

An associative learning account of how saliva becomes a cue for comfort

Deon T. Benton^{*†}, and Justine E. Hoch^{*†}

Department of Psychology and Human Development, Vanderbilt University, Nashville, TN, United States

^{*}Corresponding author: Deon T. Benton, Department of Psychology and Human Development, Vanderbilt University, Peabody College, 230 Appleton Place #552, Nashville, TN 37235, United States. Email: deon.benton@vanderbilt.edu

[†]Deon T. Benton and Justine E. Hoch share first-authorship.

Abstract

A recent study by [Thomas et al. \(2022\)](#) showed that 8.5- to 10-month-olds use saliva sharing as an indicator of relationship closeness: infants expect that someone who previously shared saliva with a partner will also comfort that partner during times of distress. The current paper offers a mechanistic explanation for this finding. We propose that through everyday experiences, saliva sharing and comforting become integrated into an enriched representation that contains information about both events. Across five simulations, we show that when this proposal is implemented in a connectionist computational model, the model reproduces infants' looking behaviors in the study by [Thomas et al. \(2022\)](#) (Simulation 1), captures the finding that infants' responses were specific to comfort (Simulation 2), predicts bidirectional expectations about saliva sharing and comforting (Simulation 3), and shows that these expectations strengthen with experience and can shift in response to changes in environmental statistics (Simulations 4 and 5). Taken together, our findings suggest that domain-general associative learning can explain how infants come to understand the social meaning of saliva sharing.

Keywords connectionist models, computational modeling, mechanisms of cognitive change, cognitive development

Lay summary

How do infants know who will help when someone is upset? Recent research shows that infants expect people who share saliva to comfort one another. The current study uses computer simulations to explain how infants develop this understanding. We propose that infants learn these social rules through everyday observations. By seeing people share saliva and provide support, infants learn to connect these two actions. Our simulations successfully explained infant behavior and showed that these expectations not only grow stronger with more experience but can also be reversed in extreme cases. This suggests that infants use general learning abilities to build an understanding of close relationships based on the patterns they observe in their daily lives.

To make sense of the social world, preverbal infants must learn to detect social relationships and evaluate their strength, or “thick-ness.” Thick relationships—such as those among immediate family members—are marked by strong levels of attachment, loyalty, obligation, and mutual responsiveness ([Ainsworth et al., 1974](#); [Clark & Mills, 1979](#); [Fiske & Fiske, 2007](#); [Miller et al., 1998](#); [Thomas, 2024](#)). Distinguishing thick relationships from thinner ones enables infants to predict how people are likely to behave. For example, a father might be expected to comfort his own child, but not another child. However, despite the clear social utility of detecting thick relationships, the developmental origins of this ability are poorly understood.

One account proposes that infants infer the strength of social relationships by observing external cues, specifically saliva sharing, which is a reliable indicator of thick relationships. Because saliva sharing is a cross-cultural and cross-species hallmark of intimacy ([Jaeggi & Gurven, 2013](#); [Silk et al., 2013](#); [Thomas, 2024](#)), [Thomas et al. \(2022\)](#) investigated whether infants use this behavior to predict social support. Their study, published in *Science*, found that infants expect an individual who previously shared saliva with a partner also to comfort that partner during times of distress. In Thomas et al.'s primary experiment (i.e., Experiment 2B), 8.5- to 10-month-old infants saw a puppet interact with two actresses. One actress took turns eating from the same orange

Handling Editor: Haley Vlach

Editor: Adam Winsler

Received: November 5, 2025. **Revised:** June 3, 2026. **Accepted:** June 8, 2026

© The Author(s) 2026. Published by Oxford University Press on behalf of the Society for Research in Child Development. This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial License (<https://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited. For commercial re-use, please contact reprints@oup.com for reprints and translation rights for reprints. All other permissions can be obtained through our RightsLink service via the Permissions link on the article page on our site—for further information please contact journals.permissions@oup.com.

slice as the puppet (implying saliva sharing). The other actress shared a ball with the puppet, but did not share saliva. Unlike saliva sharing, which is theorized to occur mainly in thick relationships, ball sharing represents a neutral interaction that can occur in both thin and thick relationships. Infants then saw the same puppet positioned between the two actresses on the screen. The puppet exhibited distress, and both actresses turned toward the puppet. Thomas et al. (2022) reasoned that if infants expect saliva-sharers to provide comfort during distress, infants should look first and longest at the actress who ate from the same orange slice as the puppet—and this is exactly what they found.

To test whether infants' expectations were specific to previous saliva-sharing partners, rather than a general expectation that saliva-sharers comfort *anyone* in distress, Thomas et al. (2022) repeated the initial study but replaced the original puppet with a new puppet at the beginning of the distress event (Experiment 2D). When the new puppet was in distress, infants looked equally long at both actresses. The findings from these two studies suggest that infants expect saliva-sharers to comfort their saliva-sharing partners but not distressed individuals generally.

