



RESEARCH ARTICLE OPEN ACCESS

Developmental Change in Structure Learning Reflects a Shift From Recency-Based to Relational Prediction

Kate Nussenbaum^{1,2} | Ari E. Kahn^{1,3} | Alice Zhang^{4,5} | Nathaniel D. Daw¹ | Catherine A. Hartley⁴

¹Princeton University, Princeton, New Jersey, USA | ²Boston University, Boston, Massachusetts, USA | ³University of Arizona, Tucson, Arizona, USA | ⁴New York University, New York City, New York, USA | ⁵University of Oxford, Oxford, UK

Correspondence: Kate Nussenbaum (katenuss@bu.edu) | Catherine A. Hartley (cate@nyu.edu)

Received: 6 March 2025 | **Revised:** 18 March 2026 | **Accepted:** 8 May 2026

Keywords: cognitive development | computational modeling | knowledge acquisition | schema formation | structure-learning | successor representation | temporal abstraction

ABSTRACT

Children are adept statistical learners, capable of parsing streams of structured input into meaningful units, but the cognitive processes they engage during learning may differ from those of adults. To date, however, it is unclear how learners of different ages predict upcoming experience when navigating environments with complex structure, as well as how changes in predictive learning mechanisms influence structured knowledge acquisition. To address this question, we tested 106 children, adolescents, and adults, ages 8–22 years, on a predictive learning task, in which they experienced sequences of stimuli with a higher-order temporal structure. After an initial learning phase, participants' explicit knowledge of the relations between stimuli was probed via two additional task measures. We used a recently introduced computational model to characterize participants' response times during learning, and found that all participants relied on simple, recency-based prediction, anticipating that they would encounter stimuli they recently encountered in the past. With increasing age, however, participants demonstrated greater evidence of additionally relying on a more sophisticated learning mechanism, which captured a predictive representation of the conditional relations between stimuli. Though predictive learning changed with age, we found only weak evidence that these changes related to the acquisition of explicit knowledge of the environment. Our results suggest that the learning mechanisms through which people parse continuous streams of experience change with age, influencing their predictions about upcoming events.

1 | Introduction

Across development, experiences in the world unfold in a continuous, temporal stream. This stream of experiences tends to have rich structure that can be exploited to guide predictions about the future and subsequent behavior. Across repetitions of structured streams of input, people form temporal associations between their experiences, grouping them into more abstract event representations or “temporal schemas” (Ghosh and Gilboa 2014; Schapiro et al. 2013; Schapiro et al. 2016). These representations in turn facilitate reasoning, inference (Pudhiviyath et al. 2022, 2020), and value-guided decision-making, enabling the rapid evaluation of

the potential consequences of different actions multiple steps into the future (Kahn and Daw 2025; Momennejad et al. 2017; Russek et al. 2017). A child's mornings, for example, may regularly involve eating breakfast, packing their backpack, getting on the bus, and eventually, sitting in their math class, such that they learn to associate these experiences with each other, and the simple act of eating breakfast elicits anticipation of and preparation for the day's math lesson.

Given the ubiquity of statistical regularities within naturalistic streams of experience—and the consequences of learning those regularities for adaptive behavior—for decades, developmental

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *Developmental Science* published by John Wiley & Sons Ltd.

Summary

- Children, adolescents, and adults all learn to segment continuous streams of structured perceptual input, but they may do so via different learning processes.
- We examined how the learning mechanisms that enable people to predict upcoming experiences change and relate to structured knowledge acquisition across development.
- Computational modeling revealed that in a graph-learning task, younger participants relied on simple, recency-based prediction, while older participants tracked temporal relations between stimuli.
- The extent to which participants engaged in this sophisticated form of predictive learning only weakly related to their knowledge of the task's structure.

researchers have investigated early-emerging sensitivity to structured input (Forest et al. 2023). This research has revealed that infants (Hay et al. 2011; Pelucchi et al. 2009; Saffran et al. 1996) and children (Arciuli and Simpson 2011; Jung et al. 2021) are excellent statistical learners, capable of parsing continuous streams of auditory and visual experiences into meaningful units based on transition probabilities between events. Despite the early emergence of statistical learning and its persistence throughout the lifespan, the mechanisms that underlie the learning of environmental regularities undergo pronounced shifts from childhood to early adulthood. Prior work has revealed that in some contexts, children demonstrate greater sensitivity to “local” task statistics, like the frequency or probability with which different events occur (Janacsek et al. 2012; Nemeth et al. 2013). However, when instructed to learn patterns in stimulus sequences, older participants demonstrate stronger engagement of more controlled, “model-based” learning processes in which they explicitly track both adjacent and nonadjacent *relations* between stimuli (Nemeth et al. 2013).

The statistical learning processes that individuals engage may influence their acquisition of explicit knowledge of their environments. In some learning contexts, children and adults demonstrate comparable knowledge acquisition—for example, in studies using linguistic stimuli, children and adults are similarly able to recognize novel words parsed from continuous, auditory streams (Moreau et al. 2022; Saffran et al. 1997; Shufaniya and Arnon 2018; Smalle and Bogaerts 2024). However, in other contexts, children demonstrate weaker explicit knowledge of the relations between stimuli in their environments than adults (Arciuli and Simpson 2011; Finn et al. 2018; Janacsek et al. 2012; Jung et al. 2021; Nemeth et al. 2013; Potter et al. 2017; Qu et al. 2024; Raviv and Arnon 2018; Schlichting et al. 2017; Shufaniya and Arnon 2018). This difference may be driven in part by stimulus modality, with children tending to demonstrate more robust signatures of knowledge in auditory versus visual learning tasks (Forest et al. 2023; Raviv and Arnon 2018; Shufaniya and Arnon 2018) as well as the complexity of the relations between stimuli—developmental differences in the acquisition of explicit knowledge may be particularly pronounced in environments where predictive relations between events are probabilistic and span multiple timesteps (Pudhiyidath et al. 2020).

Here, we sought to computationally characterize age-related change in the learning processes that children, adolescents, and adults spontaneously engage when responding to visual events with a predictable, temporal structure. Prior work has revealed that in environments defined by latent, temporal regularities, two distinct forms of learning enable people to parse continuous streams of input into distinct contexts and form expectations about events multiple steps into the future. In particular, people can rely on simple, recency-based learning mechanisms or more complex, stimulus-conditional learning (Kahn et al. 2025). In many environments, people experience repeated exposures to the same perceptual inputs, such that states or stimuli that have been encountered in the recent past are likely to be encountered again in the near future (Anderson and Schooler 1991). In these environments, people can rely on recency-based predictive learning mechanisms in which they form expectancies about the future based on the strength of traces of past experience (Reynolds et al. 2007; Zacks et al. 2011). Critically, however, these recency-based predictions are grounded in time and do not involve the formation of associations between different states.

People may also rely on a more complex, state-conditional predictive learning mechanism that supports the explicit forecasting of specific, upcoming experiences (Kahn et al. 2025). In contrast to simpler, recency-based learning, this more complex form of learning tracks dependencies between states, incrementally building up an abstracted world model of their temporal relations—or a “successor representation” (SR) (Garvert et al. 2017; Kahn and Daw 2025; Momennejad et al. 2017; Russek et al. 2021; Stachenfeld et al. 2017)—that can be used to predict future experiences. Evidence from a recent study (Kahn et al. 2025) suggests that adults track these dependencies by maintaining decaying “eligibility traces” of the states they encounter, and incrementally update their beliefs about the future stimuli they lead to over multiple time steps. The maintenance of these traces enables each past state to be associated with, and able to trigger prediction of, each, specific, future state that is later encountered, enabling the learner to form state-specific expectations that span long temporal gaps. As with recency-based learning, state-conditional learning can be harnessed to parse continuous streams of experiences into more discrete temporal contexts, with unexpected states signaling contextual shifts. However, the learned SR can also, in theory, be accessed offline when knowledge of environmental structure is probed or required for inference or choice.

Previous work has suggested that the cognitive mechanisms underlying statistical learning undergo a pronounced shift from childhood to adulthood, characterized by decreasing sensitivity to “raw” event frequencies (Janacsek et al. 2012), and increasing awareness of the structured relations between events (Nemeth et al. 2013). Here, we sought to build on this prior work to provide a computational account of age-related changes in the mechanisms that govern how individuals spontaneously parse their experiences, and to examine whether such changes lead to developmental differences in representations of complex environmental structure. We were interested in how two aspects of learning changed with age. Given past findings that suggest that children sometimes derive weaker knowledge of complex environmental structure through experience (Pudhiyidath et al. 2020), our primary aim was to test the hypothesis that to make

sense of dynamic streams of perceptual input, they may rely on simpler, recency-based predictive learning mechanisms that circumvent the need for a full representation of the relations between stimuli. With increasing age, people may rely on a more complex, state-conditional predictive learning mechanism, which supports the acquisition of explicit, structured knowledge of the environment. As a secondary, more exploratory aim, we also sought to examine how the timescales over which people predict the future vary across development. We hypothesized that children may build expectations about more immediate future events, with predictions increasingly spanning longer temporal horizons into adolescence and adulthood.

