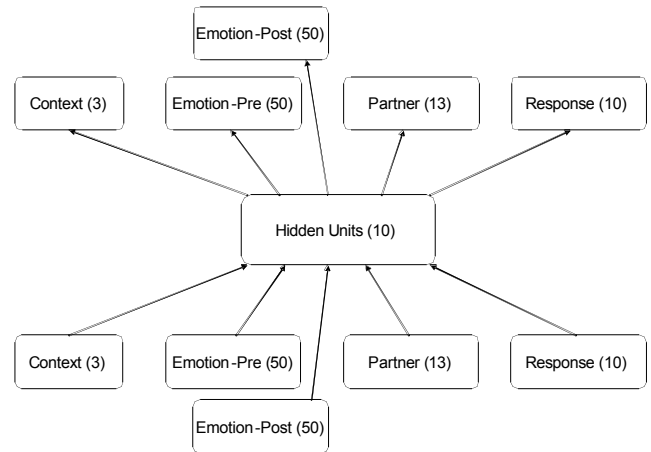


Figure 1: Network Architecture

The network has five groups of units in the input layer (Context, Emotion-Pre, Partner, Response, and Emotion-Post), one group of 10 hidden units, and five groups in the output layer, each corresponding to an input group (Context-out, Emotion-Pre-out, Partner-out, Response-out, and Emotion-Post-out). The Context groups are comprised of 3 units, the Emotion groups of 5 subgroups of 10 units each (50 total units), the Partner groups of 16 units, and the Response groups of 12 units. The hidden units receive input from each of the 5 input groups, and are unidirectionally connected to each of the 5 output groups. In addition, each hidden and output unit receives an additional “bias” connection, which can be thought of as coming from an additional unit whose activation is fixed at 1.0. Including the bias connections, the total number of connections within the network is 2,158.

Input is presented to the network by fixing the states of the input units to specified values (either 0 or 1). The activations of the hidden units are computed based on these input activations, and then the activation of the output units are computed from the hidden activations. Specifically, the

activations of the hidden and output units range between 0.0 and 1.0 and are computed according to following equations:

$$a_j = \frac{1}{1 + \exp(-n_j)}$$

where a_j is the activation of unit j , n_j is its net input, w_{ij} is the weight on the connection from unit i to unit j , and $\exp(\cdot)$ is the exponential function.

Representations The network was trained with four different sets of input and target output to simulate four different attachment histories. The training sets provided as input to the network indicate (a) the environmental context an individual might be in (Context); (b) emotion associated with that context (Emotion-Pre); (c) the identity of a relationship partner, if present (Partner); (d) the relationship partner’s responsiveness (Response); and (e) emotion associated with the relationship partner’s response (Emotion-Post). Distinct features of each type of information were coded by the activity of individual units; although these units clearly do not capture the complexity of the corresponding real-world information, their pattern of co-occurrences within and across training patterns were designed to reflect plausible alternative attachment histories. The three different attachment histories used in the simulation differed in the proportions of different types of training patterns, as discussed in the section *Training and Testing Procedures*.

Eight training patterns were created to represent eight pairs of Context and Emotion-Pre input patterns (see Table 1). The Context units correspond to situations that are “novel” (vs. familiar), “alarming” (vs. non-alarming), and “whereabouts of attachment figures (AF) unknown” (vs. known). The units of both sets of emotion groups correspond to: “calm,” “happy,” “sick,” “anxious,” “angry,” “afraid,” “curious,” “proud,” “disappointed,” and “sad”. Training patterns for the Emotion-Pre group were matched to patterns in the Context group, and the proportion of training patterns with each context/emotion-pre pair are presented in Table 1.

Ten training patterns for the Partner group correspond to ten specific individuals. The 18 units in the Partner group correspond to observable characteristics of these individuals. The specific patterns for each relationship partner are presented in Table 1. For all six sets of training patterns, 40% of the patterns in which a relationship partner was present (third context unit = 0) used the representation for “mother”; 25% used the representation for “father”; 5% used the representation for each of “grandmother”, “grandfather”, “teacher 1”, “teacher 2”, “friend 1”, and “friend 2”; and 2.5% used the representation for each of “aunt” and “uncle”.