In a third experiment (Experiment 3B), Thomas et al. (2022) tested whether infants' behavior in the previous experiments was specifically due to saliva sharing (rather than food sharing). They showed infants a video of a central actress flanked by two puppets. First, the actress placed her finger into her mouth, rotated it, and then placed it into one of the puppet's mouths. Next, the actress placed her finger on her forehead, rotated it, and placed it on the other puppet's forehead. The actress then expressed distress, and Thomas et al. (2022) measured which puppet infants looked at first and longest. Although infants' first looks were split evenly between both puppets, infants looked longer at the puppet with whom the actress shared saliva. Thomas et al. (2022) interpreted these findings to mean that preverbal infants use saliva sharing specifically to predict who should provide comfort to whom. These results (Experiment 2B) were partially replicated by Ciftci et al. (2025); these authors found that 8.5- to 10-month-old infants looked proportionally longer, but not first, at the saliva sharer than at the non-saliva sharer.

Although Thomas (2024) posits an evolutionary basis for the cues infants use to detect "thick" relationships, she remains largely agnostic about the developmental origins of these expectations. Here, we offer a mechanistic account of how saliva sharing may come to predict comfort. We return to how this mechanistic proposal can be reconciled with an evolutionary account in the General Discussion.

Our central claim is that infants' looking behaviors in Thomas et al. (2022) are the result of a domain-general associative learning mechanism. As infants repeatedly witness others sharing saliva or providing comfort, they gradually build a representation of these events. Here, "representation" refers to a pattern of activation across units in an associative, cognitive system that is elicited by different environmental events such as the above saliva-sharing event. In our view, infants represent saliva-sharing in terms of its low-level, semantically neutral kinematics and the surface features of the objects and entities that engage in them. For example, infants' representation of saliva-sharing events will include information about the perceptual appearance of the agents, their actions, and the food items involved. However, because the particular food item that is shared tends to vary across situations, the

kinematics of saliva sharing will feature more prominently in the representation that is formed than any one food item.

We propose that during saliva-sharing events, the perceptual features that embody the kinematics of saliva sharing become associated with the individuals participating in that action. If that same person later participates in comforting, that person will reactivate some of the same representational units that encoded the action of saliva-sharing itself. Once reactivated, the kinematics that comprise comforting become incorporated into the same representation. A corollary of this is that a single "enriched" representation emerges that contains information about saliva sharing and comforting, even though saliva sharing and comforting were not experienced simultaneously. As a consequence, our account predicts bidirectional expectations; when infants see saliva sharing, they will expect comforting, and when infants see comforting, they will expect saliva sharing. Importantly, we are not arguing that this enriched representation encodes the conceptual meaning of saliva sharing or comfort; rather, this representation reflects the combined (hence, "enriched") low-level, surface-based kinematic and perceptual features of these actions and entities involved in them.

In the next section, we present a connectionist computational model that instantiates this account. Although Thomas et al. (2022) also tested older children, we focus on infants who cannot rely on explicit instruction. Across five simulations, we demonstrate that the model reproduces infants' looking behaviors in Thomas et al. (2022) (Simulation 1), captures the finding that infants' responses were specific to comfort (Simulation 2), predicts bidirectional expectations between saliva sharing and comforting (Simulation 3), and shows that these expectations strengthen with experience and can shift in response to changes in environmental statistics (Simulations 4 and 5).

All simulations were run using the Lens neural network simulator (Rohde, 1999). All training and testing files, including the data from the simulations, are available for download at https://osf.io/9mqyn/overview?view_only=c09f319c3fa643d9b688cb09e0794379.

Simulation 1: Saliva sharing predicts comforting

Network architecture

The simulations reported here used a three-layer, feedforward, autoassociative neural network model (Mareschal et al., 2000) trained with backpropagation in cross-entropy error (Rumelhart et al., 1986). The hidden and output units used a sigmoid activation function. Gaussian noise was added to input activations to reflect infants' noisy sensory processing ($M=0$, $SD=0.875$). The learning rate was set to 0.1 in all simulations except for Simulation 4 (see below); momentum was not used in any of the simulations.

The model consisted of three layers (Figure 1). There were 147 input units, 30 hidden units, and 147 output units. The 147 input and output units were divided into 10 groups of units—the output groups were identical to the input groups. Saliva Sharing Identity and Comforting Identity groups captured the *identity* of being a saliva sharer or comforter; the Saliva Sharing Action and Comforting Action groups captured the *action* of engaging in saliva sharing or comforting. By separating an individual's

identity from its actions, it was possible for the model to represent situations in which a former comforter engages in saliva sharing or a former saliva sharer engages in comforting. For each input group, the model's task was to activate in the output layer exactly those units that were active in the input layer. Localist representations—in which a single unit represented a single concept—were adopted for each group to simplify the interpretation and analysis of the model. Although the names of the groups and features were conceptually laden (e.g., Distressed Action), we do not intend to suggest that infants see these features in a semantically meaningful or rich way. Our view is that infants see these actions as nothing more than semantically neutral, surface-based, action-centered visual events (see Benton, 2025 for a similar argument). The rationale for using such semantically meaningful labels to describe the model is to make the structure of the model maximally transparent. Two final points are worth

noting: First, because saliva/non-saliva sharing and comforting/non-comforting necessarily require two people, whenever either of these actions occurred in the present simulations, two people were present (i.e., a single unit was turned on in the People Left and Right groups). Second, in conditions that involved a new character (e.g., the New Puppet condition in Simulation 1 and Experiment 2D in Thomas et al.) and during preferential-looking tests (e.g., Simulation 2 and the “peek-a-boo” condition in Thomas et al.), only person identity and action information were provided to the model.