To test these hypotheses, we had a large sample of children, adolescents, and adults complete a modified version of a community-structure learning paradigm that has been used in prior work (Karuza et al. 2017; Pudhiyidath et al. 2020; Schapiro et al. 2013). In contrast to a prior developmental study using this task (Pudhiyidath et al. 2020), here, we solicited trial-wise responses during the initial, structure-learning phase, enabling us to exploit a novel, computational model recently used to characterize adult response times in these types of learning tasks (Kahn et al. 2025). The model provides an algorithmic account of how recency-based and state-conditional predictive learning mechanisms influence participants' expectations about upcoming stimuli. In line with our primary hypothesis, we found evidence of a developmental shift in the engagement of predictive learning processes: Whereas younger participants relied primarily on recency-based predictive learning, older participants increasingly engaged state-conditional predictive learning, tracking the complex temporal relations between events. In addition, we found that the temporal horizons over which participants formed predictions increased from childhood to early adulthood. In contrast to our initial hypothesis however, we did not find robust evidence for age-related change in explicit representations of environmental structure, suggesting that multiple learning mechanisms may enable the explicit representation of the coarse temporal features of the environment. Here, by computationally characterizing the dynamics of the learning process, we gained insight into how latent predictive learning mechanisms change with age and shape structured representations of continuous experience.

2 | Methods

2.1 | Participants

One hundred and six participants between the ages of 8 and 22 years ($n = 27$ children (ages 8.2–12.5 years, 11 female, 16 male), $n = 35$ adolescents (ages 13.0–17.9, 16 female, 19 male), and $n = 44$ adults (ages 18.4–22.7 years, 24 female, 18 male, 2 other) completed the experiment online, remotely and asynchronously, and were included in the analyses. While many studies of statistical learning have focused on infants and young children, here we focused on older children, adolescents, and young adults due to prior work indicating pronounced changes in predictive learning mechanisms throughout this developmental period (Janacek et al. 2012; Nemeth et al. 2013; Pudhiyidath et al. 2020) as well as methodological constraints—we thought our task would be too long and challenging for children younger than 8 years old. The same participants were originally recruited to take part

in two separate experiments in different online sessions across different days. We recruited participants until we reached our target sample size of $n = 150$ for our first experiment (Zhang et al. 2026), with the goal of including at least 30 participants per age group in our analyses here (in line with (Pudhiyidath et al. 2020)), after accounting for participant drop-out and exclusions.

An additional 49 participants ($n = 28$ Children, $n = 11$ Adolescents, $n = 10$ Adults) completed the experiment but were excluded from all analyses based on a priori defined exclusion criteria, which included: (a) interacting with their browser window more than 20 times ($n = 1$ Child, $n = 3$ Adolescents, and $n = 3$ Adults), (b) failing to respond on more than 20% learning trials ($n = 10$ Children, $n = 1$ Adolescent, and $n = 3$ Adults; see “task” below), (c) responding inaccurately on more than 25% of learning trials ($n = 16$ Children, $n = 5$ Adolescents, $n = 2$ Adults), (d) responding in under 200 ms on more than 20% of learning trials ($n = 1$ Child), and (e) pausing for more than 30 mins between the first and second learning task block ($n = 2$ Adolescents, $n = 2$ Adults).

As with our previous online studies (Nussenbaum et al. 2020, 2022), participants were recruited via ads on social media (Facebook and Instagram), as well as through word-of-mouth and events around New York City. Prior to entering our database and being invited to complete the study, participants completed a 5-min Zoom call with a researcher to verify their identities. During this call, participants, or a parent or guardian (for participants under 18 years old), were required to show a photo ID on camera, and verify their full name and date of birth (or those of their children).

According to self- or parental report, 39% of the 106 participants included in the final sample were White, 30% were Asian, 18% were Black, and 15% were two or more races. In addition, 18% of participants were Hispanic. All participants resided in the United States at the time of study participation. Of the 106 included participants, 91 reported their annual household income. Of these 91 participants, the annual household income distribution was as follows: < \$20k: 11%, \$20–\$40k: 14%, \$40–\$60k: 14%, \$60–\$80k: 14%, \$80–\$100k: 14%, \$100–\$200k: 22%, \$200–\$300k: 4%, \$300–\$400k: 2%, > \$500k: 2%. For child and adolescent participants, we collected information about the highest education level attained by one parent or guardian: 2% of minor participants had a parent or guardian with no formal educational qualifications, 2% reported receiving a GED, 18% reported receiving a high school diploma, 5% reported receiving a degree from a technical or community college, 39% reported receiving a degree from a four-year college, 23% reported receiving a masters degree, and 13% reported receiving a doctorate.

2.2 | Graph-Learning Task

The graph-learning task was based on those used in previous studies (Pudhiyidath et al. 2020; Schapiro et al. 2013) and comprised three distinct phases: learning, parsing, and graph reconstruction. The task was framed within a child-friendly narrative: Participants were told that they were scientists on board a space shuttle, whose job was to learn about 15 different, visually distinct planets.




FIGURE 1 | Task structure. (A) Unbeknownst to participants, the sequence of planets that they observed in the learning and parsing tasks had a higher-order temporal structure. All planets had an equal probability of transitioning to one of four other neighboring planets (illustrated by the thin gray lines). However, the planets were clustered into three densely interconnected “communities,” illustrated by the colored circles, such that each planet would often be preceded and succeeded by other, within-community planets. (B) In the learning phase of the task, participants viewed a sequence of planets determined by a random walk along the planet graph (depicted in A). On every trial, a vertical or horizontal gabor patch (randomly determined) was overlaid on each planet, and participants had 1.25 seconds to classify the planet’s “wind direction” by pressing “F” or “J”. Participants completed 600 learning trials. (C) In the parsing phase, participants also viewed a sequence of planets determined by a random walk along the planet graph. Here, they had 1.25 seconds to press the spacebar if they believed they had entered a new “galaxy.” Colored boxes are for visualization purposes only; the true planet community was not signaled to participants. (D) In the final, graph reconstruction task, participants saw all 15 planets placed in random locations on the screen. They had as much time as they wanted to click and drag the planets to rearrange them into a map, placing planets that they believed had occurred close together in time near each other in space.

In the learning phase of the task, participants viewed a sequence of planets. Unbeknownst to participants, the sequence of planets had a higher-order temporal structure. Specifically, the sequence they observed was a random walk along a community graph structure ((Schapiro et al. 2013), Figure 1(A)). As in prior work, the graph structure was designed such that each planet had an equal chance of transitioning to four other neighboring planets. However, the planets were arranged in three distinct, highly interconnected “communities” or clusters. This graph structure meant that, despite uniform, single-step transition probabilities, planets would most often be preceded and succeeded by other planets within the same cluster.

During learning (Figure 1(B)), participants completed a cover task that ensured they attended and responded to the planets. Specifically, the planets were overlaid with subtle, shaded gabor patches that were oriented vertically or horizontally. Participants were told that their job as a space scientist was to classify the direction of the wind on each planet, by pressing one of two keys, based on the gabor patch’s orientation. Each planet remained on the screen for 1.25 s, during which participants were able to respond. They did not receive feedback after each trial, but were told that they would earn a bonus payment depending on their number of correct responses. Participants completed two blocks of 300 learning trials, for a total of 600 trials. The specific planet images assigned to each location in the graph were randomized for each participant, and the direction of the gabor patch was randomized on every trial.

In the parsing phase of the task (Figure 1(C)), participants once again saw a sequence of planets determined by a random walk along the same community graph structure. Here, they were told that the planets they previously saw were actually from different galaxies. During parsing, they were tasked with pressing the spacebar whenever they thought they had traveled to a different galaxy. No gabor patches were overlaid on the planets; participants were told the winds had calmed. Here, participants similarly viewed each planet for 1.25 s, during which time they could make a response. Participants completed four blocks of 150 parsing trials,

for a total of 600 trials. Random walks along the graph were truly random for each participant, meaning each participant experienced a different sequence of planets in both the learning and parsing phase of the task.

Finally, in the graph reconstruction phase (Figure 1(D)), participants were tasked with explicitly mapping out how they believed the planets were spread out across the universe. Participants were told that planets they saw next to each other during their “space missions” were close together in the universe, and those that were not seen next to each other were further apart. Participants saw a screen with all 15 planets in random locations and could click and drag on each planet to place it. Participants had unlimited time to arrange the planets on the screen, after which they could press a button to submit their map.

Prior to completing the experiment, participants completed an extensive, interactive tutorial which included child-friendly instructions that were both presented in text and read aloud. Participants were unable to advance past each instruction screen until the audio track had finished playing. Before the learning and parsing phases of the task, participants also had to answer true / false comprehension questions, which repeated until they were answered correctly. To learn the task mechanics, participants also completed a short practice block of the task’s learning phase, with unique practice stimuli.

2.3 | Analysis Approach

2.3.1 | Treatment of Age

Age was treated as a continuous variable in all analyses and modeled linearly. For all reported regressions, we additionally fit a model with a quadratic age term and compared the model with linear age to the model with linear and quadratic age. In no case did the model with quadratic age fit better. Thus, we report only the results from the better-fitting simpler linear age models in the manuscript. For visualization purposes only, participants

were divided into three categorical age groups: Children (ages 8–12 years), Adolescents (ages 13–17 years), and Adults (ages 18–22 years).

2.3.2 | Mixed-Effects Models

We used the “afex” package for R (Singmann et al. 2020) to fit mixed-effects models to our data. All continuous predictor variables were z -scored prior to their inclusion in the models. Models included random intercepts for each participant and random slopes across fixed effects and their interactions (when possible) for each participant. When possible, models also included random intercepts for each planet stimulus, though we pruned planet random intercepts when models failed to converge. We include the full specification for and results from each model in the [Supporting Information](#). For linear mixed-effects models, we assessed the significance of fixed effects with F tests using the Satterthwaite approximation to estimate the degrees of freedom. For logistic mixed-effects models, we assessed the significance of fixed effects with likelihood ratio tests.

