

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Simulating variation in infant-caregiver attachment using reinforcement learning

Permalink

<https://escholarship.org/uc/item/9kp245hr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 47(0)

Authors

Zhou, Xi Jia

Doyle, Chris

Cross, Logan Matthew

et al.

Publication Date

2025

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Simulating variation in infant-caregiver attachment using reinforcement learning

Xi Jia Zhou

Graduate School of Education
Stanford University

Chris Doyle

Graduate School of Education
Stanford University

Logan Cross

Department of Computer Science
Stanford University

Michael C. Frank

Department of Psychology
Stanford University

Nick Haber

Graduate School of Education
Stanford University

Abstract

Infants' attachment to their caregivers is a central feature of their early social and emotional development. Attachment Theory posits that these relationships vary systematically across distinct styles, though there has been debate about the extent to which these differences reflect features of caregivers' responsiveness vs. infants' own temperament. We develop a simple reinforcement learning model of infant exploration that allows us to vary the characteristics of simulated infants and caregivers and analyze the resulting patterns of model behavior. A set of equilibria reliably emerges that corresponds qualitatively to canonical attachment styles; particular agents' equilibria are controlled by both caregiver and infant parameters. These simulations point the way towards a quantitative synthesis of prior theoretical debates about the nature of attachment.

Keywords: infant attachment; development simulation; reinforcement learning models; formal theory; child development

Introduction

Attachment Theory is a lifespan model that emphasizes the central role of early caregivers as attachment figures to provide safety and security for infants to explore freely in the world (M. S. Ainsworth & Bowlby, 1991). The theory posits four archetypal infant attachment styles that describe how infants respond to separation from their caregivers: secure attachment, insecure avoidant attachment, insecure anxious-resistant attachment, and disorganized attachment (M. D. S. Ainsworth, Blehar, Waters, & Wall, 2015). For example, while a securely attached infant will be distressed at separation from their caregiver, they will be comforted by reunion; in contrast, an insecurely attached infant might react differently, either not showing distress at separation (insecure avoidant) or not being able to be comforted at reunion (insecure anxious-resistant). The four attachment styles are thought to describe how infants balance proximity to their caregiver vs. exploration, with the securely attached infant treating the caregiver as a "secure base" for exploration (Salter Ainsworth & Bell, 1981).

Attachment Theory is one of the most influential frameworks in early childhood development and has been applied to understand cross-cultural variation in caregiving (Van Ijzendoorn & Kroonenberg, 1988; Mesman, Van Ijzendoorn, & Sagi-Schwartz, 2016) and children's and adults' social expectations (Main, Kaplan, & Cassidy, 1985; Dykas & Cassidy, 2011; Waters & Waters, 2024). Yet it has also been the subject of substantial debates (Sroufe, 1985; van Ijzendoorn & Bakermans-Kranenburg, 2004). Debates about the

origins of attachment styles have centered on two major factors: caregiver sensitivity and infant temperament. While some researchers emphasize caregiver sensitivity and responsiveness as the primary determinants of attachment styles (Isabella, 1993; Dunst & Kassow, 2008), others argue for the role of infant temperamental characteristics, such as emotional reactivity and regulatory capacity "explaining away" the formation of attachment patterns (Thomas & Chess, 1984; Kagan, Reznick, Clarke, Snidman, & Garcia-Coll, 1984; Lewis & Feiring, 1989; Calkins & Fox, 1992; Cassidy, 1994).

One challenge in resolving these debates is that it is very difficult to study the causal mechanisms underlying attachment. Attachment is conceptualized as a process of learning that extends over time: infants are posited to learn from repeated interactions with their caregivers how sensitive and responsive they are. Although a wide variety of interventions have been used to increase caregiver sensitivity (Bakermans-Kranenburg, Van Ijzendoorn, & Juffer, 2003), it is challenging to move from these costly and time-consuming practical interventions to precise inferences about the learning mechanisms at work in infants.