Pretraining

The pretraining phase—which approximated infants' real-world experience—lasted between 600 and 800 epochs, in increments of 100. Different numbers of epochs were compared to ensure that




Figure 1 Schematic of an untrained model architecture used in the current simulations. Note. Depicted here is a saliva-sharing event. The input and output groups are identical. The names of the input groups are in bold. The model consisted of 10 structurally identical input and output groups: People Left/Right, Saliva Sharing Left/Right Identity, Comforting Left/Right Identity, Comforting Action, Distressed Action, Food Item, and Saliva Sharing Action. The specific feature that a given unit in a group codes for is shown above that unit in italics. Note that the figure is color-coded to make it easier to connect network training (shown in Tables 1 and 2) to the model architecture. The borders around each output unit correspond to that unit's target activation value. Black borders correspond to target activations of “0” (i.e., the unit should be “off”); white borders correspond to target activations of “1” (i.e., the unit should be “on”). The size of a given output unit's inner region corresponds to that unit's actual activation value. Larger inner regions correspond to greater activations. Darker inner regions for the output units correspond to activation values ≤ 0.5 , whereas lighter inner regions correspond to activation values > 0.5 .

the network's performance was not idiosyncratic to the length of the pretraining phase, and different food items (represented by one of five units in the purple box in Figure 1) were used in each saliva-sharing event to instantiate the idea that the crucial link is between saliva sharing and comforting rather than between food sharing and comforting. Networks experienced ten kinds of events during training. Half of the events depicted saliva and non-saliva sharers engaging in different actions (Table 1); the remaining half depicted comforters and non-comforters engaging in different actions (Table 2). These events instantiated our assumptions about the kinds of real-world, saliva-sharing, and comforting experiences that influenced infants' looking behaviors in Thomas et al. (2022). Each event was implemented by activating the relevant units in the Person, Saliva Sharing/Comforting Identity, Distressed Action, and Saliva Sharing/Comforting Action groups (see Tables 1 and 2).

Familiarization

Following the pretraining phase, networks were randomly assigned to one of three conditions (20 networks per condition): (1) Food Sharing; (2) New Puppet; and (3) Saliva Only. The conditions corresponded, respectively, to Experiments 2B, 2D, and 3B in Thomas et al. (2022). A "participant" corresponded to an untrained model whose weights were initialized to small, random values (sampled uniformly between ± 0.1). Note that the Saliva Only condition was identical to Food Sharing and New Puppet conditions except that food information was missing (as was the case in Study 3B in Thomas et al., 2022). The familiarization phase for networks in all three conditions involved an event in which food/saliva sharing occurred between a central character and another

individual and an event in which non-food/non-saliva sharing (e.g., ball passing) occurred between the same central character and a second individual. Because our theory does not depend on who shared versus received saliva, we did not model the direction of sharing. All familiarization events used novel people and were implemented in a similar manner to the pretraining events (see Tables 1 and 2). Finally, the familiarization phase could last for either 5 or 10 epochs. Different values ensured that the model's performance was robust to the length of the familiarization phase.

Test

Following the familiarization phase, networks participated in a test phase. The nature of this test phase depended on the condition to which the networks were assigned. Networks in the Food Sharing and Saliva Only conditions saw an event in which the same central character expressed distress. In contrast, networks in the New Puppet condition saw the same event except that the character that was in distress differed from the one shown during the familiarization phase.

For all conditions, we measured the cross-entropy error produced when the network saw the saliva-sharer provide comfort and the non-saliva-sharer provide comfort. Typically, greater model error corresponds conceptually to longer looking and mirrors the logic of violation of expectation paradigms. Thus, the model should produce greater error if it had learned that non-saliva sharers do not provide comfort, but then observed at test that one did. However, because Thomas et al. (2022) measured infants' looking behaviors before any comforting occurred, they interpreted longer looks as "expectations" rather than "surprise." In other words, Thomas et al. inferred that if infants

Table 1 Training events depicting saliva and non-saliva sharers engaging in different actions and their model instantiation.

