Macaques preferentially attend to intermediately surprising information

Shengyi Wu (shengyiwu@berkeley.edu)
Department of Psychology, University of California, Berkeley

Tommy Blanchard (blanchard.tommy@gmail.com)
Klaviyo, 125 Summer St, Floor 6 Boston, MA 02111 USA

Emily Meschke (emily_meschke@berkeley.edu)
Helen Wills Neuroscience Institute, University of California, Berkeley

Richard N. Aslin (richard.aslin@yale.edu)
Haskins Laboratories, Yale University

Benjamin Y. Hayden (benhayden@gmail.com)
Department of Neuroscience, University of Minnesota

Celeste Kidd (celestekidd@berkeley.edu)
Department of Psychology, University of California, Berkeley

Abstract

Normative learning theories dictate that we should preferentially attend to informative sources, but only up to the point that our limited learning systems can process their content. Humans, including infants, show this predicted strategic deployment of attention. Here we demonstrate that rhesus monkeys, much like humans, attend to events of moderate surprisingness over both more and less surprising events. They do this in the absence of any specific goal or contingent reward, indicating that the behavioral pattern is spontaneous. We suggest this U-shaped attentional preference represents an evolutionarily preserved strategy for guiding intelligent organisms toward material that is maximally useful for learning.

Keywords: Attention; statistical learning; eye tracking.

Introduction

Intelligent organisms acquire knowledge through their experiences in the world; however, there is far more information in the world than any explorer could hope to exhaustively explore, much less encode (Oudeyer & Smith, 2016; Wang & Hayden, 2020). Thus, intelligent organisms must have a method for organizing their search for information in the world.

Adaptive theories of curiosity posit that learners' exploration may be guided by their uncertainty in the absence of any explicit reward (Berlyne, 1960; Loewenstein, 1994; Schulz & Bonawitz, 2007; Kang et al., 2009; Bonawitz, van Schijndel, Friel, & Schulz, 2012; Kidd & Hayden, 2015; Dubey & Griffiths, 2020). An inverse-U-shaped curve that correlates attention with information complexity is a key signature of strategic information-seeking. That is, adaptive learners should preferentially attend to information of intermediate complexity—overly simple information offers little to learn from and overly complex information is beyond the learners' capacities to process (Dember & Earl, 1957; Berlyne, 1960; Fantz, 1964; Piaget, 1970; Kinney & Kagan, 1976; Hunter & Ames, 1988; Aslin, 2007; Kidd, Piantadosi, & Aslin, 2012, 2014; Cubit, Canale, Handsman, Kidd, & Bennetto, 2021; Kidd & Hayden, 2015). If intelligent organisms possess this general, curiosity-driven mechanism that guides them towards material that is intermediately surprising, it would prevent them from wasting time on information which is already known, and also information which is unpredictable or overly complex. As a result, the information overload problem is resolved elegantly.

Human infants are drawn to material that is intermediately surprising (Kidd et al., 2012, 2014; Piantadosi, Kidd, & Aslin, 2014), as are older children (Cubit et al., 2021). While this pattern has never been observed outside of humans, there is some evidence that strategic information-seeking applies to monkeys as well when there is no reward tied to a specific task. For example, macaques are willing to sacrifice some amount of liquid reward in exchange for information that has no clear strategic benefit (Blanchard, Hayden, & Bromberg-Martin, 2015; Wang & Hayden, 2019b) and engage in directed exploration (Ebitz, Tu, & Hayden, 2020; Pearson, Hayden, Raghavachari, & Platt, 2009). These data raise the possibility that the strategic information-seeking patterns observed in humans may be shared with other species, possibly reflective of an evolutionarily ancient capacity for adaptive regulation of information gathering. Extending these observations from humans to animals would demonstrate that these are general principles of advanced evolved learners rather than uniquely specialized human skills.

Here, we employ a variation on the infant paradigm with rhesus macaques, in the interest of testing the hypothesis that adaptive regulation of information seeking is a cognitive skill shared with our common ancestor. Unlike most previous work on curiosity in macaques, we employ a free-viewing paradigm without rewards tied to any particular response or behavior. The benefit of this approach is that it tests for spontaneous preference and avoids possible learning effects. We