An exception to this generalization is a series of infant habituation studies that probed infant expectations about animated characters that are shown to be separated from a caregiver (Johnson, Dweck, & Chen, 2007; Johnson et al., 2010). Infants who were classified as securely attached appeared to have different expectations than insecurely attached infants, with the securely attached infants predicting that the schematic "caregiver" would move to reunite with the "child". This provided the first experimental evidence that infants are forming abstract expectations about behavior, giving insight into the underlying cognitive mechanisms of attachment. Yet this work is challenging, especially given that it relies on prior classifications of children into attachment categories (which cannot itself be manipulated).

In contrast, computational models offer several key advantages as a method for studying developmental phenomena, including attachment (Smaldino, 2023; Frankenhuis, Borsboom, Nettle, & Roisman, 2023). First, they require that theoretical assumptions be explicit and constructs be precise. For example, infants' mental representations of caregivers might involve only their actions and the felt rewards without a more elaborative model, and infant temperaments can be constructed as levels of sociability vs. inhibition and ex-

1991

citability vs. irritability, formulated as parameters that are independent of cultural context. Second, they enable precise operationalization of constructs and manipulations, such as the amount of caregiver sensitivity, which are challenging to control or manipulate in real life. Third, they allow examination of parameter interactions, such as the interaction of caregiver sensitivity and infant emotion regulation ability. Our hope is that applying a computational modeling approach to infant attachment might reveal how caregiver sensitivity and infant temperamental characteristics might combine to produce different attachment patterns.

To date, there has been limited computational work that simulates attachments. One exception is Chumbley & Steinhoff, 2019, who provide a computational model that simulates the child's process of learning about a caregiver and shows how different attachment styles emerge from this learning process. This work explores the process of social reward learning in the caregiver-child dyad, but does not touch on one other aspect of attachment, namely how it trades off with the child's exploration of the broader world. The connection to exploration is critical for understanding the role of infant temperament. Critiques of attachment from the perspective of temperament theory suggest that high-arousal infants who tend to be more distressed by external stimuli might also stay closer to a caregiver as a source of comfort (Kagan, 1995). We explore this idea in our work.

Here, we develop a computational model of infant-caregiver attachment using the reinforcement learning (RL) framework. This framework allows us to examine how the simulated infant's internal states and interaction history inform its behaviors. We put these simulated infants in a simplified one-dimensional environment that captures the fundamental tension in attachment theory: the balance between exploration and proximity-seeking. Based on parameters that control its "temperament", the agent can remain close to the caregiver agent or explore further away. Further, we vary the parameters of the simulated caregiver that control whether it comforts, ignores, or punishes the infant agent. Even in this relatively simple setting, distinct behavior patterns emerge that resemble several canonical attachment styles in the real world. Our results support an integrative perspective, showing how both caregiver parenting styles and infant temperamental characteristics contribute to attachment outcomes.

Methods

Inspired by the Strange Situation Test (M. D. S. Ainsworth et al., 2015), we created an abstract virtual environment with a caregiver agent and an infant agent (See Figure 1). The environment is a finite 1-dimensional grid world. The infant agent has an internal emotional state and receives environmental and social rewards as it navigates the environment. Through reinforcement learning, it learns a policy to maximize its cumulative reward. The behavior of the caregiver agent is a parameter of the simulation; the caregiver agent can choose to comfort, ignore, or punish the infant agent. The




Figure 1: The reinforcement learning environment set up and agent design. The infant agent's exploration reward functions and the probability of encountering a hazard increase with distance from the caregiver (top). The infant agent's state can be affected by environmental hazards, interaction with the caregiver, or by self-regulation (bottom).

caregiver agent behaviors are intended to be analogies to different parenting styles that the Attachment Theory literature has suggested as causing factors for different infant attachment types (Isabella, 1993; Dunst & Kassow, 2008).

Environment and Environment Dynamics

The 1-dimensional environment has 10 total cells (see Figure 1). The infant agent can move between these cells and the caregiver agent is fixed in cell 0. The caregiver agent has an action space of size 3: it can comfort, ignore, or punish. These only have an effect when the infant is also in cell 0.