Event	Model instantiation of Person/Identity		Model instantiation of Action	
(1) Saliva sharers sharing saliva and food	Saliva sharer  Saliva Sharing Left Identity	Saliva sharer  Saliva Sharing Right Identity	Saliva sharing  Saliva Sharing Action	Food item 1 
(2) Saliva sharers sharing something else (e.g., ball)	Saliva sharer  Saliva Sharing Left Identity	Saliva sharer  Saliva Sharing Right Identity	Non-saliva sharing  Saliva Sharing Action	
(3) Saliva sharers providing comfort during distress	Saliva sharer  Saliva Sharing Left Identity	Saliva sharer  Saliva Sharing Right Identity	Comforting  Comforting Action	Distressed  Distressed Action
(4) Non-saliva sharers sharing something else (e.g., ball)	Non-saliva sharer  Saliva Sharing Left Identity	Non-saliva sharer  Saliva Sharing Right Identity	Non-saliva sharing  Saliva Sharing Action	
(5) Non-saliva sharers not providing comfort during distress	Non-saliva sharer  Saliva Sharing Left Identity	Non-saliva sharer  Saliva Sharing Right Identity	Non-comforting  Comforting Action	Distressed  Distressed Action

Note. Colors correspond to network architecture in Figure 1.

Table 2 Training events depicting comforters and non-comforters engaging in different actions and their model instantiation.

Event	Model instantiation of Person/Identity		Model instantiation of Action	
(1) Comforters providing comfort during distress	 Comforting Left Identity	 Comforting Right Identity	 Comforting Action	 Distressed Action
(2) Comforters sharing saliva and food	 Comforting Left Identity	 Comforting Right Identity	 Saliva Sharing Action	 Food item 1
(3) Comforters sharing something else (e.g., ball)	 Comforting Left Identity	 Comforting Right Identity	 Non-saliva sharing Saliva Sharing Action	
(4) Non-comforters not providing comfort during distress	 Comforting Left Identity	 Comforting Right Identity	 Non-comforting Comforting Action	 Distressed Action
(5) Non-comforters sharing something else (e.g., ball)	 Comforting Left Identity	 Comforting Right Identity	 Non-saliva sharing Saliva Sharing Action	

Note. Colors correspond to network architecture in Figure 1.

detected a thick relationship, they would look at the person they expected to comfort the puppet. To convert model error into expectations, we subtracted a given network's error score produced by a test event from the maximum error for that test event. This transformation means that lower network error now corresponds to longer looking or greater expectation for that outcome.

Results

Figure 2 shows networks' expectations of comfort from the Saliva and Non-Saliva Sharers across Conditions and Pretraining Epochs. Familiarization Epoch (or the number of familiarization trials) was not included as an additional effect because it was colinear with Pretraining Epochs and created singularity issues for the model. Data were entered into a linear mixed-effects model with Test Trial as the within-subjects fixed effect (Saliva Sharer vs. Non-saliva Sharer), Condition (Food Sharing, New Puppet, and Saliva Only) and Pretraining Epoch (600 vs. 700 vs. 800) as the between-subjects fixed effects, and participants as a random effect. The results revealed main effects of Test Trial, Pretraining Epochs, and Condition, which were qualified by an interaction between Condition and Pretraining Epochs, $F(4, 351) = 11.06$, $p < .001$. As can be seen in Figure 2, this interaction reflected the fact that networks' expectations for comfort differed across conditions and pretraining epochs. Perhaps most crucially, there was a significant interaction between Condition and Test Trial, $F(2, 351) = 37.46$, $p < .001$. Follow-up analyses were conducted for each Condition to assess the nature of this interaction. Paired t -tests for the Food Sharing and Saliva Only conditions showed that networks, like infants in the corresponding experiments in

Thomas et al. (i.e., Experiments 2B and 3B), showed a greater expectation of comfort from the Saliva Sharer (Food Sharing: $M = 11.41$, $SD = 2.19$; Saliva Only: $M = 11.23$, $SD = 2.24$) than from the Non-Saliva Sharer (Food Sharing: $M = 8.28$, $SD = 2.94$; Saliva Only: $M = 8.20$, $SD = 3.49$), both $t_s > 9.17$, both $p_s < .001$. In contrast, and like infants in the corresponding experiment in Thomas et al. (i.e., Experiment 2D), networks in the New Puppet condition showed an equivalent expectation of comfort from the Saliva Sharer ($M = 2.78$, $SD = 1.17$) than the Non-Saliva Sharer ($M = 2.57$, $SD = 1.09$) when the central character was replaced with someone new, $t(119) = 1.79$, $p = .08$.

Discussion

Simulation 1 suggests that associative learning can reproduce infants' looking patterns in Thomas et al. (2022). However, because these results could reflect a simple actress preference, Thomas et al. (2022) conducted an additional preferential looking test to confirm that infants' looking patterns reflected comfort-specific expectations. During this test, infants in Experiments 2B and 2D saw both actresses smiling and playing peek-a-boo while the central character was absent. Thomas et al. (2022) found that, during this non-distress non-comforting event, infants looked equally long at both actresses, suggesting that neither actress was preferred. Likewise, prior work shows that when agents laugh together rather than engage in comforting, infants do not make inferences about social affiliations (Spokes & Spelke, 2017). Thus, Simulation 2 implements a model analog of a non-distress, non-comfort event to test whether networks' responses in Simulation 1 similarly reflected comfort-specific expectations.




Figure 2 Simulation 1: networks' mean expectations of comfort from the saliva and non-saliva sharers across conditions and pretraining epochs. Note. Food Sharing, New Puppet, and Saliva Only conditions correspond, respectively, to Experiments 2B, 2D, and 3B in [Thomas et al. \(2022\)](#).