Each Context and Emotion-Pre pair could be combined with one of three Response training patterns from the

Table 1: Context, Emotion-Pre, and Partner Training Patterns

Number of Training Patterns	Context			Emotion-Pre												
	Novel	Alarming	Whereabouts of AF Unknown	Calm	Happy	Sick	Angry	Anxious	Afraid	Curious	Proud	Disappointed	Sad			
180	0	0	0	1	1	0	0	0	0	0	0	0	0			
20	0	0	1	0	0	0	0	1	0	0	0	0	0			
36	0	1	0	0	0	0	0	0	1	0	0	0	0			
4	0	1	1	0	0	0	0	1	1	0	0	0	0			
72	1	0	0	0	1	0	0	0	0	1	0	0	0			
8	1	0	1	0	0	0	0	1	0	1	0	0	0			
72	1	1	0	0	0	0	0	0	1	1	0	0	0			
8	1	1	1	0	0	0	0	1	1	1	0	0	0			

Relationship Partner	Proportion of Experience with Relationship Partner	Partner																
		Female	Same Age as Self	Older Than Self	Much Older Than Self	Blonde Hair	Red Hair	Brown Hair	Black Hair	Grey Hair	Curly Hair	Freckles	Glasses	Beard	Mustache	Thin	Average Weight	Overweight
Mother	40%	1	0	1	0	0	1	0	0	0	1	1	1	0	0	0	1	0
Father	25%	0	0	1	0	0	0	0	1	0	0	0	0	0	1	0	1	0
Grandmother	5%	1	0	0	1	0	0	0	0	1	0	1	0	0	0	0	0	1
Grandfather	5%	0	0	0	1	0	0	0	0	1	1	0	1	1	1	0	0	1
Teacher 1	5%	1	0	1	0	0	1	0	0	0	0	0	0	0	0	1	0	0
Teacher 2	5%	0	0	1	0	0	0	1	0	0	1	1	0	1	1	0	0	1
Friend 1	5%	1	1	0	0	1	0	0	0	1	1	1	0	0	1	0	0	0
Friend 2	5%	1	1	0	0	0	1	0	0	0	1	0	0	0	1	0	0	1
Aunt	2.5%	1	0	1	0	0	1	0	0	0	1	0	1	0	0	1	0	0
Uncle	2.5%	0	0	1	0	0	0	1	0	0	0	1	0	0	0	0	0	1

relationship partner: sensitive, hostile, or indifferent. The patterns for these responses for each Context/Emotion-Pre pair are presented in Table 2. Training sets for the Emotion-Post group then depended on the Response pattern. When a sensitive response was received, the Emotion-Post was either an activation of “happy” and “calm” (contexts 000 and 010); “happy”, “curious”, and “proud” (context 100); or “happy” and “curious” (context 110). When a hostile response was received, the Emotion-Post was an activation of “anger”, “disappointment”, and any negative emotions that were active in the Emotion-Pre pattern. When an indifferent response was received, the Emotion-Post was an activation of “disappointment” and any negative emotions that were active in the Emotion-Pre pattern.

Training and Testing Procedures The patterns described above were used to train the network to anticipate outcomes of attachment situations. The network was provided with a series of patterns, each of which represented a specific attachment experience. Although variability could exist within a set of patterns (i.e., although individuals may experience others as sensitively responsive in one pattern, but hostilely responsive in another pattern), the cumulative effect of providing this series of patterns was such that the network was able to extract general tendencies regarding how features of attachment experiences were typically related. The network was tested by providing it with a pattern of input, and then observing the produced response across the output layer. Four different sets of training

patterns were used in order to simulate differences in attachment history.

The current project explored cognitive sequelae of individual differences in attachment experiences. Although it is likely that individuals with more positive interpersonal experiences during times of adversity also experience qualitatively different interactions with attachment figures during less aversive conditions, the four training patterns during these less aversive conditions did not differ. This was intended to isolate and explore the effects of the context of attachment-relevant situations. In order to avoid adding unnecessary complexity to the current model, each of the four training sets included the same patterns for experiences that were not attachment-relevant. In each set, 90% of the responses received to positive emotion (contexts 000 and 100) were sensitive; 5% were hostile; and 5% were indifferent. The responsiveness from different relationship partners did not differ in these contexts.

In order to explore consequences of individual differences in attachment-relevant experiences, the networks did, however, differ in experiences with relationship partners (and Emotion-Post) in response to negative Emotion-Pre input (contexts 010 and 110). Training patterns varied in terms of: (a) percent of sensitive vs. hostile responses received, (b) continuity of sensitive vs. hostile responses over time, and (c) similarity of responses received from different relationship partners.