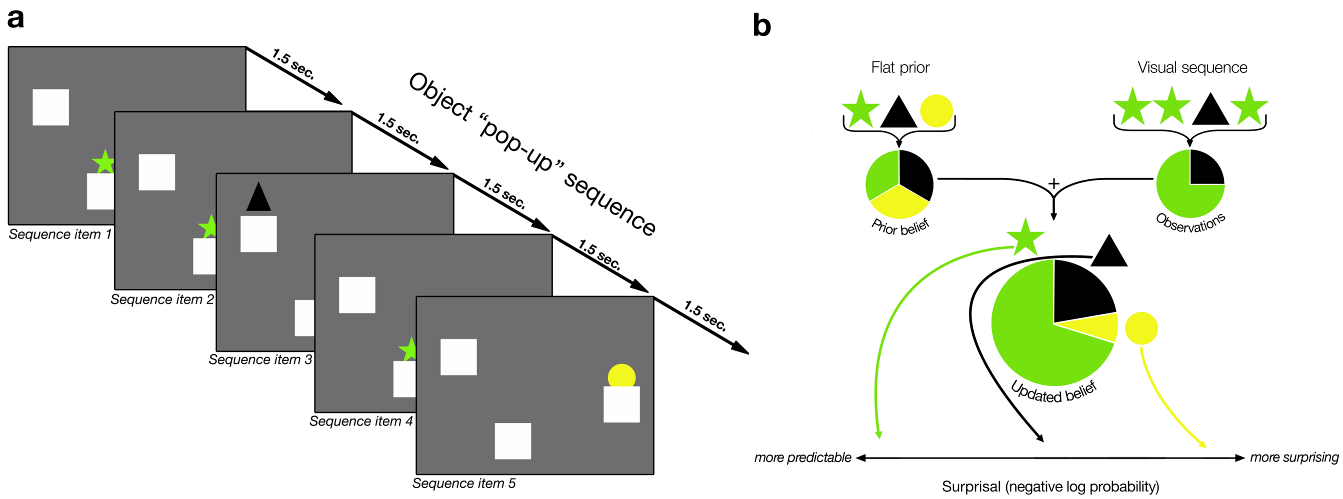

Figure 1: a) Example of Sequential Visual Display. The illustration shows five different time points in a sequence. Each display featured three boxes, each occluding a unique geometric object (e.g., a green star). At each event in the sequence, one of the three objects popped up from behind one of three boxes. b) Idealized Learning Model. The schematic shows an example of how the idealized learning model forms probabilistic expectations about the surprisingness of the next event in a sequence.

find that macaques' visual attention is strikingly similar to that of human infants.

Methods

Subjects

All animal procedures were performed at the University of Rochester (Rochester, NY, USA) and were approved by the University of Rochester Animal Care and Use Committee. All experiments were conducted in compliance with the Public Health Service's Guide for the Care and Use of Animals. Five male rhesus macaques (*Macaca mulatta*) served as subjects [1]. Each subject had a small head-holding prosthesis for collecting high-resolution measurements of eye movements. Subjects had full access to standard chow while in their home cages. Subjects received at minimum 20 mL/kg water per day, although in practice they received close to double this amount in the lab as a result of our experiments. Subjects had been trained to perform oculomotor tasks for liquid rewards through positive-reward-only reinforcement training.

Stimuli

Visual stimuli were colored shapes on a computer monitor (see Figure 1a). Stimuli were controlled by Matlab with Psychtoolbox. Eye positions were measured with Eyelink Toolbox (Cornelissen, Peters, & Palmer, 2002). A solenoid valve controlled the delivery duration of fluid rewards. Eye positions were sampled at 1,000 Hz by an infrared eye-monitoring camera system (SR Research, Osgoode, ON, Canada).

[1] Using single-sex macaques would prevent fighting and overmating opportunities among subjects. We also do not expect sex differences in macaques' behaviors based on the results from the infant version of this study (Kidd et al., 2012).

We designed the displayed stimuli to be easily captured by a simple statistical model (as in Kidd et al. 2012, 2014; Cubit et al., 2021). Each trial featured one of 80 possible visual-event sequences. Sequences were designed to vary in terms of their statistical properties. All sequences were presented to all subjects in a different randomized order. Only one sequence was presented per trial, and each was presented in the form of a unique animated display generated by a Matlab script. An example video can be seen at haydenlab.com/surprisal.

Each animated display featured three identical boxes in three distinct, randomly-chosen spatial locations on the screen that remained static throughout the sequence. Each box concealed one unique geometric object, which was randomly selected from a set of 32 colored geometric shapes including 4 different shapes in 8 colors (e.g., a yellow triangle, a red star, or a blue circle). Geometric objects remained associated with their respective gray boxes throughout the sequence, but were chosen randomly from the set across trials. (following the methods in Kidd et al., 2012, 2014; Piantadosi et al., 2014; Cubit et al., 2021).