Simulation 2: Saliva sharing predicts comfort-specific expectations

Network architecture, pretraining, familiarization, and test

The architecture, pretraining, familiarization, and test phases were identical to the Food Sharing condition of Simulation 1 (corresponding to Experiment 2B in the study by [Thomas et al. \(2022\)](#)) with one key exception: networks were tested on a non-distressed, non-comfort-giving event (i.e., peek-a-boo) rather than a distress-based, comfort-giving event.

Results

[Figure 3](#) shows networks' expectations for the Former Saliva and Non-Saliva Sharers across Pretraining Epochs. Data were entered into a linear mixed-effects model with Test Trial as the within-subjects fixed effect (Former Saliva Sharer vs. Former Non-Saliva Sharer), Pretraining Epochs (600 vs. 700 vs. 800) as between-subjects fixed effects, and participants as the random effect. This analysis revealed a main effect of Pretraining Epochs, which was qualified by an interaction between Test Trial and Pretraining Epochs, $F(2, 117) = 3.72, p = .03$. As can be seen in [Figure 3](#), this interaction reflected the fact that networks' pattern of looking differed across pretraining epochs. Crucially, networks did not show a greater expectation for the Former Saliva Sharer over the Former Non-Saliva Sharer, all $ps > .29$.

Discussion

Simulation 2 tested whether models looked equally long at the former Saliva and Non-Saliva Sharers in a non-distressed, non-comfort-giving event. As was true for infants, networks looked equally long at both actresses. This suggests that networks' performance in the Food Sharing and Saliva Only conditions in Simulation 1 reflected comfort-specific expectations rather than a preference for a specific individual. Together, Simulations 1 and 2 suggest that infants' looking behaviors in [Thomas et al. \(2022\)](#) may be underpinned by an associative-learning mechanism.

Given that the goal of Simulations 1 and 2 was to capture infants' looking behaviors across the infant experiments in [Thomas et al. \(2022\)](#), we tested whether saliva sharing would predict comfort. However, our associative-learning account also implies that comfort will predict saliva sharing. Simulation 3 addressed this issue.

Simulation 3: Comforting predicts saliva sharing

Network architecture, pretraining, familiarization, and test

The architecture and pretraining used in Simulation 2 were identical to those in Simulation 1. Following pretraining, networks were randomly assigned to 1 of 2 conditions (20 networks per condition): (1) Comfort; (2) New Puppet. All networks were




Figure 3 Simulation 2: networks' mean expectations of peek-a-boo from former saliva and former non-saliva sharers across pretraining epochs. Note. The peek-a-boo condition corresponds to a non-distressed, non-comfort-giving event and mirrors the preferential looking test following Experiments 2B and 2D in Thomas et al. (2022).

familiarized with an event in which the central character was comforted by one individual and an event in which the central character was not comforted by a second individual. Networks in the Comfort condition were then tested on an event in which the central character from the familiarization phase engaged in saliva sharing with both the former comforter and former non-comforter. Networks in the New Puppet condition were tested on an identical event, except that the central character was replaced with a new one.

Results

Figure 4 shows networks' expectations of saliva sharing from the Comforter and Non-Comforter across Conditions and Pretraining Epochs. Data were entered into a linear mixed-effects model with Test Trial as the within-subjects fixed effect (Comforter vs. Non-Comforter), Condition (Comfort vs. New Puppet) and Pretraining Epoch (600 vs. 700 vs. 800) as the between-subjects fixed effects, and participants as the random effect. The results revealed main effects for Condition and Test trial, both of which were qualified by an interaction between Condition and Test Trial, $F(1, 234) = 50.64$, $p < .001$. Paired t -tests for the Comfort condition showed that networks had a greater expectation of saliva sharing from the Comforter ($M = 14.22$, $SD = 3.12$) than from the Non-comforter ($M = 10.79$, $SD = 4.73$), $t(119) = 7.73$, $p < .001$. In contrast, networks in the New Puppet condition looked equally long at the Comforter ($M = 2.30$, $SD = 0.92$) and Non-comforter ($M = 2.13$, $SD = 0.98$), $t(119) = 1.69$, $p = .09$.

Discussion

The results of Simulation 3 suggest that an enriched representation supports bidirectional expectations and provides a testable

prediction for future behavioral research. However, behavioral data confirming this prediction would be insufficient to distinguish between associative learning and an evolutionary account, as infants may have an evolved predisposition to interpret both actions as cues to relationship thickness. Thus, to tease these two accounts apart, we examined their diverging *developmental* predictions. If infants' expectations stem from associative learning, we predict that they should increase in strength with experience (Simulation 4) and reflect environmental statistics (Simulation 5). In both simulations, we manipulated networks' pretraining as a proxy for "real-world" experience: Simulation 4 varied the total amount of experience, while Simulation 5 held total experience constant and varied the observed ratio of comfort to non-comfort.