Training: Individual Differences in Attachment History To explore contributions of connectionist modeling to understanding accessibility of attachment knowledge, three sets of training patterns were used. The first network, the *sensitive* network, simulated an attachment history in which attachment figures had generally been available to provide sensitive care in response to attachment needs. This network consistently received 90% sensitive responses (and

5% hostile responses and 5% indifferent responses) to negative Emotion-Pre input from each relationship partner. The *hostile* network simulated an attachment history in which attachment figures were primarily hostilely responsive to and rejecting of attachment needs throughout development. The *hostile* network was trained with 90% hostile responses (and 5% sensitive and 5% indifferent responses) to negative Emotion-Pre input. The *exception* network simulated an attachment history in which the relationship partner with whom the network had the most experience (mother; 40% of total relationship partner inputs) responded to distress in a way that differed from responses of the other nine relationship partners. This exception partner responded sensitively to experiences of distress, while all other relationship partners responded hostilely to experiences of distress. Each network was trained for 1,000 trials.

Testing Procedure Following training, each network was provided with increasing amounts of input information, and expectations of sensitive responsiveness were assessed. In this way, we examined the type and amount of input necessary in order to elicit expectations of maximal sensitive responsiveness from others and associated emotion. Networks were initially given (1) no input; (2) “alarming” Context input and “afraid” Emotion-Pre input; (3) “alarming” Context input, “afraid” Emotion-Pre input, and “mother” Partner input; (4) “alarming” Context input, “afraid” Emotion-Pre input, “mother” Partner input, and “sensitive” Response input; and (5) “alarming” Context input, “afraid” Emotion-Pre input, “mother” Partner input, and “hostile” Response input.

Expectations of sensitive responsiveness were operationalized as the similarity between the response produced by the network across Response and Emotion-Post-Out groups and the prototypical sensitive response to

Table 2: Sensitive and Hostile Response and Emotion-Post Patterns

Emotion-Pre	Sensitive Response											Emotion-Post										
	Attentive	Understanding	Reassuring	Accepting	Rejecting	Critical	Hostile	Annoyed	Proud	Disinterested	Encouraging	Assisting	Calm	Happy	Sick	Angry	Anxious	Afraid	Curious	Proud	Disappointed	Sad
Anxious	1	1	1	1	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0
Afraid	1	1	1	1	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0
Anxious & Afraid	1	1	1	1	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0
Anxious & Curious	1	1	1	1	0	0	0	0	1	0	1	1	1	1	0	0	0	0	0	1	0	0
Afraid & Curious	1	1	1	1	0	0	0	0	1	0	1	1	1	1	0	0	0	0	0	1	0	0
Anxious, Afraid, & Curious	1	1	1	1	0	0	0	0	1	0	1	1	1	1	0	0	0	0	0	1	0	0
	Hostile Response											Emotion-Post										
Anxious	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	1	0	0	0	1	0
Afraid	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	0	1	0	0	1	0
Anxious & Afraid	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	1	1	0	0	1	0
Anxious & Curious	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	1	0	0	0	1	0
Afraid & Curious	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	0	1	0	0	1	0
Anxious, Afraid, & Curious	0	0	0	0	1	1	1	1	0	0	0	0	0	0	0	1	1	1	0	0	1	0
	Indifferent Response											Emotion-Post										
Anxious	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	1	0
Afraid	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	0	1	0	0	0	1
Anxious & Afraid	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1	0	0	1	0
Anxious & Curious	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	1	0
Afraid & Curious	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	0	0	1	0	1
Anxious, Afraid, & Curious	0	0	0	0	0	0	0	0	0	1	0	0	0	0	0	0	1	1	1	0	1	0

“afraid” input (see Table 2). Similarity was calculated using the following equation:

$$S = 1 - \frac{\sum_{i=1}^n (P_i - O_i)^2}{n}$$

where S represents the similarity between the produced output and the prototypic sensitive pattern of activation; P_i represents the prototypic activation value of the i th unit in the sequence of n Response and Emotion-Post units; and O_i represents the activation value of the i th unit in the sequence of n corresponding Response-Out and Emotion-Post-Out units. In the current calculations of similarity, $n = 62$ because a total of 62 Response and Emotion-Post output units were included in the model.