For our analysis of participants’ response times during learning, we excluded trials with inaccurate responses, as well as those with responses that were faster than 200 ms. Because we were interested in modeling participant response times as a function of the number of within-block trials they had completed since they had last encountered the same planet, we excluded all trials in which participants encountered a planet for the first time within a block. Response times were log-transformed. For our analysis of participants’ parsing responses, we excluded trials with responses that were faster than 200 ms, as well as the first trial within each block. In addition, we excluded two participants from all parsing analyses: one adult, who made *no* parsing responses throughout the task, and one adolescent, who made more than 530 parsing responses throughout the 600-trial task. We included all other participants, who made between 20 and 470 parsing responses.

2.3.3 | Graph-Reconstruction Behavior

We first confirmed that all participants moved at least one planet to create their final “map” of the planet locations. Participants moved between 6 and 15 planets from their initial random locations ($M = 13.7$ planets, $SD = 1.9$ planets). To analyze how closely participants’ final maps of planets that they created in the graph reconstruction task reflected the true relations between them, we used representational similarity analysis (RSA; (Shepard 1962)) to compare their map with the ground truth. For each participant, we constructed a matrix reflecting the Euclidean distances between all pairs of planets, based on their final placement. To account for individual differences in how participants used the screen, we normalized these distances by dividing all of them by the maximum distance with which they spaced planets. We then compared this matrix to the ground-truth matrix of geodesic distances between planets by computing the Pearson correlation between the lower-triangular elements of the two matrices as an index of how faithfully each participant’s reconstructed planet map captured the true graph structure.

2.3.4 | Computational Model

To characterize the dynamics of participants’ learning, we fit participants’ response times (RTs) from the learning task with a computational model introduced in a recent adult paper (Kahn et al. 2025). Within the computational model, stronger expectations that a planet is likely to appear facilitate faster response times to the gabor patch tilt. Specifically, the model assumes that on every trial, a participant’s response time (rt_t) is drawn from a log-normal distribution with mean μ_t and variance σ^2 , such that:

$$\log(rt_t - shift) \sim N(\mu_t, \sigma^2)$$

where, *shift* reflects the minimum possible RT predicted by the model (fit as a fraction between 0 and 1 of each subject’s minimum RT), and μ_t is a predicted value for trial t , that depends on learned expectations. When the model more strongly expects a particular planet stimulus, μ_t decreases, leading the model to predict faster RTs for more strongly expected stimuli.

The model dynamically learns to “expect” particular planets based on experience. Critically, the model learns via two different mechanisms: A recency-based learning mechanism, and a more complex, state-conditional, ‘successor-representation-based’ learning mechanism.

2.3.5 | Recency-Based Learning Mechanism

In the recency-based learning algorithm, response times are a function of individual stimulus expectancies, such that μ_t is defined as:

$$\mu_t = \beta_R W[s_t]$$

where, β_R is an inverse temperature parameter that reflects the strength with which the recency-weighted expectation term influences RTs and $W[s_t]$ is the recency-weighted expectation of the upcoming stimulus. Values of β_R that are more negative (i.e., further away from zero), indicate that stronger recency-weighted expectations lead to faster reaction times; values of β_R that are close to zero indicate little influence of recency-weighted expectations. In the model, $W[s_t]$ is initialized at 1/15, reflecting equal beliefs about encountering every stimulus before any learning has taken place. On every trial, $W[s_t]$ is updated such that:

$$W[s_t] \leftarrow (1 - \alpha_R)W_t + \alpha_R I(s_t)$$

where, α_R is a learning rate, and I is 1 for the stimulus that was experienced on that trial (s_t) and 0 otherwise. The recency-based model thus maintains expectations about how likely each of the 15 unique stimuli are to appear on any given trial, increasing such expectations each time a particular stimulus is encountered. The learning rate parameter controls the timescale of predictions—high learning rates will strongly increase the expectancy weight on the most recently experienced stimulus, reducing the relative expectation for encountering stimuli that appeared in the more distant past.

2.3.6 | State-Conditional Learning Mechanism

In the state-conditional learning algorithm, response times are a function of individual stimulus expectancies *conditional* on the previously experienced stimulus. In brief, the model maintains decaying memory traces of every unique stimulus it has encountered, and tracks the probability that each of these stimuli will ultimately lead to every other stimuli. The state-conditional model learns these expectancies via a temporal difference λ update rule. Here, μ_t is defined as:

$$\mu_t = \beta_C M [s_{t-1}, s_t]$$

where β_C is an inverse temperature parameter that reflects the strength with which expectations derived from the learned successor representation (SR) influences RTs. As with β_R , values of β_C that are more negative (i.e., further away from zero), indicate that stronger state-conditional expectancies lead to faster reaction times; values of β_C that are close to zero indicate little influence of state-conditional expectancies. Here, $M[s_{t-1}, s_t]$ is the entry in the learned 15x15 SR matrix (M) that reflects the likelihood of transitioning from s_{t-1} to s_t . M was initialized to $(I - \gamma T)^{-1} * T$ for a uniform transition matrix, T, reflecting an initial belief that all transitions between stimuli are equally likely. M is updated on every trial such that:

$$e[s_{t-1}] \leftarrow (1 - \alpha_C) e[s_{t-1}] + 1$$

$$\delta = I [s_t] + \gamma M [s_t, :] - M [s_{t-1}, :]$$

$$\text{for } s' \text{ in } 1 : 15; M [s', :] \leftarrow M [s', :] + \alpha_C e [s'] \delta$$

$$e \leftarrow \gamma \lambda e$$

Here, e is a vector of “eligibility traces” that permits learning over delays (implemented as a Dutch trace (van Seijen et al. 2015)) and α_C is a learning rate parameter that determines how much the eligibility trace is updated after each stimulus encounter. On every trial, δ reflects the difference between the experienced stimulus encounter and the previous trial’s expectancies of future stimuli ($M[s_{t-1}, :]$). Importantly, here, the experienced stimulus encounter takes into account not only the trial’s stimulus ($I[s_t]$), but also the temporally discounted subsequent stimulus encounters ($\gamma I[s_{t+1}], \gamma^2 I[s_{t+2}], \text{etc}$). These are approximated by the model estimate *conditional* on the trial’s stimulus ($\gamma M[s_t, :]$), a process known as bootstrapping. For each of the 15 stimuli, expectations about future stimulus encounters are then updated based on the product of this difference and its eligibility (scaled by a learning rate). Finally, on every trial, eligibility traces are updated based on λ , a trace parameter that determines the rate at which the eligibility of a particular stimulus to predict an upcoming stimulus decays, and γ , a discount factor that determines the rate at which predictions decay, and therefore, the temporal horizon over which future stimuli are predicted. These two parameters imbue the state-conditional learning algorithm with flexibility: The same algorithm can capture a “simpler” learning process through which participants only track a small number of steps of

temporal dependency, or a more complex process through which participants track a longer history of past experience and make predictions further into the future.

2.3.7 | Hybrid Model

Finally, the full hybrid model combines expectations computed by both the recency-based and state-conditional learning algorithms, while controlling for the effects of additional task variables, such that:

$$\begin{aligned} \mu_t = & \mu_0 + \beta_R W [s_t] + \beta_C M [s_{t-1}, s_t] + \beta_{block} b + \beta_{trial} t \\ & + \beta_{stimulus} s + \beta_{response} r \end{aligned}$$

where μ_0 reflects baseline RTs, $\beta_{block} b$ accounts for changes in RTs across learning blocks, $\beta_{trial} t$ accounts for changes in RTs across trials (within each block), $\beta_{stimulus} s$ (where s is a vector of dummy variables for stimulus identity and $\beta_{stimulus}$ a vector of coefficients) accounts for differences in RTs across the 15 planet stimuli, and $\beta_{response} r$ accounts for differences in RTs across keyboard responses corresponding to horizontal and vertical wind directions.

2.3.8 | Model-Fitting Methods

Individual subject parameters were estimated using an expectation-maximization algorithm (Huys et al. 2011) implemented in Julia. Subject-level parameters were modeled as arising from population-level Gaussian distributions over subjects, where each Gaussian distribution was parameterized by its mean and variance. To estimate developmental differences in parameters, we also included an age covariate, which allowed the population-level mean to vary linearly with age.