Simulation 4: Comfort-based expectations depend on experience

Network architecture, pretraining, familiarization, and test

Simulation 4 was identical to Simulation 1 except that networks were tested at three different "ages" that reflected different levels of pretraining experience: young, intermediate, and old. The pretraining phase lasted for 100 epochs for young networks, 300 for intermediate networks, and 800 for old networks. This parameter implemented the assumption that infants at different ages necessarily enter experiments with differing amounts of real-world experience. Additionally, the learning rate parameter for young networks was 0.001, 0.01 for intermediate networks, and 0.1 for old networks. This parameter controls the magnitude of the




Figure 4 Simulation 3: networks' mean expectations of saliva sharing from the comforter and non-comforter across conditions and pretraining epochs. Note. In the comfort condition, comforting occurs before saliva sharing; the same central character appears in both events. In the New Puppet condition, comforting occurs before saliva sharing; a different central character appears in each event.

weight changes in the model and corresponds theoretically to the notion that processing speed improves with age. Finally, the number of hidden units was 10 for young networks, 20 for intermediate networks, and 30 for old networks. This parameter captures the idea that infants' processing capacity improves with age. All three parameters implemented a simple model of development that is consistent with existing developmental models (Benton, 2025; Benton et al., 2024; Rakison & Lupyan, 2008).

Results

Figure 5 shows networks' expectations of comfort from the Saliva and Non-Saliva Sharers across Conditions, Age, and Test Trials. Networks showed increasingly greater expectations of comfort from the Saliva Sharer than from the Non-Saliva Sharer with age (see the grey bar in the Food Sharing condition). This conclusion was supported by a linear model in which Age (young vs. intermediate vs. old) was treated as an ordered factor. The outcome measure for this analysis was the difference between the network's expectation of comfort from the Saliva Sharer and that of the Non-Saliva Sharer. This analysis indicated that in the Food Sharing condition, the difference in expectations between the Saliva and Non-saliva Sharers increased monotonically with simulated age, $t(117) = 3.82, p < .001$; networks showed an equivalent expectation for both sharers across ages in the New Puppet condition, $t(117) = -0.18, p = .86$.

Discussion

Simulation 4 suggests that the expectation for comfort following saliva sharing strengthens with experience. Likewise, because the association between saliva sharing and comfort is bidirectional,

our model predicts that the expectation for saliva sharing following comfort should also increase with experience. Although Simulation 4 offers testable predictions for future behavioral research, data showing age-related increases in this expectation would still not distinguish associative learning from an evolutionary "triggering" account, in which experience merely activates a latent, domain-specific mechanism whose primary job is to strengthen a preexisting link between saliva sharing and comfort. Instead, a more definitive test would be to measure whether infants' expectations can be modified or reversed by environmental input—a result unexpected for links formed over evolutionary "deep time" (Buss, 2005), but predicted by associative learning. Simulation 5, therefore, tested whether networks' expectations are graded by the statistical structure of experience rather than just total exposure. To do this, we manipulated whether Saliva Sharers provided comfort consistently, intermittently, or not at all.

Simulation 5: Comfort-based expectations can be reversed

Network architecture, pretraining, familiarization, and test

This simulation was identical to Simulation 1 with one crucial exception: We parametrically manipulated the relative amount of distress-based comfort that networks observed during pretraining. We used the following ratios: 100% comfort vs. 0% non-comfort; 90% comfort vs. 10% non-comfort; 80% comfort vs. 20% non-comfort; 20% comfort vs. 80% non-comfort; 10% comfort vs. 90% non-comfort; and 0% comfort vs. 100% non-comfort. To




Figure 5 Simulation 4: networks' mean expectations of comfort from the saliva and non-saliva sharers across conditions and model age. Note. Food Sharing and New Puppet conditions correspond, respectively, to Experiments 2B and 2D in [Thomas et al. \(2022\)](#).

understand this manipulation, consider the 100% comfort vs. 0% non-comfort ratio: Networks exposed to this ratio always observed comfort from the Saliva Sharer. In contrast, networks exposed to 20% comfort vs. 80% non-comfort observed the Saliva Sharer providing comfort just 20% of the time; they did not observe the Saliva Sharer providing comfort the remaining 80% of the time. These ratios were chosen to form a graded series with mirrored extremes (i.e., 100/0 through 80/20, on the one hand, and 20/80 through 0/100, on the other hand). This enabled us to test whether expectations for comfort and non-comfort are graded by the statistical structure of experience.

Results

[Figure 6](#) shows networks' expectations of comfort from the Saliva and Non-Saliva Sharers across conditions and pretraining ratios. In the Food Sharing condition, expectations of comfort from Saliva Sharers declined steadily as non-comfort became more frequent. In contrast, in the New Puppet condition, networks maintained equivalent expectations across all ratios. A linear model with pretraining ratio as an ordered factor showed that for networks in the Food Sharing condition, the difference in comfort expectations between Saliva and Non-Saliva sharers decreased monotonically with pretraining ratio, $t(714) = -11.44$, $p < .001$; networks showed an equivalent expectation for both sharers across ages in the New Puppet condition, $t(714) = -0.46$, $p = .65$.