Each of the 80 sequences was conveyed by the order in which objects appeared from boxes on the displays. Each event within a sequence consisted of one of the three unique objects popping out from behind one of the three boxes (750 ms), and then back into the box (750 ms), with these pop-ups presented sequentially with no overlap or delay. 32 out of the 80 sequences contained 30 pop-up events, while the rest sequences had 60 pop-up events. The 80 unique sequences were generated to maximize the difference of their theoretical information property, such that the pop-up probabilities of each geometric object varied. Some sequences contained many predictable pop-up events (e.g., ★★★★★) while others contained more unpredictable ones (e.g., ★★▲★●).

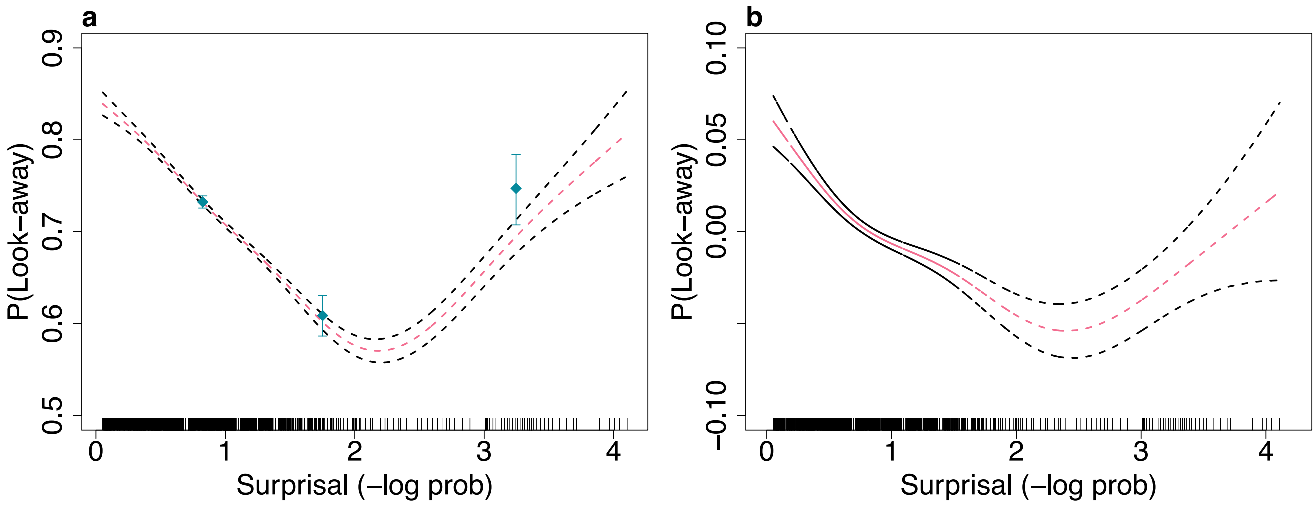

Figure 2: Look-Away Probability as a Function of Unigram Surprisal. a) Subjects probability of looking away (y-axis) as a function of surprisal (x-axis) as measured by the unigram model. The points and error bars show the raw probability of looking away; the smooth curve shows the fit of a generalized additive model with 95% confidence interval. Vertical tick marks show values of surprisal attained in the experiment. b) The relationship between look-away probability (y–axis) and unigram surprisal (x-axis), while controlling for all covariate factors.

Exhaustive randomization and counterbalancing of extraneous variables (e.g., sequence order, object identity, color, shape, spatial location) across trials and subjects served to control for uncertainty about—and variation across—subjects' existing mental representations, processing speeds, and biases for stimulus salience.

Procedure

We recorded eye movements as subjects watched the visual displays. The system delivered a 53 µL water reward when each object was at its peak (every 1.5 sec.), regardless of where or whether the subject was looking. The intermittent and fully predictable reward is a standard procedure in primate behavior studies in which there are no task-specific rewards (See in Azab & Hayden, 2017, 2018). This reliable, non-contingent reward was designed to increase general task participation and arousal without making any particular task events reward-associated. Regardless of subjects' gaze behavior, each sequence (one per trial) was displayed in full. The rate of presentation was between 0 and 2 trials per day, interspersed within between 400 and 2,000 unrelated trials for other studies (Strait, Sleezer, & Hayden, 2015; Blanchard, Wolfe, Vlaev, Winston, & Hayden, 2014; Azab & Hayden, 2018). We recorded all spatial and temporal details of the randomized, Matlab-generated sequential displays we presented, as well as all macaque visual fixations during the stimulus presentations.