3 | Results

3.1 | Age-Related Changes in Learning Dynamics

We first examined whether participants’ response times during the initial, learning phase of the task showed evidence of multistep predictive learning, by examining their sensitivity to both planet recency and the planets’ higher-order temporal structure. Though the gabor patch tilts were randomized on every trial and therefore unpredictable, we hypothesized that participants’ expectations of particular planet stimuli would facilitate perception of the overlaid gabor patches and lead to faster responses. Therefore, we expected that participants would make faster responses both when they had encountered the same planet more recently, as well as when they transitioned from one planet to another within-community planet versus when they transitioned between communities. To test this hypothesis, we ran a linear mixed-effects model estimating the effects of age, block, within-block-trial, within-block-trials since the last same-planet encounter (planet lag), transition type, and their interactions on participants’ log-transformed response times. To control for potential differences in the ease of the two motor responses, we further included wind direction as a predictor. Log response times decreased as a function of age, $\beta = -0.05$,

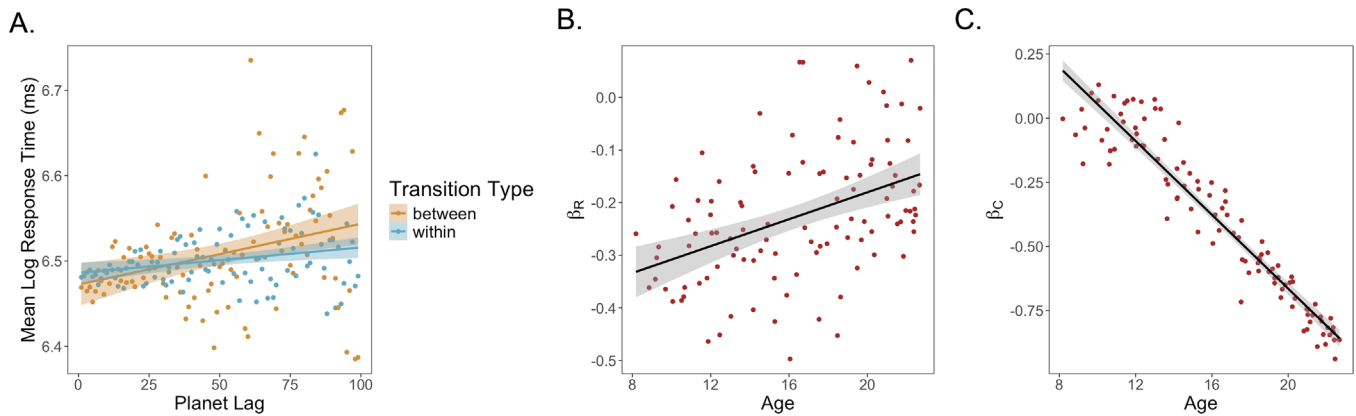


FIGURE 2 | (A) Participants responded faster to planets that they had encountered more recently ($p < 0.001$), an effect that was more pronounced following between-community transitions ($p = 0.001$). (B) Evidence for use of the recency-based learning mechanism did not vary significantly with age ($p = 0.127$), (C) whereas evidence for use of the state-conditional learning mechanism (more negative β_C values) increased across development ($p = 0.001$). Points reflect fitted parameters estimated for each participant via a hierarchical model. The line reflects the best-fitting linear trend line through the points, with the shaded region reflecting 95% confidence intervals.

$SE = 0.01$, $F(1, 106) = 23.3$, $p < 0.001$, and block $\beta = -0.02$, $SE = 0.003$, $F(1, 134) = 44.0$, $p < 0.001$, and were faster for horizontal (right-button) versus vertical (left-button) gabor patches, $\beta = 0.005$, $SE = 0.002$, $F(1, 105) = 11.3$, $p = 0.001$. In addition, participants were significantly slower to respond when the lag between the last same-planet encounter was greater (Figure 2(A)), $\beta = 0.011$, $SE = 0.001$, $F(1, 470.5) = 62.9$, $p < 0.001$, which suggests that they may have engaged in recency-based predictive learning. In contrast to our initial hypothesis, we did not observe a main effect of transition type on response times, $\beta = -0.003$, $SE = 0.002$, $F(1, 114) = 3.4$, $p = 0.069$. However, we observed a transition type x planet lag interaction, $\beta = 0.003$, $SE = 0.001$, $F(1, 51307) = 4.2$, $p = 0.040$, such that participants demonstrated a stronger effect of planet lag after between- versus within-community transitions (Figure 2(A)). In addition, we observed a transition type x trial interaction effect, $\beta = -0.005$, $SE = 0.001$, $F(1, 40364) = 10.4$, $p = 0.001$, such that over the course of each task block, participants' response times decreased to a greater extent for planets following between- versus within-community transitions. We did not observe any other significant effects ($ps > 0.20$).

Taken together, these results indicate that across age, participants made predictions about upcoming planets based on the recency with which they had encountered them. This planet recency effect inherently implies that participants treat between- versus within-community transitions differently—when participants spend time within a community, they encounter the same stimuli repeatedly, leading to shorter planet lags, whereas when they transition between communities, they typically encounter a planet they have not encountered in the recent past. The transition effects in the model thus capture any *additional* influence of higher-order structure that are not already captured by planet lag. Though we did not observe a main effect of transition type, participants' different response time trajectories for between- versus within-community transitions across trials suggest that participants may have gradually acquired sensitivity to the planets' structure over the course of the experiment, perhaps via a separate predictive learning process. In addition, the interaction between planet lag and transition type indicates

that recency had a stronger effect on response times for between-community transitions, suggesting that expectations for these transitions were more sensitive to the immediate past than those for within-community transitions, consistent with the involvement of multiple predictive mechanisms. However, this simple regression model uses a rough proxy for recency (trials since the last same-planet encounter), and therefore cannot provide an algorithmic account of how expectations are learned and updated based on the specific sequence of planets participants experienced across trials. Therefore, to better characterize participants' trial-by-trial updating of planet expectations, we fit our data with a computational model that was developed to explain adult behavior across similar graph-learning tasks (Kahn et al. 2025).

Briefly, within the computational model, stronger expectations that a planet is likely to appear facilitate faster response times. As described in detail in our methods section, the model assumes that participants update their expectations for which planet is likely to appear on any given trial via two learning mechanisms: simpler, recency-based learning of planet probabilities and more complex learning of a “successor representation (SR)” — a world model of temporally discounted predictions of which planets are likely to succeed other planets. The recency-based learning mechanism simply increases the model's expectancy of a particular planet each time that planet is encountered. The nature of the update rule means that the model more strongly “expects” to encounter planets that it has more recently encountered in the past. In contrast, rather than learning simple, global planet expectancies (i.e., a single vector of planet probabilities), the more complex SR-learning mechanism learns conditional planet predictions that depend on the last encountered planet (i.e., a full 15×15 matrix of transition probabilities). It accomplishes this via temporal-difference ($TD(\lambda)$) learning—the model maintains a decaying eligibility trace for each visited planet, which controls the extent to which its “successor” predictions are updated after visits to future planets, even multiple steps in the future.

The two predictive learning mechanisms make different predictions for participants' response times on specific sequential patterns of stimuli (Kahn et al. 2025). In a recency-based account,

participants' response time for the appearance of a target stimulus T should depend on their prior history of encounters with T. Thus, if participants see two sequences: S-A-T-S-T and A-B-T-S-T, where each letter represents a different stimulus, then a recency-based account predicts equivalent response times for the final appearance of T: In both cases, participants would have seen the target stimulus once before, two trials before the trial of interest. However, if participants are relying on an SR-based learning mechanism and see the sequence: S-A-T, then they should learn that S predicts T. Thus, when they see the second appearance of S, they should increase their expectation of encountering T. An SR-based account thus predicts that participants will respond faster to the final target stimulus during the S-A-T-S-T sequence than during the A-B-T-S-T sequence, regardless of how recency influences reaction times.

Within the model, the recency-based learning mechanism and the more complex, state-conditional SR-learning mechanism are weighted by free parameters (β_R, β_C) fit to each participant, to determine the overall expectancies for each planet. In addition, both mechanisms account for individual differences in different aspects of learning affecting the timescales over which participants form and make predictions—the recency-based model includes a learning rate parameter that reflects how much the participant's most recent experience (versus more distant, past experiences) shapes their expectations, while the state-conditional model includes a discount parameter, which reflects how much predictions wane over long temporal distances.

In line with the prior adult study (Kahn et al. 2025), parameter estimation revealed that both β_R and β_C significantly differed from 0 (*mean* $\beta_R = -0.22$, $SE = 0.03$, $p < 0.001$; *mean* $\beta_C = -0.41$, $SE = 0.08$, $p < 0.001$); see [Supporting Information](#) for full model results) meaning both forms of learning contributed to participants' expectations about which planets were likely to appear, as evidenced by the speed with which they responded to them. Interestingly, we found that β_R and β_C followed different trajectories of age-related change. We did not observe a significant effect of age on β_R , ($\beta_{age} = 0.013$, $SE = 0.008$, $p = 0.127$; Figure 2(B)). β_C , however, significantly decreased as a function of age ($\beta_{age} = -0.072$, $SE = 0.014$, $p < 0.001$), indicating that the state-conditional, SR-based learning mechanism more strongly influenced RTs in older participants (Figure 2(C)). Thus, results from the model suggest that recency-based prediction similarly influenced planet expectations across development, but with increasing age, participants' predictions were increasingly shaped by the multistep, conditional relations between stimuli. In line with our regression results, we additionally observed age-related shifts in participants' mean response time parameter, μ_0 ($\beta_{age} = -0.050$, $SE = 0.006$, $p < 0.001$), *shift* parameter ($\beta_{age} = 0.21$, $SE = 0.025$, $p < 0.001$) and σ parameter ($\beta_{age} = 0.023$, $SE = 0.004$, $p < 0.001$), reflecting slower though less variable RTs in younger participants.

The computational model can also provide insight into whether the timescales over which participants predict future events changed systematically with age. Recency-based planet predictions are updated on every planet encounter based on a learning rate (α_R), which determines the extent to which the most recent planet—versus a longer history of planet encounters—influences each participant's expectations. Initially, we hypothesized that

younger participants may demonstrate higher learning rates, such that their predictions are shaped by the more recent past. However, in contrast to our initial hypothesis, we did not observe evidence for significant age-related change in α_R ($\beta_{age} = 0.012$, $SE = 0.017$, $p = 0.477$). The timescale of state-conditional planet predictions are reflected in a different model parameter—the temporal discount parameter γ . Higher γ values reflect reduced temporal discounting, meaning that predictions of upcoming experiences extend further into the future. As such, we expected γ to increase with age. In line with our hypothesis, we observed that γ significantly increased with age ($\beta_{age} = 0.216$, $SE = 0.042$, $p < 0.001$), indicating that not only did older participants more strongly engage state-conditional predictive learning, they also made predictions that extended further into the future.