Discussion

Consistent with an associative learning account, Simulation 5 showed that expectations can be modified or reversed by environmental input. Taken together, Simulations 1–5 suggest that infants' expectations about comfort can emerge from and be

shaped by a domain-general associative learning mechanism that operates over the statistics of everyday experience and that yields comfort-specific, bidirectional expectations that undergo a developmental progression and can be reversed.

General discussion

[Thomas et al. \(2022\)](#) showed that 8.5- to 10-month-old infants expect an individual who previously shared saliva with a partner to also comfort that partner during times of distress. The goal of the current paper was to offer a mechanistic explanation for these findings and to test whether a connectionist computational instantiation of our account could explain them. Simulation 1 showed that an associative learning mechanism that links saliva sharing to comforting via an enriched representation was sufficient to explain infants' looking behavior in the study by [Thomas et al. \(2022\)](#). Simulation 2 revealed that this expectation is specific to comfort rather than to the identity of a particular individual. Simulation 3 extended these findings by showing that expectations of saliva sharing and comforting could be bidirectional, and Simulations 4 and 5 showed that these expectations are contingent upon the statistics of visual experience. Together, these simulations provide an experience-based, associative-learning account of how preverbal infants come to use saliva as a cue for comfort. In doing so, this work adds to a broader tradition in the developmental cognitive sciences that uses connectionist models to demonstrate that domain-general associative processes are sufficient to account for infant behavior in looking-time tasks ([Benton, 2023, 2025](#); [Benton & Lapan, 2022](#); [Christiansen & Curtin, 1999](#); [Cohen et al., 2002](#); [Munakata et al., 1997](#); [Schlesinger & Casey, 2003](#); [Vong et al., 2024](#)).

Although our model demonstrates that associative learning is sufficient to explain how infants come to link saliva sharing to




Figure 6 Simulation 5: networks' mean expectations of comfort from the saliva and non-saliva sharers across conditions and pretraining ratios. Note. Food Sharing and New Puppet conditions correspond, respectively, to Experiments 2B and 2D in [Thomas et al. \(2022\)](#).

comfort, it does not rule out the possibility that this expectation has an evolutionary basis. Indeed, research in non-human primates shows that food sharing and reciprocity are predictive of kinship and social affiliation (e.g., [Jaeggi & Gurven, 2013](#)). For example, chimpanzees preferentially share food with kin ([Silk et al., 2013](#)), and food sharing in bonobos predicts grooming behavior ([de Waal, 1989](#)). These observations suggest that evolutionary processes may have established the co-occurrence of saliva sharing and comfort. Thus, one way to reconcile the current account with an evolutionary one is to maintain that domain-general associative learning operates on the environmental statistics of the local ecologies that evolution provides. In other words, evolutionary history may help to explain *why* certain cues are informative; our account explains *how* infants come to detect them in the first place.

Although the current simulations demonstrate that associative learning can link saliva sharing to comfort, they do not show that infants' knowledge necessarily derives from this process. To validate this proposed mechanism empirically, behavioral research must also show that these expectations are bidirectional and learned. For instance, an associative account would be supported by evidence that infants expect comfort following saliva sharing and that this expectation strengthens with experience. Future research might also test the effects of experience on infants' expectations about thick relationships by examining children who have experienced neglect or maltreatment. In line with evidence that early caregiving environments fine-tune attention and social cognition ([Pollak & Sinha, 2002](#); [Pollak & Tolley-Schell, 2003](#)), an associative learning account predicts that the inconsistent social cues present in neglectful environments will make children less likely to associate saliva sharing with comfort than children from typical caregiving backgrounds. Evidence of experience-dependent developmental trajectories would provide strong support for an associative account.

More broadly, a key advantage of the current modeling approach is that it can naturally accommodate variation across individuals and environments. For example, in addition to manipulating the statistics of experience, our framework can be extended to model individual differences and cultural variation by altering how quickly the model learns (i.e., the learning rate), how much information it can represent (i.e., the number of hidden units), and how much of that information is forgotten (e.g., weight decay). Ultimately, we hope that this generative framework facilitates dialogue between modeling and behavioral data that generates deeper insights into the mechanisms that shape social cognition.

Moving forward, we expect that integrating findings from behavioral work will necessitate refinements to the model's current architecture and assumptions. For instance, the current model used localist, semantically neutral representations to facilitate model interpretation; however, this approach simplifies the richness of infants' real-world experiences. A natural next step would be to instantiate the current proposal into neural networks that better capture naturalistic social inputs. In addition, the relationship between training epochs and real-world experience remains an open question in the developmental cognitive sciences. Although it could be argued that the model requires a significant amount of training, the goal of our pretraining phase was not to suggest that infants observe hundreds of explicit saliva-sharing events. Rather, it was to ensure that saliva sharing and comfort were reliably correlated, given that it is the presence of such correlations—rather than the sheer amount of experience—that is the theoretical linchpin of the current proposal. Future work might explore the minimum amount of experience required to form social expectations in both infants and models. Notably, Simulation 4 suggests that the link between saliva sharing and comfort can be established in as few as 300 epochs, suggesting

that the structure of the environment, rather than the sheer volume of experience, is central to the development of infants' social expectations.