Prior to training, connection weights are very small, random values, so that testing the network produces little change in activation in output units. Thus, the baseline activation in each output unit (O_i) is approximately 0.5. Because the value for P_i is always 0 or 1, $(P_i - O_i)^2 \approx 0.25$. The approximate baseline sensitivity and hostile similarity scores can be calculated as follows:

$$S = 1 - \frac{\sum_{i=1}^n (P_i - O_i)^2}{n} \approx 1 - \frac{62(.25)}{62} = .75$$

Results

The sensitive similarity score produced by the *sensitive* network did not vary substantially based on input to the network (see Figure 2). When no input was provided, the sensitive similarity score of the *sensitive* network was .816; in each of the instances of receiving input, the sensitive similarity score was slightly greater, and ranged between .87-.89. This network’s hostile similarity score was .671 when no input was provided, and ranged between .71-.72 in each of the input conditions.

The *hostile* network showed slightly more variability in response to different input patterns. Sensitive similarity in the absence of input (.748) was the lowest in this network, relative to the other two networks. Sensitive similarity was lowest when the *hostile* network received “alarming” and

“afraid” input (.637), and highest (.795) when the network received only “mother” input. The “alarming” and “afraid” input appeared to be particularly influential: even when the network additionally received “sensitive” Response input, the sensitive similarity score was still relatively low, .727 – lower than the sensitive similarity score when no input was provided and when only “mother” input was provided. Finally, the impact of providing “hostile” Response input, in addition to “alarming,” “afraid,” and “mother” input, was relatively small. The sensitive similarity score in the former input condition was .699, only slightly lower than the sensitive similarity score of the latter, .713.

The *high-frequency exception* network produced high sensitive similarity scores in response to no input (.816), “mother” Partner input (.903), and the combination of “alarming” Context “afraid” Emotion-Pre, “mother” Partner, and “sensitive” Response input (.902). The sensitive similarity scores were substantially lower in response to “alarming” Context and “afraid” Emotion pre input (.677), and “alarming” Context, “afraid” Emotion-Pre, and “mother” Partner input (.668). The sensitive similarity to the Context, Emotion-Pre, and Partner input (.539). score was lowest when “hostile” Response input was added

Discussion

The results of the current simulation suggest that individuals with different attachment experiences could be expected to differ in terms of the accessibility of expectations of types of responses to attachment situations and emotional outcomes. The *sensitive* network, which had experienced very few negative outcomes of attachment experiences, learned that sensitive responses are generally likely, especially when these responses are needed (i.e., “alarming” and “afraid” input), and that “mother” is generally a sensitive partner, and does provide sensitive care during attachment experiences. Interestingly, even when her responses were hostile (i.e., “alarming,” “afraid,” “mother,” and “hostile response” input), the *sensitive* network nonetheless perceived the care to be sensitive, and produced expectations of favorable emotional outcomes. Thus, to the extent that predicting and perceiving positive/not negative attachment outcomes is indicative of the accessibility of positive working models of attachment, the *sensitive* network’s positive working models of attachment were highly accessible in each input situation.

The *hostile* network, however, received sensitive responses and experienced positive emotional outcomes virtually exclusively when contexts and associated initial emotion were not attachment-relevant. In other words, in attachment-relevant contexts when the *hostile* network initially experienced negative emotion (30% of its total experience), the outcomes were predominantly negative. This network learned that although sensitive responses were generally likely (in the absence of input indicating attachment-relevant situations; 70% of its experiences), alarming contexts in which fear was experienced were

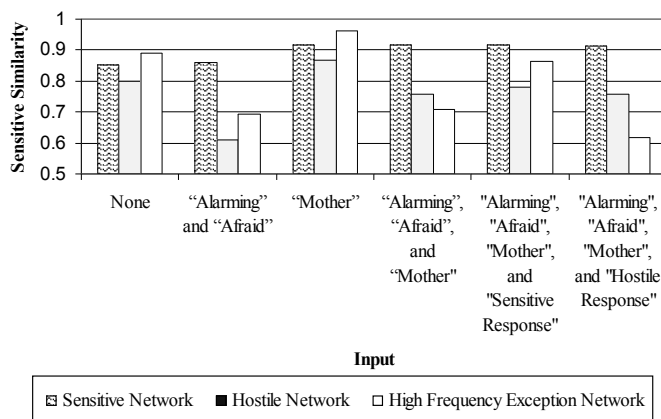


Figure 2: Accessibility of Expectations of Sensitive Care

incompatible with sensitive responses from others and favorable emotional outcomes. The presence of the mother, represented as someone likely to respond sensitively (and not necessarily hostilely), diminished the negative impact of the attachment context to some extent. The actual response of the mother, however, did not substantially impact the expectations generated by the network: when “mother” input was combined with “alarming” and “afraid” input, the network expected her to respond less sensitively and more hostilely, regardless of whether her response was objectively sensitive or hostile. In this network’s experience, the presence of a relationship partner during an attachment situation was sufficient to generate negative expectations. Because the expectation that attachment-relevant situations generally result in negative outcomes was largely accurate in this network’s experience, the network did not need to further differentiate knowledge regarding types of responses.