Analysis

We computed three dependent behavioral measures from the data: look-aways, reaction times, and predictive-looking. ***Look-aways*** are defined as the first point in the sequence when the macaque looked off-screen for 0.75 sec (50 % of the total pop-up event duration, as in Kidd et al., 2012, 2014 and Cubit et al., 2021). ***Reaction time (RT)*** measures the subjects' latency to shift gaze to the object after it appeared. ***Predictive-looking*** [2] is a binary variable that indicates whether the subject was already looking at the current object when it first became active but before the object actually popped up. We analyzed these three behavioral measures as a function of the surprisal value of each event in the sequence, which is simply the negative log probability of the event's occurrence, according to unigram and bigram Markov Dirichlet-Multinomial (ideal observer) models (again following the analysis methods of Kidd et al. 2012, 2014 and Cubit et al., 2021). The unigram model treats each event as statistically independent, while the bigram model assumes event-order dependence and tracks the conditional probability on the immediately preceding event. The models begin with a simple prior corresponding to the implicit beliefs a learner possesses before beginning to make any observations. By using a flat (or uninformative) prior, we assume that the learner begins the sequence presentation with the implicit belief that each of the three possible objects is equally likely to pop-up from behind their occluding boxes. Once the sequence pre-

[2] In this paradigm, the pop-up events occurred temporally predictable (every 1,500 ms), with no breaks between, and thus a predictive look could not be expected to yield any early information beyond enabling the subject to view the entirety of a new pop-up event, without the costs that would be incurred by attentional switch. Thus, these predictive looks differ from those elicited in most predictive-looking paradigms, where delays between events specifically encourage predictive looking. Regardless, predictive looks may be taken to indicate some degree of attentional allocation to inactive boxes in advance of their opening.

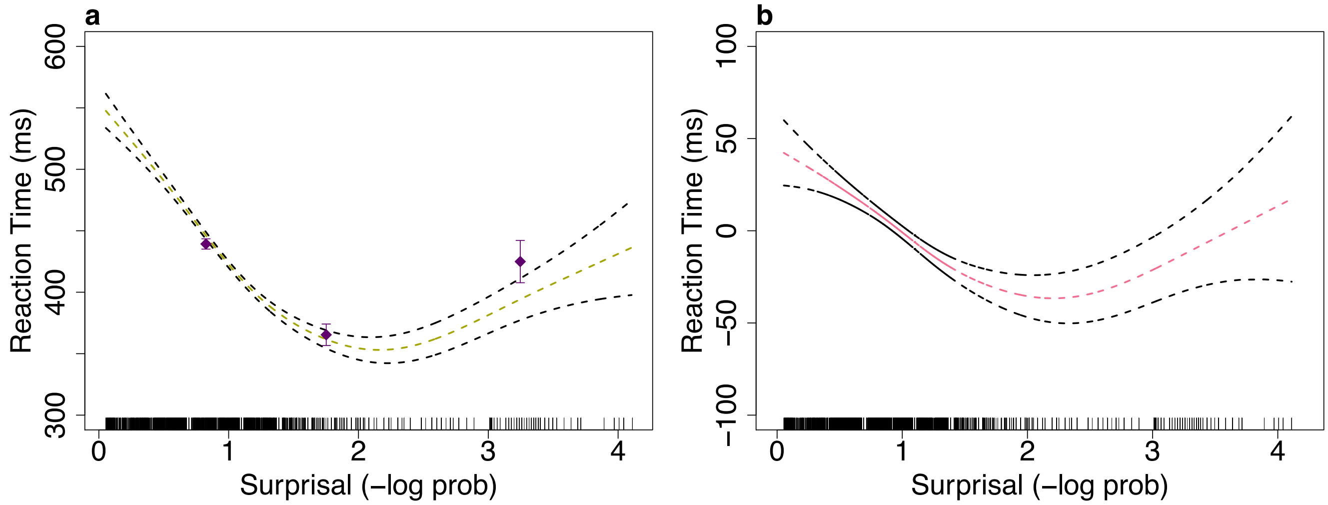

Figure 3: Reaction Time (ms) as a Function of Unigram Surprisal. a) Subjects reaction time (latency) to fixate the active object (y-axis) as a function of surprisal (x-axis) as measured by the unigram model; the smooth curve shows the fit of a generalized additive model with standard errors. b) Reaction Time (y–axis) and unigram surprisal (x-axis), while controlling for all factors.

sentation begins, the model estimates the surprisal value of the current event at each item in the sequence. To do this, it combines the simple prior with the learner's previous observations from the sequence in order to form a posterior or updated belief. The next object pop-up event then conveys the surprisal value according to the probabilistic expectations of the updated belief (see Figure 1b). We also evaluated the statistical significance of each variable using mixed effect linear or logistic regressions with random intercepts, and linear and quadratic surprisal slopes. A generalized additive model (GAM) was used to visualize the relationship between the surprisal estimate from the computational model and the behavioral data.