In sum, by demonstrating that an enriched mental representation that arises from associative learning is sufficient to explain infants' responses in the study by [Thomas et al. \(2022\)](#), we show how abstract social expectations can emerge from simple learning mechanisms and experience. In doing so, our findings underscore the value of computational modeling for elucidating mechanisms of early social reasoning.

Data availability

The modeling scripts and data are available for download at https://osf.io/9mqyn/overview?view_only=c09f319c3fa643d9b-688cb09e0794379.

Author contributions

Deon T. Benton (Conceptualization, Formal analysis, Investigation, Methodology, Visualization, Writing—original draft, Writing—review & editing [equal], Software [lead]) and Justine E. Hoch (Conceptualization, Formal analysis, Investigation, Methodology, Visualization, Writing—original draft, Writing—review & editing [equal])

Funding

This research did not receive funding.

Conflicts of interest

None declared.

Acknowledgments

We would like to thank Veronica Hearne, Rebecca Beaton, Tairan Liu, Jessica Ye, and four anonymous reviewers for helpful comments on earlier versions of this manuscript.

References

- Ainsworth, M. D. S., Bell, S. M., & Stayton, D. F. (1974). Infant-mother attachment and social development: Socialization as a product of reciprocal responsiveness to signals. In M. P. M. Richards (Ed.), *The integration of a child into a social world* (Vol. 316, pp. 99–135). Cambridge University Press. <https://doi.org/10.4324/9780203753873-4>
- Benton, D. T. (2023). ‘But what is the mechanism?’: Demystifying the ever elusive ‘developmental mechanism’. *Infant and Child Development*, 32, e2355. <https://doi.org/10.1002/icd.2355>
- Benton, D. T. (2025). An associative-learning account of how infants learn about causal action in animates and inanimates: A critical reexamination of four classic studies. *Journal of Experimental Psychology: General*, 154, 497–521. <https://doi.org/10.1037/xge0001656>
- Benton, D. T., Kamper, D., Beaton, R. M., & Sobel, D. M. (2024). Don’t throw the associative baby out with the Bayesian bathwater: Children are more associative when reasoning retrospectively
- another person suggest intimacy; two studies of American college students. *European Journal of Social Psychology*, 28, 423–436. [https://doi.org/10.1002/\(SICI\)1099-0992\(199805/06\)28:3%3C423::AID-EJSP874%3E3.0.CO;2-V](https://doi.org/10.1002/(SICI)1099-0992(199805/06)28:3%3C423::AID-EJSP874%3E3.0.CO;2-V)
- Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 104, 686–713. <https://doi.org/10.1037/0033-295X.104.4.686>
- Pollak, S. D., & Sinha, P. (2002). Effects of early experience on children’s recognition of facial displays of emotion. *Developmental Psychology*, 38, 784–791. <https://doi.org/10.1037/0012-1649.38.5.784>
- Pollak, S. D., & Tolley-Schell, S. A. (2003). Selective attention to facial emotion in physically abused children. *Journal of Abnormal Psychology*, 112, 323–338. <https://doi.org/10.1037/0021-843X.112.3.323>

- Rakison, D. H., & Lupyan, G. (2008). *Developing object concepts in infancy: An associative learning perspective*. *Monographs of the Society for Research in Child Development*, 1–127. <https://doi.org/10.1111/j.1540-5834.2008.00454.x>
- Rohde, D. L. (1999). *LENS: The light, efficient network simulator*. School of Computer Science, Carnegie Mellon University.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323, 533–536. <https://doi.org/10.1038/323533a0>
- Schlesinger, M., & Casey, P. (2003). Where infants look when impossible things happen: Simulating and testing a gaze-direction model. *Connection Science*, 15, 271–280. <https://doi.org/10.1080/09540090310001655093>
- Silk, J. B., Brosnan, S. F., Henrich, J., Lambeth, S. P., & Shapiro, S. (2013). Chimpanzees share food for many reasons: the role of kinship, reciprocity, social bonds and harassment on food transfers. *Animal Behaviour*, 85, 941–947. <https://doi.org/10.1016/j.anbehav.2013.02.014>
- Spokes, A. C., & Spelke, E. S. (2017). The cradle of social knowledge: Infants' reasoning about caregiving and affiliation. *Cognition*, 159, 102–116. <https://doi.org/10.1016/j.cognition.2016.11.008>
- Thomas, A. J. (2024). Cognitive representations of social relationships and their developmental origins. *Behavioral and Brain Sciences*, 49, e25. <https://doi.org/10.1017/S0140525X24001328>
- Thomas, A. J., Woo, B., Nettle, D., Spelke, E., & Saxe, R. (2022). Early concepts of intimacy: Young humans use saliva sharing to infer close relationships. *Science*, 375, 311–315. <https://doi.org/10.1126/science.abh1054>
- Vong, W. K., Wang, W., Orhan, A. E., & Lake, B. M. (2024). Grounded language acquisition through the eyes and ears of a single child. *Science*, 383, 504–511. <https://doi.org/10.1126/science.adi1374>