The state of the environment is defined by the position of the infant agent: x , the emotional state of the infant agent: e , which can be happy or distressed, and the duration the infant agent has been distressed: d . The system is initialized with $x_0 = 0$, $e_0 = \text{happy}$, and $d_0 = 0$.

The infant agent has an action space of size 3: it can go left, go right, or stay at the current location. The infant agent's emotional state can be changed by three factors: environmental hazards can cause a happy infant agent to become distressed, a distressed infant agent can self-regulate, or a caregiver agent can comfort or punish an infant agent when it is in cell 0. The frequency of hazards is denoted as $h(x)$, and the probability of the infant agent spontaneously moving from a distressed state to a happy state at each timestep is $p_{\text{self-regulate}}$. The probability of moving between emotional states can hence be expressed as follows.

$$P(e_{t+1} = \text{distressed} \mid e_t = \text{happy}) = \begin{cases} 1, & \text{if } x_t = 0 \text{ and } a_t^{\text{caregiver}} = \text{punish}, \\ h(x_t), & \text{otherwise.} \end{cases} \quad (1)$$

$$P(e_{t+1} = \text{happy} \mid e_t = \text{distressed}) = \begin{cases} 1, & \text{if } x_t = 0 \text{ and } a_t^{\text{caregiver}} = \text{comfort}, \\ p_{\text{self-regulate}}, & \text{otherwise.} \end{cases} \quad (2)$$

Rewards for Infant Agent The rewards the infant agent receives are a combination of “environmental” rewards away from the caregiver and “social” rewards close to the caregiver.

The infant agent can receive positive or negative environmental rewards. Positive environmental rewards are “exploration” rewards, where, the further the infant moves away from the caregiver, the higher the rewards. Negative rewards come when the infant agent is in the distressed state. The longer an infant agent is in distress, the more negative the rewards, which caps out at a ceiling at 30 timesteps. At positions far away from the caregiver agent, there are higher exploration rewards but also a higher probability of encountering an environmental hazard, leading to a distressed state and negative distress rewards. Environmental rewards are calculated as follows,

$$R_{\text{env}}(t) = \beta \cdot \mathbf{1}_{e=\text{happy}} \cdot E(x_t) + (1 - \beta) \cdot \mathbf{1}_{e=\text{distressed}} \cdot D(d_t), \quad (3)$$

where $E(x)$ and $D(d)$ are exploration and distress reward functions, and d_t is the duration that the infant agent is distressed (since it was last happy, not cumulatively). The parameter β defines the infant agent’s relative sensitivity to exploration and distress.

The caregiver agent can interact with the infant agent when its in cell 0. If the infant is distressed and the caregiver chooses the “comfort” action, the infant receives a positive reward and becomes happy. If the caregiver punishes the infant, it receives a negative reward and becomes or remains distressed. If the caregiver agent ignores the infant agent, it receives no reward. Social rewards are as follows,

$$R_{\text{social}}(t) = \mathbf{1}_{x=0} \cdot C(a_t^{\text{caregiver}}), \quad (4)$$

where $C(a^{\text{caregiver}})$ maps caregiver agent actions to reward:

$$C(a^{\text{caregiver}}) = \begin{cases} -1, & \text{if } a^{\text{caregiver}} = \text{punish}, \\ \mathbf{1}_{e=\text{distressed}}, & \text{if } a^{\text{caregiver}} = \text{comfort}, \\ 0, & \text{if } a^{\text{caregiver}} = \text{ignore}. \end{cases} \quad (5)$$

The environmental and social rewards are combined as follows, with the parameter α defining the infant agent’s relative sensitivity to environmental and social rewards,

$$R(t) = \alpha \cdot R_{\text{env}} + (1 - \alpha) \cdot R_{\text{social}}(t) \quad (6)$$

Function and Parameter Choices A linear-exponential combination is used to parameterize exploration rewards $E(x)$, distress rewards $D(d)$, and the probability of hazards $h(x)$. All are normalized such that the maximum value in their range is 1.