Finally, we also examined whether the eligibility trace parameter, λ , changed with age. Higher λ values reflect eligibility traces that decay more slowly, such that experiences further in the past are assigned greater “credit” for predicting future events. Here, however, we did not observe evidence for age-related change ($\beta_{age} = 0.080$, $SE = 0.107$, $p = 0.453$).

3.2 | Structure Learning Improves Across Development

We next asked whether, in accordance with prior work (Pudhiyidath et al. 2020), there were age-related changes in participants' representations of the underlying temporal structure of the planets. To do so, we first examined behavior in the parsing task, in which participants responded every time they believed they encountered a planet from a different galaxy. If participants learned the community structure of the graph, then they should be more likely to make a parsing response after transitioning to a planet from a different community versus after transitioning to a different, within-community planet. To test whether participants showed evidence of structured knowledge acquisition, we ran a logistic mixed-effects model estimating the effects of age, trial, transition type, and their interactions on participants' parsing responses. Participants' parsing responses did indeed reflect the graph's community structure — they were more likely to respond when transitioning between versus within communities, $\beta = 0.26$, $SE = 0.04$, $\chi^2(1) = 39.1$, $p < 0.001$. In addition, this effect grew stronger with age, $\beta = 0.11$, $SE = 0.04$, $\chi^2(1) = 8.8$, $p = 0.003$, such that older participants' parsing responses best reflected the task's underlying structure (Figure 3(A)). Participants were also increasingly likely to parse in earlier task blocks, $\beta = -0.08$, $SE = 0.04$, $\chi^2(1) = 4.0$, $p = 0.046$, as well as on earlier trials within each block, $\beta = -0.06$, $SE = 0.02$, $\chi^2(1) = 9.0$, $p = 0.003$, and younger participants made more parsing responses than older participants, $\beta = -0.33$, $SE = 0.10$, $\chi^2(1) = 9.5$, $p = 0.002$.

The sensitivity of participants' parsing responses to between-community versus within-community transitions could reflect explicit, acquired knowledge of the graph structure, derived from SR-based learning, or as was the case in the task's initial learning phase, could reflect “online” or more implicit recency-based expectations. Here, large trial lags between same-planet encounters could have signaled to participants that they were in a

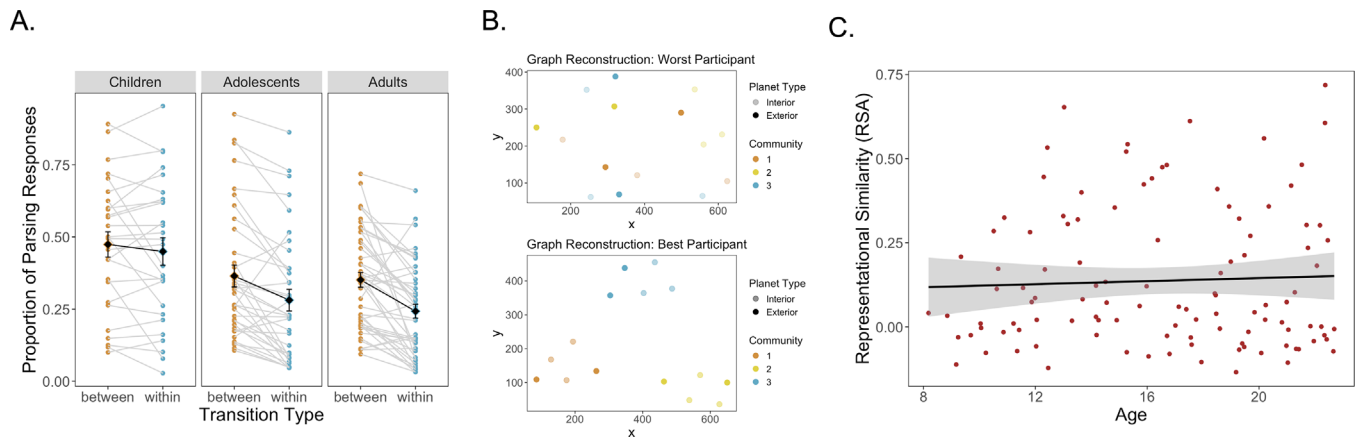


FIGURE 3 | Participants' parsing responses (A) and graph reconstruction behavior (C) indicate that they incidentally acquired knowledge of the planets' community structure during learning. During parsing, participants were more likely to indicate that planets from different communities were from different "galaxies" ($p < 0.001$). With increasing age, participants' parsing responses increasingly reflected the graph's community structure ($p = 0.003$). Colored points indicate individual participant means, while black points show age group averages. Error bars reflect standard errors across participant means. (B) We used RSA to quantify the similarity between participants' reconstructed planet graphs and the planets' true, underlying structure. Figure panels depict the graph reconstruction arrangements of the participant with the lowest (top panel) and highest (bottom panel) RSA scores. (C) On average, participants' RSA scores were above 0 ($p < 0.001$), indicating above chance-level knowledge of the relations between the planets. RSA scores did not vary significantly with age ($p = 0.52$).

new galaxy, eliciting parsing responses, even in the absence of any explicit knowledge of the transition probabilities between planets. To characterize the contributions of recency-based expectations to parsing, we added the planet lag term as an interacting predictor to our parsing model. We observed a main effect of planet lag, $\beta = 0.19$, $SE = 0.02$, $\chi^2(1) = 51.0$, $p < 0.001$, such that participants were more likely to parse when encountering planets that they had not encountered in the recent past. In line with results from our computational model, this term did not interact with age ($p = 0.479$), though it did get weaker across trials within each block, $\beta = -0.04$, $SE = 0.02$, $\chi^2(1) = 5.4$, $p = 0.021$. Importantly, however, even when including the planet lag term, we continued to observe a main effect of transition type on parsing responses, $\beta = 0.19$, $SE = 0.04$, $\chi^2(1) = 23.3$, $p < 0.001$. We also continued to observe a transition type x age interaction, such that older participants' parsing responses better reflected the task's structure, $\beta = 0.11$, $SE = 0.04$, $\chi^2(1) = 7.5$, $p = 0.006$, even when controlling for planet lag.

One possible explanation for the strong, age-related increase in the effect of transition type on parsing is that children simply did not understand the parsing task. Indeed, when we re-ran our parsing model on children only, we did not observe evidence that their parsing behavior was influenced by transition type, $\beta = 0.03$, $SE = 0.07$, $\chi^2(1) = 0.20$, $p = 0.654$, suggesting that they did not differentially recognize that they were in "a new galaxy" after transitioning between planet communities. However, children were not responding randomly: their responses were strongly influenced by planet lag, $\beta = 0.16$, $SE = 0.05$, $\chi^2(1) = 10.6$, $p = 0.001$, indicating that they increasingly believed they were in a new galaxy when they encountered a planet that they had not encountered recently. While the planet lag term only accounts for a linear effect of a relatively simple proxy for recency, together, our parsing findings align with those from our computational model: Children's expectations about upcoming planets were shaped by recent experience, but increasingly with age, participants tracked

the conditional, temporal relations between planets, and used such knowledge to segment streams of continuous experience.

Participants' acquisition of structured knowledge about the relations between planets was also reflected in their final, graph reconstruction behavior (Figure 3(A)). A single-sample t -test revealed that on average, RSA scores reflecting the similarity between participants' final planet maps and the ground-truth geodesic distances between planets ($M = 0.14$, $SD = 0.21$) were significantly above 0, $t(105) = 7.1$, $p < 0.001$, 95% CI = [0.10, 0.18]. This indicates that participants acquired knowledge of the relations between planets, such that the similarity between their maps and their true underlying structure was greater than what would be expected if they had placed the planets randomly. In contrast to our initial hypothesis however, we did not observe evidence that structured knowledge improved with age, $\beta = 0.01$, $SE = 0.02$, $F(1, 104) = 0.43$, $p = 0.52$.

Together, results from the parsing task show that participants learned key features of the graph's structure, while results from the graph reconstruction task demonstrate that through learning, participants acquired structured knowledge they could explicitly report. Results from the parsing task further suggest that, in line with prior work (Pudhiyidath et al. 2020), the extent to which participants learned the true temporal relations between task stimuli increased from childhood to early adulthood.

We hypothesized that participants' increasing reliance on a state-conditional, SR-based learning mechanism across development may provide an account of *why* participants' parsing responses more strongly aligned with the planet's true temporal structure with increasing age. Within the learning model, the recency-based learning mechanism captures only the statistics of recent experience—the representation that it learns is dependent on the current temporal context, and does not reflect the underlying community structure of the graph. The SR-based learning

mechanism, however, maintains and updates a matrix that captures the transitions between all the planets, thereby containing the necessary information to reconstruct the graph's community structure. Thus, it may be the case that the acquisition of this transition matrix during learning facilitated participants' reports of community boundaries during parsing and their ability to cluster within-community planets closer together during graph reconstruction. To test this possibility, we examined how each participant's state-conditional learning weight, as indexed by their fitted β_C parameter from the initial structure-learning phase of the task, related to their parsing and graph reconstruction behavior.