The current account suggests that the concept of “accessibility” of a (stored) representation is something of a misnomer. The networks did not have stored representations of relationships or particular experiences or prototypes available to them. Instead, the networks stored knowledge of tendencies of features of attachment experiences to be related with each other and to predict positive/negative responses and outcomes. Thus, rather than calling to mind a particular model, the networks generated predictions of attachment outcomes de novo, based on currently available information and knowledge of how this information tends to be related. What is thought of as high accessibility of a particular model is, in our formulation, a description of the overall or default likelihood of generating expectations consistent with a positive/negative “model.” The same knowledge structure, or network, however was used to process all information.

The current work can be extended to generate novel predictions regarding individual differences in accessibility of attachment knowledge and resultant behavior. For example, individuals who have experienced others as primarily sensitively responsive should consistently perceive others positively across a variety of different contexts, and should behave in ways that are consistent with these perceptions (e.g., seek support when support is needed). Conversely, individuals who have experienced caregivers positively only when they are not distressed, and individuals whose caregivers have responded inconsistently, should learn any number of contingencies that bias their perceptions of others, depending on the particular situation. Thus, these individuals should also behave consistently with their expectations and perceptions (e.g., seek support only under certain conditions, approach others only when not distressed).

Despite the important theoretical advances and novel implications of the current simulations, the current work has several significant limitations. First, the modeling of extremely complex human cognition is currently unfortunately, but necessarily, oversimplified. Real

individuals interact much more extensively with a greater number of relationship partners and to experience more varied instances and outcomes of adversity than could be modeled. Furthermore, the features that constituted attachment experiences in the current simulations were already relatively high-level concepts. A more complicated model could have allowed a network to discover relationships among more elemental features of experience, but our goal was to demonstrate relatively simply how the *principles* of connectionist modeling can contribute to understanding the structure of attachment knowledge.

The current model, therefore, simulates attachment processes and provides a foundation upon which to generate more specific predictions regarding important emerging themes in attachment research.

Acknowledgments

This work was supported by a National Science Foundation Interdisciplinary Graduate Education and Research Training Grant provided to the first author by the Center for the Neural Basis of Cognition.

References

- Bowlby, J. (1969/1982). *Attachment and loss: Attachment*. New York: Basic Books.
- Collins, N.L., Guichard, A.C., Ford, M.B., & Feeney, B.C. (2004). Working models of attachment: new developments and emerging themes. In W.S. Rholes & J.A. Simpson (Eds.), *Adult attachment: Theory, research, and clinical implications*. New York, NY, US: Guilford Publications.
- McClelland, J. L. and Rogers, T. T. (2003). The parallel distributed processing approach to semantic cognition. *Nature Reviews Neuroscience*, 4, 310-322.
- McLeod, P., Plunkett, K., & Rolls, E., (1998). *Introduction to Connectionist Modelling of Cognitive Processes*. Oxford: Oxford University Press.
- Mikulincer, M. & Arad, D. (1999). Attachment Working Models and Cognitive Openness in Close Relationships A Test of Chronic and Temporary Accessibility Effects. *Journal of Personality and Social Psychology*, 4, 710-725.
- Mikulincer, M., Gillath, O. & Shaver, P.R. (2002). Activation of the Attachment Systems in Adulthood Threat - Related Primes Increase the Accessibility of Mental Representations of Attachment Figures. *Journal of Personality and Social Psychology*, 4, 881-895.
- Mikulincer, M., Shaver, P., & Pereg, D. (2003). Attachment Theory and Affect Regulation: The Dynamics, Development, and Cognitive Consequences of Attachment-Related Strategies. *Motivation and Emotion*, 2, 77-102.
- Rumelhart, D. E., Smolensky, P., McClelland, J. L., & Hinton, G. E. Parallel Distributed Processing Models of Schemata and Sequential Thought Processes. In McClelland, J. L., Rumelhart, D. E., and the PDP research group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume II*. Cambridge, MA: MIT Press.