$$\begin{aligned} E(x) &\propto x + e^{0.4x} \\ -D(d) &\propto d + e^{0.2d} \\ h(x) &\propto x + e^{0.4x} \end{aligned} \quad (7)$$

We make these function choices to satisfy a few criteria. First, rewards for exploration and the probability of encountering an environmental hazard should increase with distance, as novel environments can be more rewarding but also more risky (Gopnik et al., 2017), and the effect of being distressed should increase in magnitude (more negative) over time (Feldman, Singer, & Zagoory, 2010). Second, the scale of rewards and hazard probability must allow for a diversity of behavior such that caregiver behavior or agent parameters can affect the optimal strategy. Third, we sought a function class that was simple and flexible. We used linear and exponential terms $ax + e^{bx} + c$, which can produce a wide range of shapes. In preliminary experiments, we found that many combinations yielded a similar diversity of behaviors as our chosen function.

Agent Implementation

Caregiver Within one simulation, the caregiver agent has a fixed policy. We implement three policies based on patterns identified in attachment literature (M. D. S. Ainsworth et al., 2015): the “comforting” policy, where the caregiver agent consistently comforts the infant; the “punishing” policy, where the caregiver agent consistently punishes the infant; and the “inconsistent” policy, where the caregiver agent sometimes comforts and sometimes ignores the infant (with a 10 percent chance to comfort and a 90 percent chance to ignore on each step). These policies represent simplified models of caregiving patterns that research has associated with different attachment classifications: consistent responsive caregiving is associated with secure attachment patterns; consistent unresponsive or rejecting caregiving is associated with insecure-avoidant attachment patterns; and inconsistent caregiving is associated with insecure anxious-resistant attachment patterns in infants.

Infant The system can be described as a Markov decision process, defined by the states, actions, transition probabilities, and reward functions we have described above. Given that the state and action spaces are discrete and finite, we can apply a tabular reinforcement learning (RL) method to optimize the infant agent’s behavior policy π_{infant} to maximize cumulative reward. The learning algorithm used is ϵ -greedy Q-learning, a common approach for a tabular setting (Sutton, 2018), with a learning rate α decreasing linearly over training from 0.2 to 0.01. The discount factor γ for future reward is set as 0.99.

$$a_{\text{infant}} \sim \pi_{\text{infant}}(x, e, d) \quad (8)$$

Simulation Details

A parameter combination for a simulation is defined by the choice of $\alpha, \beta, p_{\text{self-regulate}}$, and $\pi_{\text{caregiver}}$. For each parameter combination, we ran ten simulations, each with a different random seed. Each simulation included 10,000 episodes. At the end of each episode the environment is reinitialized. Episodes are a fixed length: 500 timesteps. We linearly decayed the random component of the policy (ϵ) from 1.0 to




Figure 2: Average Infant Location in Emotional States: interaction effects between caregiver parenting style and infant emotion regulation. This figure shows 270 simulated infants’ averaged locations at the end of training (200 episodes; 27 parameter sets: 3 levels of caregiver style x 9 levels of infant emotion regulation, each run with 10 random seeds). Connected areas are convex hulls containing each point in a parameter set. Colors represent the level of $p_{\text{self-regulate}}$, and data shapes represent different caregiver styles.

0.05 over the first 80% training episodes, which was then held constant at 0.05. This training length was selected based on preliminary experiments, which indicated that it was sufficient for the observed stabilization of infant behaviors.

Results

We report two experiments exploring how caregiver styles interact with aspects of the infant, including temperament and preferences for different kinds of reward, in guiding infant agents’ behavior. In each, we examine the infant agent’s average location in the happy state vs. distressed states. This analysis allows us to characterize the distance infant agents are willing to explore in the happy state and to observe their strategy while distressed. By allowing us to identify clusters of agents, this analysis also allows us to draw connections to traditional classifications in the attachment literature.