To examine how β_C related to parsing, we ran an additional logistic mixed-effects model examining how age, β_C , planet lag, trial, and transition type influenced parsing responses. We additionally modeled interactions between our subject-level variables (age and β_C) and transition type, planet lag, and trial. We hypothesized that participants who relied more on the state-conditional learning weight (more negative β_C values) would demonstrate a stronger effect of transition type on parsing, and that this interaction may attenuate the influence of our previously observed age \times transition type effect. Contrary to our predictions, however, we did not observe a significant $\beta_C \times$ transition type interaction effect ($p = 0.695$), though when we included β_C in our parsing model, we no longer observed a significant age \times transition type effect ($p = 0.625$). In contrast, the age \times transition type effect on parsing persisted when we included β_R in our model ($p = 0.001$), instead of β_C . These results indicate that age and β_C may account for some of the same variance in the effect of transition type on parsing.

Finally, we did not observe relations between parameters derived from the computational model and graph reconstruction behavior; Neither β_C nor β_R predicted RSA scores ($ps > 0.60$).

4 | Discussion

The ability to predict upcoming events based on past experience underpins many essential cognitive capacities, including reasoning, inference, and decision-making (Kahn and Daw 2025; Momennejad et al. 2017; Morton et al. 2020; Pudhiyidath et al. 2022). However, while children and adults are equally adept statistical learners in many contexts (Forest et al. 2023), recent evidence suggests that in complex task environments, the strategies that individuals use to predict upcoming states and the knowledge that individuals extract from their experiences exhibits pronounced developmental changes (Finn et al. 2018; Janacek et al. 2012; Nemeth et al. 2013; Potter et al. 2017; Pudhiyidath et al. 2020; Qu et al. 2024). Here, we used a community-structure learning task coupled with a novel computational model (Kahn et al. 2025) to gain insight into how predictive learning processes change with age and whether such changes influence the structured representations that people form during learning. We found that across age, people form expectancies about future experiences over multiple timescales, but they do so using different predictive learning mechanisms. Younger children primarily rely on recency-based prediction, which enables online segmentation of continuous streams of perceptual input. However, this form of learning does not build up a representation of the temporal dependencies between different experiences. With increasing age, people demonstrated greater

reliance on a state-conditional, SR-based learning algorithm, in which they predicted upcoming experiences by tracking the likelihood of encountering future stimuli based on each specific stimulus they encountered in the past.

Here, based on the findings from a prior adult study (Kahn et al. 2025), we fit a specific SR-based learning model that proposes that participants learn to track relations between stimuli in the environment via an eligibility trace mechanism. In particular, the model posits that each time a stimulus is encountered, it becomes "eligible" for predicting upcoming stimuli, with the strength of its eligibility decaying across trials. This means that stimuli will be assigned the most predictive "credit" for future stimuli that are encountered in closest temporal proximity. Over the course of learning, this mechanism enables the construction of the full SR matrix, which, for each stimulus, tracks the likelihood of future encounters with all other stimuli. Two parameters centrally govern the nature of this learning process—the eligibility trace parameter λ , which determines the timescale over which stimulus associations are learned, and the discount parameter, γ , which determines the timescale into the future over which learned predictions, ultimately, extend. We observed that γ increased with age, suggesting that across development, individuals make state-conditional predictions further into the future. This finding accords with prior research on statistical learning and decision making that has found that children *are* highly capable of explicitly learning and using one- or two-step stimulus contingencies (Finn et al. 2018; Potter et al. 2017; Raviv and Arnon 2018; Saffran et al. 1997; Zhang et al. 2026), and that age-related changes in knowledge acquisition are more evident in tasks that require tracking contingencies over longer timescales (Pudhiyidath et al. 2020). Future work could bolster our preliminary findings by employing multiple tasks that vary the timescales over which stimuli predict each other, and examining whether age-related changes in online learning and explicit knowledge acquisition are greater in environments with longer temporal lags between predictive and predicted stimuli.

Originally, we hypothesized that age-related changes in predictive learning mechanisms would relate to participants' acquisition of structured knowledge. However, our data did not strongly support this hypothesis. We observed only weak evidence for a relation between participants' initial online learning and their responses to galaxy boundaries in the parsing task. During parsing, older participants were more accurate in explicitly indicating between-community boundaries. The extent to which participants relied on the state-conditional learning mechanism in the prior learning task accounted for some of the age-related variance in the accuracy of these responses. This suggests that, from childhood to early adulthood, people increasingly build detailed, state-conditional representations of their environments that they can access to guide their behavior.

We did not observe age-related change in performance in the graph-reconstruction task, nor did graph reconstruction behavior relate to the use of either predictive learning mechanism. We initially hypothesized that graph reconstruction performance would demonstrate greater age-related change than parsing behavior. Prior work has suggested that age-related change in statistical learning is most apparent when participants must explicitly report relational knowledge (Forest et al. 2023).

Participants could solve the parsing task explicitly, by accessing stored representations of environmental structure that they had built up during learning, *or* more implicitly, by relying on the same predictive learning mechanisms as in the initial learning phase. However, the graph reconstruction task was fully explicit, requiring participants not only to reflect on the temporal patterns of planets they experienced during learning and parsing, but to translate that stream of temporal experience into a spatial map. Across age, participants demonstrated relational knowledge of the planets, and the accuracy of their spatial maps was not related to the extent to which they engaged a state-conditional mechanism in the initial learning phase. This suggests that the online construction of an SR during learning may not be necessary to facilitate accurate map-making; participants could have relied on alternative memory mechanisms to determine the spatial arrangement of the planets. For example, while resting between the different phases of the task, participants could have engaged in “replay”, mentally reactivating sequences of planets that they experienced to construct an SR offline (Momennejad 2020), though recent work (Wittkuhn et al. 2025) suggests reactivation may play a larger role in online SR learning versus the acquisition of explicit knowledge. Alternatively, even without the explicit tracking of the long-range conditional relations between planets, participants may have been able to “chunk” temporally proximal planets together, and access these “chunks” to reconstruct the planet map (Slone and Johnson 2018; Thiessen 2017).

There may also be subtle differences in participants’ explicit representations of the structure of the planets that were not reflected in their graph reconstruction behavior. Recent work has found that by late childhood, children remember events with a shared context as occurring closer together in time, but it is not until early adolescence that they remember events that occur in different contexts as having occurred farther apart in time (Coughlin et al. 2024). Thus, children’s ability to accurately reconstruct the spatial map of planet contexts may have been driven by their ability to cluster within-community planets, whereas older participants may have additionally attempted to separate between-community planets. This may relate to the extent to which they make long-range predictions during learning—younger participants who discount the future to a greater extent and focus only on temporally proximal events may tightly link within-community planets to each other, whereas those who make predictions farther into the future may additionally represent the greater temporal distance between communities. Future work could ask participants to explicitly report the average temporal distance between planets (versus relying on this spatial measure) to test this possibility.

While our modeling findings provide evidence that the construction of a state-conditional, predictive representation during learning increases across development, they leave open the important question of why this is the case. What drives increased weight on the SR-based learning mechanism across development? One possibility is that age-related improvements in working memory (Bunge and Wright 2007; Luna et al. 2015) facilitate tracking of multiple temporal contingencies between stimuli in the environment, though it is unclear how or indeed whether the maintenance of eligibility traces interacts with or depends on the engagement of active working memory processes. Alternatively, people may “rationally” balance recency-based and

state-conditional predictive learning mechanisms, weighing the benefits of carrying out more complex cognitive computations with its costs (Kahn and Daw 2025; Lieder and Griffiths 2019), and the evaluation of these benefits and costs may itself change with age. In our initial learning task, participants were simply instructed to respond rapidly to the unpredictable “wind direction” overlaid on each planet stimulus—they were not told that there *was* structure in the stream of planets they encountered, nor were they told that they would later need to report this structure. Thus, here, there were no obvious benefits to engaging a costlier form of predictive learning, but it is possible that children and adults began the task with different learned prior beliefs about the usefulness of extracting structured knowledge. In other words, across multiple learning experiences over their lifetimes, adults may have developed a stronger bias toward structure learning (Nussenbaum and Hartley 2024), such that they were more likely to engage in this form of learning even when the benefits of doing so were unclear. This interpretation may explain why prior work in which structure learning was explicitly instructed did not observe evidence for age-related change (Nemeth et al. 2013).

The underlying neural mechanisms that support the learning and use of predictive representations like the SR may also change across development. The hippocampus likely plays a central role in tracking environmental statistics (Schapiro et al. 2014, 2016; Stachenfeld et al. 2017)—after structure-learning, the similarity of representations within the hippocampus has been shown to recapitulate the temporal relations between stimuli. With increasing age, the hippocampus has been shown to better represent finer-grained distinctions between stimuli (Keresztes et al. 2018), perhaps facilitating the representations of stimulus relations that span multiple, distinct timescales. Indeed, in young children, greater hippocampal volumes correlated with better statistical learning (Finn et al. 2018), suggesting that age-related changes in hippocampal structure may facilitate increasingly accurate representations of relations between stimuli. Beyond the hippocampus, regions of the prefrontal cortex, including the inferior frontal gyrus and medial prefrontal cortex may play a central role in “reading out” representations of stimulus similarity to identify boundaries between different events or temporal contexts (Schapiro et al. 2013, 2016). Developmental increases in connectivity between the hippocampus and prefrontal cortex (PFC) (Calabro et al. 2019) may thus also support the online use of hippocampal predictive “maps” and explicit access to them offline. By coupling neural measures with computational modeling, future work could directly test whether age-related changes in hippocampal structure and hippocampal-PFC connectivity explain changes in reliance on the SR-based learning mechanism, as well as whether the strength of evidence for reliance on the SR-based learning algorithm relates to the strength with which representations within the hippocampus recapitulate the task’s true, underlying graph structure.

Our finding of age-related increases in state-conditional learning only partially aligns with prior work that has demonstrated increases in explicit relational learning throughout late childhood and early adolescence (Nemeth et al. 2013). Interestingly, this prior work found that age-related improvements in “model-based” statistical learning may occur alongside a corresponding *decrease* in the sensitivity of implicit statistical and procedural learning processes (Janacsek et al. 2012; Juhasz et al.

2019; Nemeth et al. 2013). Though past work has found that children and young adolescents may be particularly sensitive to raw stimulus probabilities (Nemeth et al. 2013), here, we did not observe evidence that sensitivity to “local” task statistics—namely, the recency with which a prior stimulus had been encountered—declined with age. This discrepancy may be due to inherent differences in tracking stimulus frequency versus stimulus recency. Tracking whether an event occurred may be simpler than tracking when an event occurred, and therefore may more easily facilitate learning in younger children.

In addition, due to our task’s difficulty, we had to exclude half of the child participants who completed the study from our analyses, largely due to their failures to provide accurate responses during the learning task within the relatively short response window. While these criteria were necessary for effectively modeling learning, we may have inadvertently selected for children with more “adult-like” cognitive abilities, limiting our ability to detect age-related change in recency-based learning. Future studies could overcome this limitation by reducing the demands of the learning task by making the gabor patches easier to perceive and lengthening the response window. Future work could also vary the incentive for accurate responses and ensure that observed age differences in learning did not emerge from differences in motivation.

In the present study, we begin to gain insight into developmental changes in structure learning by unveiling the hidden algorithms that support online prediction. By using a recently introduced, computational model of learning (Kahn et al. 2025), we moved beyond coarse behavioral response time measures that obscure important learning dynamics to construct a process-level account of learning. The current study thus provides insight into the developmental changes in learning dynamics that may underpin age-related change in the acquisition of structured knowledge, and sets the stage for future work that can ultimately address the fundamental questions of why those changes occur and how they may be shaped by experience.

Author Contributions

Kate Nussenbaum: conceptualization, investigation, writing – original draft, methodology, validation, visualization, writing – review and editing, software, formal analysis, project administration, data curation. **Ari Kahn:** conceptualization, methodology, writing – review and editing, software, formal analysis. **Alice Zhang:** investigation, writing – review and editing, data curation, conceptualization. **Nathaniel D. Daw:** conceptualization, funding acquisition, writing – review and editing, methodology, supervision. **Catherine Hartley:** supervision, conceptualization, funding acquisition, writing – review and editing.

Acknowledgments

We thank Julie Lee for programming an initial version of the graph-learning task and Alana Jaskir for helpful discussion.

Funding

This work was supported by the National Institute of Mental Health (R01MH126183 to C.A.H., R01MH136875 and R01MH135587 to N.D.) and the CV Starr Foundation (Fellowship to K.N.).

Ethics Statement

This study was approved by New York University’s Institutional Review Board (IRB-FY2021-5654). Written informed consent was obtained from all participants over the age of 18 and from parents or legal guardians of child participants. In addition, participants under the age of 18 provided written informed assent prior to participating.

Conflicts of Interest

The authors have no conflicts of interest to disclose.

Data Availability Statement

All task code, anonymized data, and analysis code is available on the Open Science Framework: <https://osf.io/6zwe2/>.

References

- Anderson, J. R., and L. J. Schooler. 1991. “Reflections of the Environment in Memory.” *Psychological Science* 2, no. 6: 396–408. <https://doi.org/10.1111/j.1467-9280.1991.tb00174.x>.
- Arciuli, J., and I. C. Simpson. 2011. “Statistical Learning in Typically Developing Children: The Role of Age and Speed of Stimulus Presentation.” *Developmental Science* 14, no. 3: 464–473. <https://doi.org/10.1111/j.1467-7687.2009.00937.x>.
- Bunge, S. A., and S. B. Wright. 2007. “Neurodevelopmental Changes in Working Memory and Cognitive Control.” *Current Opinion in Neurobiology* 17, no. 2: 243–250. <https://doi.org/10.1016/j.conb.2007.02.005>.
- Calabro, F. J., V. P. Murty, M. Jalbrzikowski, B. Tervo-Clemmens, and B. Luna. 2019. “Development of Hippocampal–Prefrontal Cortex Interactions Through Adolescence.” *Cerebral Cortex* 30, no. 3: 1548–1558. <https://doi.org/10.1093/cercor/bhz186>.
- Coughlin, C., A. Pudhiyidath, H. E. Roome, N. L. Varga, K. V. Nguyen, and A. R. Preston. 2024. “Asynchronous Development of Memory Integration and Differentiation Influences Temporal Memory Organization.” *Developmental Science* 27, no. 2: e13437. <https://doi.org/10.1111/desc.13437>.
- Finn, A. S., M. Kharitonova, N. Holtby, and M. A. Sheridan. 2018. “Prefrontal and Hippocampal Structure Predict Statistical Learning Ability in Early Childhood.” *Journal of Cognitive Neuroscience* 31, no. 1: 126–137. https://doi.org/10.1162/jocn_a_01342.
- Forest, T. A., M. L. Schlichting, K. D. Duncan, and A. S. Finn. 2023. “Changes in Statistical Learning Across Development.” *Nature Reviews Psychology*. <https://doi.org/10.1038/s44159-023-00157-0>.
- Garvert, M. M., R. J. Dolan, and T. E. Behrens. 2017. “A Map of Abstract Relational Knowledge in the Human Hippocampal–Entorhinal Cortex.” *Elife* 6: e17086. <https://doi.org/10.7554/eLife.17086>.
- Ghosh, V. E., and A. Gilboa. 2014. “What is a Memory Schema? A Historical Perspective on Current Neuroscience Literature.” *Neuropsychologia* 53: 104–114. <https://doi.org/10.1016/j.neuropsychologia.2013.11.010>.
- Hay, J. F., B. Pelucchi, K. Graf Estes, and J. R. Saffran. 2011. “Linking Sounds to Meanings: Infant Statistical Learning in a Natural Language.” *Cognitive Psychology* 63, no. 2: 93–106. <https://doi.org/10.1016/j.cogpsych.2011.06.002>.
- Huys, Q. J. M., R. Cools, M. Gölzer, et al. 2011. “Disentangling the Roles of Approach, Activation and Valence in Instrumental and Pavlovian Responding.” *PLoS Computational Biology* 7, no. 4: e1002028. <https://doi.org/10.1371/journal.pcbi.1002028>.
- Janacek, K., J. Fiser, and D. Nemeth. 2012. “The Best Time to Acquire New Skills: Age-Related Differences in Implicit Sequence Learning Across the Human Lifespan.” *Developmental Science* 15, no. 4: 496–505. <https://doi.org/10.1111/j.1467-7687.2012.01150.x>.
- Juhasz, D., D. Nemeth, and K. Janacek. 2019. “Is There More Room to Improve? The Lifespan Trajectory of Procedural Learning and Its Relationship to the Between- and Within-Group Differences in Average