Experiment 1

In Experiment 1, we study the interaction between caregiver parenting style and infant emotion regulation. We varied both the caregiver agent’s policy (comforting, inconsistent, or punishing; see Methods) and the infant agent’s probability of self-regulation (0.02 to 0.2, with 9 total levels). We kept α (infant agents’ sensitivity to environmental vs. social rewards) and β (infant agents’ sensitivity to exploration vs. distress) fixed at 0.5. Therefore, in this model, we have 3 levels of different caregiver parenting styles and 9 levels of infant emotion regulation probabilities varying, which gives us 27 total different simulated infants (3 Caregiver Parenting Style x 9 Infant Emotion Regulation). We ran 10 seeds for each combination

of parameters.

Emergent Clusters Emulate Attachment Styles Our simulations reveal distinct behavioral clusters that emerge from the interaction between the infant emotion regulation and caregiver styles (Figure 2). In what follows, we interpret them in light of styles described in classical Attachment Theory. A primary division occurs between “Low Self-regulation” infants ($p_{\text{self-regulate}} < 0.09$) and “High Self-regulation” infants (> 0.09), with the Low Self-regulation group in the bottom left of the plot exhibiting patterns that closely mirror traditional attachment style behaviors.

For Low Self-regulation, we observe distinct behaviors.

Secure: With Comforting caregivers (circles in Figure 2), these infant agents demonstrate balanced exploration and proximity-seeking behaviors characteristic of secure attachment. This leads to a behavior where infants frequently move away from the caregiver to explore and return when distressed. A visualization of the learned policy is in Figure 3 (bottom left). In the Comforting caregiver condition, comfort happens immediately, and the infant can explore once more. This leads to a subgroup of infants with approximately equal average positions in the happy and distressed states.

Insecure Anxious: Inconsistent caregivers (squares in Figure 2) lead Low Self-regulation infants to a similar pattern, except the infant must spend more time next to the unreliable caregiver to receive comfort. This causes the average position while distressed to decrease compared to the secure group. Figure 3 shows an example of their learned policy (top right).

Insecure Avoidant: With Punishing caregivers (Crosses in Figure 2), this subtype maintains a substantial distance from the caregiver when distressed, actively avoiding caregiver interaction during these periods. This leaves Low Self-regulation infants to potentially long bouts of distress. To avoid this long distress period, these infants are very conservative in their exploration so as to expose themselves to only minimal risk of encountering an environmental hazard. We notice that upon becoming distressed, infant agents in this condition “wander” rather than remain in place, potentially due to the receiving the same distress reward at all points away from the caregiver. Figure 3 shows an example of their learned policy (top left).

The **High Self-regulation** group exhibits a qualitatively different pattern that is less sensitive to caregiver styles and doesn’t map cleanly onto traditional attachment styles. With a high capacity for self-regulation, their strategy is to gain exploration reward, and, when distressed, just wait to self-regulate. Self-regulation is likely to occur before the largest magnitude distressed rewards. Thus, these infant agents consistently maintain significant distance from the caregiver (positions 5-9) regardless of emotional state or caregiver type.

Multiple combinations of parameter sets result in similar infant agent behavior Even with two parameters in our simulations, we see that different combinations of parameter




Figure 3: Visualizations of sample of infant agents’ learned policies, one from each of the subgroups found in Experiment 1. There are two sections to each subplot. To the left of the vertical dashed line, the optimal actions are shown for the happy state. Attractor states, both fixed points and cycle points, are noted with a colored square to indicate where the infant could reside in equilibrium while in a happy state. The best actions are visualized at arrows, pointing in the preferred direction of movement. Sideways arrows indicate remaining in the same position. To the right of the dotted line, the best actions are visualized for each point that the infant is distressed. Lines are drawn over these actions to indicate the path the infant would take if it remains distressed.

sets can result in similar average locations. This is visible in Figure 2 in three areas: around the points (3.5, 3.5), (5.5, 7), and (7, 8). In the first area, there is overlap of infants with Comforting and Inconsistent caregivers. In the other two areas there is overlap of Comforting, Inconsistent, and Punishing caregivers. This observation is further supported by Figure 3, where the Low Self-regulation, Inconsistent Caregiver and Low Self-regulation, Comforting Caregiver highest value actions are similar, and by the fact that High Regulation infants share a similar strategy to explore and wait to self-regulation upon distress. These results imply that, even in a simple system, it can be challenging to infer the parameters of the generating process, as the interaction between parameters can cause convergent behavior across different caregiver types.

Experiment 2

In Experiment 2, we study how infant agents’ temperamental factors – their preferences for environmental reward over social reward and their temperamental reactivity on exploration reward over distress reward – impact their behaviors. To do this, we vary α and β values: α represents how the infant agent weighs environmental reward against social (caregiver) reward (higher α increases the relative weight on environment), and β represents how much the infant agent weighs positive exploration reward against negative distress reward

(higher β puts less emphasis on distress). We use 3 levels of α and β : 0.2, 0.5, 0.8. We reduce infant emotion regulation to 2 levels (0.02 and 0.2) from 9 levels in experiment 1, resulting in 54 total different infants (3 Caregiver Parenting Style x 2 Infant Emotion Regulation x 3 α x 3 β). For each parameter combination, we run 10 random seeds.

Low relative sensitivity to exploration can induce qualitatively different behaviors A new cluster of infant agents emerges in a combination of parameters with Inconsistently rewarding caregivers, low (0.02) emotion regulation probability, and low β (0.2), that is, relatively higher sensitivity to distress rewards. These infant agents have an average position in happy and distress states near zero – any possible exploration rewards are not worth the chance of distress.

Increasing the relative sensitivity to exploration pushes the average distance from caregiver higher Figure 4 shows the effect of changing β from 0.2 to 0.8, increasing the sensitivity to exploration rewards. Most infants move farther away from the caregiver with higher β , in both happy and distressed states. This shows when we increase infant agents’ sensitivity to exploration rewards, they explore further.

Altering the relative sensitivity to environmental and social rewards did not have a clear directional effect on average location Figure 4 shows the same impact of increasing α from 0.2 to 0.8 on the infant agents’ location in the grid world. We do not see a large effect or pattern of effect when increasing or decreasing α . Most of the average distances appear similar for values of 0.2 and 0.8. Together with the large effect when varying β , this suggests that infant decision making may be dominated by environmental rewards, and social rewards may only play a small part in our simulations.

Discussion

In this paper, we described a computational model that uses reinforcement learning to simulate infant attachment behaviors via the tradeoff between their exploration of the environment and their proximity-seeking with a caregiver agent. We used this model to probe how the characteristics of the simulated infants and caregivers impact their attachment behaviors. In two experiments, we explored the interaction of caregiver sensitivity with infant emotion regulation ability and the effects of infant temperamental factors on behavior. In Experiment 1, we show an interaction effect of caregiver sensitivity and infant emotion regulation ability, with four clusters emerging: Secure, Insecure Anxious, Insecure Avoidant, and High Self-regulation. When infant agents have high emotion regulation, they self-soothe and explore far in the environment, less sensitive to caregiver parenting styles; but when they have low emotion regulation ability, their proximity-seeking behaviors are dominated by the caregiver agents’ parenting styles. In Experiment 2, we show that increasing infant agents’ reactivity to exploration reward, relative to distress reward, allows them to explore further. In contrast, increasing infant agents’ reactivity to environmental rewards does not




Figure 4: Average Infant Location in Emotional States: location changes when β (left) and α (right) changes from 0.2 to 0.8. Each plot shows 180 simulated infants’ averaged locations in emotional states at the end of training (200 episodes; 18 combinations of parameters: 3 Caregiver Styles x 2 Infant Emotion Regulation x 3 α or 3 β levels, across 10 seeds). For each parameter combination, we run 10 random seeds. Blue and red dots show simulated infants’ average locations when $\beta = .2$ and $.8$, respectively (left); yellow and green dots show simulated infants’ average locations when $\alpha = .2$ and $.8$, respectively (right); gray arrows indicate infants with otherwise matched parameters. Data shape represents different caregiver styles.

make infant agents explore further.