- Response Times.” *PLoS ONE* 14, no. 7: e0215116. <https://doi.org/10.1371/journal.pone.0215116>.
- Jung, Y., D. B. Walther, and A. S. Finn. 2021. “Children Automatically Abstract Categorical Regularities During Statistical Learning.” *Developmental Science* 24, no. 5: e13072. <https://doi.org/10.1111/desc.13072>.
- Kahn, A. E., D. S. Bassett, and N. D. Daw. 2025. “Trial-by-Trial Learning of Successor Representations in Human Behavior.” *PLoS Computational Biology* 21, no. 11: e1013696. <https://doi.org/10.1371/journal.pcbi.1013696>.
- Kahn, A. E., and N. D. Daw. 2025. “Humans Rationally Balance Detailed and Temporally Abstract World Models.” *Communications Psychology* 3, no. 1. <https://doi.org/10.1038/s44271-024-00169-3>.
- Karuza, E. A., A. E. Kahn, S. L. Thompson-Schill, and D. S. Bassett. 2017. “Process Reveals Structure: How a Network is Traversed Mediates Expectations About Its Architecture.” *Scientific Reports* 7, no. 1: 12733. <https://doi.org/10.1038/s41598-017-12876-5>.
- Keresztes, A., C. T. Ngo, U. Lindenberger, M. Werkle-Bergner, and N. S. Newcombe. 2018. “Hippocampal Maturation Drives Memory From Generalization to Specificity.” *Trends in Cognitive Sciences* 22, no. 8: 676–686. <https://doi.org/10.1016/j.tics.2018.05.004>.
- Lieder, F., and T. L. Griffiths. 2019. “Resource-Rational Analysis: Understanding Human Cognition as the Optimal Use of Limited Computational Resources.” *The Behavioral and Brain Sciences* 43: e1. <https://doi.org/10.1017/S0140525x1900061X>.
- Luna, B., S. Marek, B. Larsen, B. Tervo-Clemmens, and R. Chahal. 2015. “An Integrative Model of the Maturation of Cognitive Control.” *Annual Review of Neuroscience* 38: 151–170. <https://doi.org/10.1146/annurev-neuro-071714-034054>.
- Momennejad, I. 2020. “Learning Structures: Predictive Representations, Replay, and Generalization.” *Current Opinion in Behavioral Sciences* 32: 155–166. <https://doi.org/10.1016/j.cobeha.2020.02.017>.
- Momennejad, I., E. M. Russek, J. H. Cheong, M. M. Botvinick, N. D. Daw, and S. J. Gershman. 2017. “The Successor Representation in Human Reinforcement Learning.” *Nature Human Behaviour* 1, no. 9: 680–692. <https://doi.org/10.1038/s41562-017-0180-8>.
- Moreau, C. N., M. F. Joannisse, J. Mulgrew, and L. J. Batterink. 2022. “No Statistical Learning Advantage in Children Over Adults: Evidence From Behaviour and Neural Entrainment.” *Developmental Cognitive Neuroscience* 57: 101154. <https://doi.org/10.1016/j.dcn.2022.101154>.
- Morton, N. W., M. L. Schlichting, and A. R. Preston. 2020. “Representations of Common Event Structure in Medial Temporal Lobe and Frontoparietal Cortex Support Efficient Inference.” *Proceedings of the National Academy of Sciences of the United States of America* 117, no. 47: 29338–29345. <https://doi.org/10.1073/pnas.1912338117>.
- Nemeth, D., K. Janacek, and J. Fiser. 2013. “Age-Dependent and Coordinated Shift in Performance Between Implicit and Explicit Skill Learning.” *Frontiers in Computational Neuroscience* 7: 147. <https://doi.org/10.3389/fncom.2013.00147>.
- Nussenbaum, K., and C. A. Hartley. 2024. “Understanding the Development of Reward Learning Through the Lens of Meta-Learning.” *Nature Reviews Psychology* 3, no. 6: 424–438. <https://doi.org/10.1038/s44159-024-00304-1>.
- Nussenbaum, K., M. Scheuplein, C. V. Phaneuf, M. D. Evans, and C. A. Hartley. 2020. “Moving Developmental Research Online: Comparing In-Lab and Web-Based Studies of Model-Based Reinforcement Learning.” *Collabra* 6, no. 1: 17213. <https://doi.org/10.1525/collabra.17213>.
- Nussenbaum, K., J. A. Velez, B. T. Washington, H. E. Hamling, and C. A. Hartley. 2022. “Flexibility in Valenced Reinforcement Learning Computations Across Development.” 1601–1615. <https://doi.org/10.31234/osf.io/5f9uc>.
- Pelucchi, B., J. F. Hay, and J. R. Saffran. 2009. “Statistical Learning in a Natural Language by 8-Month-Old Infants.” *Child Development* 80, no. 3: 674–685. <https://doi.org/10.1111/j.1467-8624.2009.01290.x>.
- Potter, T. C. S., N. V. Bryce, and C. A. Hartley. 2017. “Cognitive Components Underpinning the Development of Model-Based Learning.” *Developmental Cognitive Neuroscience* 25: 272–280. <https://doi.org/10.1016/j.dcn.2016.10.005>.
- Pudhiyidath, A., N. W. Morton, R. Viveros Duran, et al. 2022. “Representations of Temporal Community Structure in Hippocampus and Precuneus Predict Inductive Reasoning Decisions.” *Journal of Cognitive Neuroscience* 34, no. 10: 1736–1760. https://doi.org/10.1162/jocn_a_01864.
- Pudhiyidath, A., H. E. Roome, C. Coughlin, K. V. Nguyen, and A. R. Preston. 2020. “Developmental Differences in Temporal Schema Acquisition Impact Reasoning Decisions.” *Cognitive Neuropsychology* 37, no. 1–2: 25–45. <https://doi.org/10.1080/02643294.2019.1667316>.
- Qu, Y., J. Ou, L. Pang, et al. 2024. “Development of Non-Spatial Grid-Like Neural Codes Predicts Inference and Intelligence.” *BioRxiv*. <https://doi.org/10.1101/2024.11.20.624569>.
- Raviv, L., and I. Arnon. 2018. “The Developmental Trajectory of Children’s Auditory and Visual Statistical Learning Abilities: Modality-Based Differences in the Effect of Age.” *Developmental Science* 21, no. 4: e12593. <https://doi.org/10.1111/desc.12593>.
- Reynolds, J. R., J. M. Zacks, and T. S. Braver. 2007. “A Computational Model of Event Segmentation From Perceptual Prediction.” *Cognitive Science* 31, no. 4: 613–643. <https://doi.org/10.1080/15326900701399913>.
- Russek, E. M., I. Momennejad, M. M. Botvinick, S. J. Gershman, and N. D. Daw. 2017. “Predictive Representations Can Link Model-Based Reinforcement Learning to Model-Free Mechanisms.” 083857. <https://doi.org/10.1101/083857>.
- Russek, E. M., I. Momennejad, M. M. Botvinick, S. J. Gershman, and N. D. Daw. 2021. “Neural Evidence for the Successor Representation in Choice Evaluation.” *BioRxiv*. <https://doi.org/10.1101/2021.08.29.458114>.
- Saffran, J. R., R. N. Aslin, and E. L. Newport. 1996. “Statistical Learning by 8-Month-Old Infants.” *Science* 274, no. 5294: 1926–1928. <https://doi.org/10.1126/science.274.5294.1926>.
- Saffran, J. R., E. L. Newport, R. N. Aslin, R. A. Tunick, and S. Barrueco. 1997. “Incidental Language Learning: Listening (and Learning) Out of the Corner of Your Ear.” *Psychological Science* 8, no. 2: 101–105. <https://doi.org/10.1111/j.1467-9280.1997.tb00690.x>.
- Schapiro, A. C., E. Gregory, B. Landau, M. McCloskey, and N. B. Turk-Browne. 2014. “The Necessity of the Medial Temporal Lobe for Statistical Learning.” *Journal of Cognitive Neuroscience* 26, no. 8: 1736–1747. https://doi.org/10.1162/jocn_a_00578.
- Schapiro, A. C., T. T. Rogers, N. I. Cordova, N. B. Turk-Browne, and M. M. Botvinick. 2013. “Neural Representations of Events Arise From Temporal Community Structure.” *Nature Neuroscience* 16, no. 4: 486–492. <https://doi.org/10.1038/nn.3331>.
- Schapiro, A. C., N. B. Turk-Browne, K. A. Norman, and M. M. Botvinick. 2016. “Statistical Learning of Temporal Community Structure in the Hippocampus.” *Hippocampus* 26, no. 1: 3–8. <https://doi.org/10.1002/hipo.22523>.
- Schlichting, M. L., K. F. Guarino, A. C. Schapiro, N. B. Turk-Browne, and A. R. Preston. 2017. “Hippocampal Structure Predicts Statistical Learning and Associative Inference Abilities During Development.” *Journal of Cognitive Neuroscience* 29, no. 1: 37–51. https://doi.org/10.1162/jocn_a_01028.
- Shepard, R. N. 1962. “The Analysis of Proximities: Multidimensional Scaling With an Unknown Distance Function. II.” *Psychometrika* 27, no. 3: 219–246. <https://doi.org/10.1007/bf02289621>.
- Shufaniya, A., and I. Arnon. 2018. “Statistical Learning is Not Age-Invariant During Childhood: Performance Improves With Age Across Modality.” *Cognitive Science* 42, no. 8: 3100–3115. <https://doi.org/10.1111/cogs.12692>.
- Singmann, H., B. Bolker, J. Westfall, F. Aust, and M. S. Ben-Shachar. 2020. *Afex: Analysis of Factorial Experiments (Version .27-2)*. Retrieved from <https://CRAN.R-project.org/package=afex>.

Slone, L. K., and S. P. Johnson. 2018. “When Learning Goes Beyond Statistics: Infants Represent Visual Sequences in Terms of Chunks.” *Cognition* 178: 92–102. <https://doi.org/10.1016/j.cognition.2018.05.016>.

Smalle, E. H. M., and L. Bogaerts. 2024. “Sensitive Periods in Language Development: Do Children Outperform Adults on Auditory Word-Form Segmentation?” *Cortex; A Journal Devoted to the Study of the Nervous System and Behavior* 179: 35–49. <https://doi.org/10.1016/j.cortex.2024.07.001>.

Stachenfeld, K. L., M. M. Botvinick, and S. J. Gershman. 2017. “The Hippocampus as a Predictive Map.” *Nature Neuroscience* 20, no. 11: 1643–1653. <https://doi.org/10.1038/nn.4650>.

Thiessen, E. D. 2017. “What’s Statistical About Learning? Insights From Modelling Statistical Learning as a Set of Memory Processes.” *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 372, no. 1711:20160056. <https://doi.org/10.1098/rstb.2016.0056>.

van Seijen, H., A. R. Mahmood, P. M. Pilarski, M. C. Machado, and R. S. Sutton. 2015. *True Online Temporal-Difference Learning*. arXiv. Retrieved from <http://arxiv.org/abs/1512.04087>.

Wittkuhn, L., L. M. Krippner, C. Koch, and N. W. Schuck. 2025. “Replay in the Human Visual Cortex During Brief Task Pauses is Linked to Implicit Learning of Successor Representations.” *Proceedings of the National Academy of Sciences of the United States of America* 122, no. 34: e2507516122. <https://doi.org/10.1073/pnas.2507516122>.

Zacks, J. M., C. A. Kurby, M. L. Eisenberg, and N. Haroutunian. 2011. “Prediction Error Associated With the Perceptual Segmentation of Naturalistic Events.” *Journal of Cognitive Neuroscience* 23, no. 12: 4057–4066. https://doi.org/10.1162/jocn_a_00078.

Zhang, A., A. E. Kahn, N. D. Daw, K. Nussenbaum, and C. A. Hartley. 2026. “Children Leverage Predictive Representations for Flexible, Value-Guided Choice.” *Cognition* 266: 106340. <https://doi.org/10.1016/j.cognition.2025.106340>.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Supporting File 1: desc70227-sup-0001-SuppMat.pdf