Our model reproduces major parts of the existing attachment theory literature. We show that clusters of behavior corresponding to the archetypal attachment styles emerge from our model. However, a long-standing debate in the literature concerns the role of infant temperament in modulating caregiver behavior. Our model is consistent with the basic role of temperament in modulating attachment style: different attachment styles only emerge in our model when parameters corresponding to infant temperament are taken into account. Our simulation results also align with the theory that fearful (“high-arousal”) infants – infants who are more easily distressed by external stimuli – might stay closer to a caregiver as a source of comfort in distress (Thomas & Chess, 1984).

Not all of our model’s conclusions support claims about temperament, however. In contrast to the claims of temperament theory (Lewis & Feiring, 1989), decreasing social reward does not make infant agents more avoidant of caregivers. Most infant agents in our simulation did not change their average locations in either emotional state with different values of alpha. In addition, temperament theory often implies that temperamental factors could determine infants’ insecure attachment style. However, in our model, for the most part, agents do not move between clusters based on temperament. Overall, our simulation results support an integrative perspective on attachment, showing how both caregiver parenting styles and infant temperamental characteristics contribute to attachment outcomes.

Like any model of a rich set of social behaviors, our model

has a number of limitations. Our environment is far removed from the real world, and our agents abstracted away from many of the child or caregiver factors identified in prior work (Esposito, Setoh, Shinohara, & Bornstein, 2017; Karakaş & Dağlı, 2019). In addition, while we tried to vary a range of environment design choices and parameter settings, our search was not exhaustive for the constructs under consideration, and we could reach different conclusions under different choices.

Indeed, what makes our framework exciting are the possibilities for extension. We could increase the complexity of both the environment and the agents, allowing, for example, more human-like agents with multi-modal sensory inputs and mental representations. The RL framework also provides us with unique opportunities to look into infant agents’ minds – such as observing their learned strategies – and to conduct experiments that are impossible in the real world – such as putting the infant agent in severely adverse environments.

Since Bowlby’s initial formulation of Attachment Theory (Bowlby, Ainsworth, & Bretherton, 1992), researchers have built a rich empirical foundation through careful observations, experiments, and interventions in labs, homes, and clinics. With new computational tools, we gain the power to translate abstract theoretical constructs into computational parameters, systematically test alternative hypotheses, discover emergent patterns, and generate novel predictions. As we bridge classical psychological theories with modern computational approaches, we open new pathways for the theories of tomorrow.

References

- Ainsworth, M. D. S., Blehar, M. C., Waters, E., & Wall, S. N. (2015). *Patterns of attachment: A psychological study of the strange situation*. Psychology press.
- Ainsworth, M. S., & Bowlby, J. (1991). An ethological approach to personality development. *American psychologist*, 46(4), 333.
- Bakermans-Kranenburg, M. J., Van IJzendoorn, M. H., & Juffer, F. (2003). Less is more: meta-analyses of sensitivity and attachment interventions in early childhood. *Psychological bulletin*, 129(2), 195.
- Bowlby, J., Ainsworth, M., & Bretherton, I. (1992). The origins of attachment theory. *Developmental Psychology*, 28(5), 759–775.
- Calkins, S. D., & Fox, N. A. (1992). The relations among infant temperament, security of attachment, and behavioral inhibition at twenty-four months. *Child development*, 63(6), 1456–1472.
- Cassidy, J. (1994). Emotion regulation: Influences of attachment relationships. *The development of emotion regulation: Biological and behavioral considerations. Monographs of the society for Research in Child Development/Fox (Ed)*.
- Chumbley, J., & Steinhoff, A. (2019). A computational perspective on social attachment. *Infant Behavior and Development*, 54, 85–98.
- Dunst, C. J., & Kassow, D. Z. (2008). Caregiver sensitivity, contingent social responsiveness, and secure infant attachment. *Journal of Early and Intensive Behavior Intervention*, 5(1), 40.
- Dykas, M. J., & Cassidy, J. (2011). Attachment and the processing of social information across the life span: theory and evidence. *Psychological bulletin*, 137(1), 19.
- Esposito, G., Setoh, P., Shinohara, K., & Bornstein, M. H. (2017). *The development of attachment: Integrating genes, brain, behavior, and environment* (Vol. 325). Elsevier.
- Feldman, R., Singer, M., & Zagoory, O. (2010). Touch attenuates infants' physiological reactivity to stress. *Developmental science*, 13(2), 271–278.
- Frankenhuis, W. E., Borsboom, D., Nettle, D., & Roisman, G. I. (2023). Formalizing theories of child development: Introduction to the special section. *Child development*, 94(6), 1425–1431.
- Gopnik, A., O'Grady, S., Lucas, C. G., Griffiths, T. L., Wente, A., Bridgers, S., ... Dahl, R. E. (2017). Changes in cognitive flexibility and hypothesis search across human life history from childhood to adolescence to adulthood. *Proceedings of the National Academy of Sciences*, 114(30), 7892–7899.
- Isabella, R. A. (1993). Origins of attachment: Maternal interactive behavior across the first year. *Child development*, 64(2), 605–621.
- Johnson, S. C., Dweck, C. S., & Chen, F. S. (2007). Evidence for infants' internal working models of attachment. *PSYCHOLOGICAL SCIENCE-CAMBRIDGE-*, 18(6), 501.
- Johnson, S. C., Dweck, C. S., Chen, F. S., Stern, H. L., Ok, S.-J., & Barth, M. (2010). At the intersection of social and cognitive development: Internal working models of attachment in infancy. *Cognitive science*, 34(5), 807–825.
- Kagan, J. (1995). On attachment. *Harvard Review of Psychiatry*, 3(2), 104–106.
- Kagan, J., Reznick, J. S., Clarke, C., Snidman, N., & Garcia-Coll, C. (1984). Behavioral inhibition to the unfamiliar. *Child development*, 2212–2225.
- Karakaş, N. M., & Dağlı, F. Ş. (2019). The importance of attachment in infant and influencing factors. *Turkish Archives of Pediatrics/Türk Pediatri Arşivi*, 54(2), 76.
- Lewis, M., & Feiring, C. (1989). Infant, mother, and mother-infant interaction behavior and subsequent attachment. *Child development*, 831–837.
- Main, M., Kaplan, N., & Cassidy, J. (1985). Security in infancy, childhood, and adulthood: A move to the level of representation. *Monographs of the Society for Research in Child Development*, 50(1/2), 66–104.
- Mesman, J., Van IJzendoorn, M. H., & Sagi-Schwartz, A. (2016). Cross-cultural patterns of attachment. *Handbook of attachment: Theory, research, and clinical applications*, 3, 852–877.
- Salter Ainsworth, M. D., & Bell, S. M. (1981). Attachment, exploration, and separation: Illustrated by the behavior of one-year-olds in a strange situation. In *The life cycle: Readings in human development* (pp. 57–71). Columbia University Press.
- Smaldino, P. (2023). *Modeling social behavior: Mathematical and agent-based models of social dynamics and cultural evolution*. Princeton University Press.
- Sroufe, L. A. (1985). Attachment classification from the perspective of infant-caregiver relationships and infant temperament. *Child development*, 1–14.
- Sutton, R. S. (2018). Reinforcement learning: An introduction. *A Bradford Book*.
- Thomas, A., & Chess, S. (1984). Genesis and evolution of behavioral disorders: from infancy to early adult life. *The American journal of psychiatry*, 141(1), 1–9.
- van IJzendoorn, M. H., & Bakermans-Kranenburg, M. J. (2004). Maternal sensitivity and infant temperament in the formation of attachment. *Theories of infant development*, 231–257.
- Van IJzendoorn, M. H., & Kroonenberg, P. M. (1988). Cross-cultural patterns of attachment: A meta-analysis of the strange situation. *Child development*, 147–156.
- Waters, E., & Waters, T. E. (2024). Developmental change, bricolage, and how a lot of things develop: Mechanisms and changes in attachment across the lifespan. *Development and Psychopathology*, 36(5), 2256–2275.