Results

Preferential gaze towards events of intermediate surprisal

Estimated by the unigram GAM analysis, subjects were more likely to terminate attention to highly predictable (low surprisal) events and also highly unexpected (high surprisal) events (Figure 2a). The GAM's estimated relationship between the unigram surprisal measure and the look-away probability exhibits a clear U-shape. In the regression that considers only surprisal and squared surprisal measures, both the linear term ($\beta = -0.45, z = -5.87, p < 0.001$) and the quadratic term are statistically significant ($\beta = 0.11, z = 3.679, p < 0.001$). The GAM visualization with other covariates being controlled for still shows a clear U-shaped relationship (Figure 2b). The logistic regression reveals a statistically robust linear term ($\beta = -0.16, z = -1.78, p < 0.08$), but not a statistically robust quadratic term ($\beta = 0.03, z = 0.86, p < 0.40$), most likely due to data sparsity in the highest surprisal range (right side) of the U shape. Results from the transitional model shows that there is a U-shaped relationship in the raw model and model fits, with the quadratic trend being statistically significant ($\beta = 0.06, z = 2.69, p < 0.008$). However, this pattern disappears when other variables are controlled for. This suggests that the transitional model does not well-predict subjects' lookaway patterns, in contrast to the robust linear and likely quadratic trends exhibited by the unigram model. Our results also show that all five subjects exhibit preference for stimuli of intermediate surprisal, suggesting that the U-shape relationship holds within rhesus macaques and is not due to subject average. This consistent pattern observed in each macaque subject was also found within individual human infants who reserve attention for events that are moderately predictable (Piantadosi et al., 2014).

Quicker deployment of gaze for events of intermediate surprisal

The unigram GAM analysis shows that the relationship between reaction times and subjects' expectations about stimulus predictability is U-shaped, with subjects exhibiting the fastest RTs for intermediately predictable stimuli (Figure 3a). The regression reveals both a significant linear term ($\beta = -64.68, t = -9.42, p < 0.0002$) and a significant quadratic term ($\beta = 14.372, t = 6.37, p < 0.002$). The relationship holds for each individual subject. The U-shape relationship also holds when other variables are controlled in the GAM model, as well as revealed by the significant linear ($\beta = -26.19, t = -3.02, p < 0.02$) and quadratic ($\beta = 6.00, t = 2.62, p < 0.03$) terms in the controlled regression (Figure 3b). The significance of the quadratic term likely corresponds to a genuine U over the range of surprisal, especially in light of the fact that the significance holds even in the controlled GAM. However, in the GAM analysis for the transitional surprisal

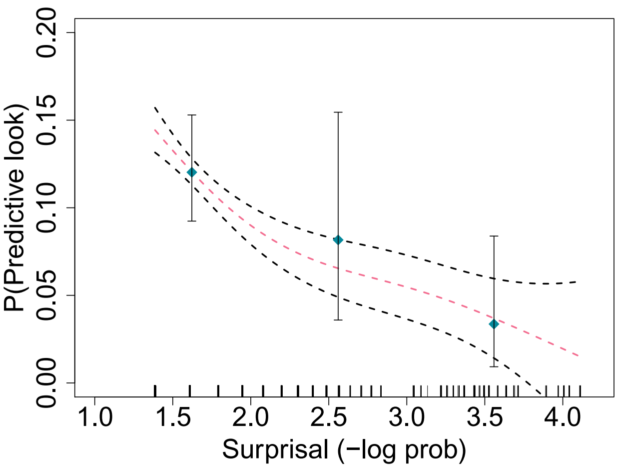

Figure 4: Predictive-Look Probability as a Function of Unigram Surprisal. Subjects probability of predictively looking to the currently active object (y-axis) on their first appearance as a function of surprisal (x-axis) as measured by the unigram model; the smooth curve shows the fit of a generalized additive model with 95% confidence interval.

measures, it shows a much shallower U-shape, with only the linear trend being significant in the raw model. Once all predictors are included, the curve becomes mostly flat. This shows that the unigram model is more robust than the transitional model to capture the relationship between subjects' RTs and the surprisingness of stimuli.

Predictive looks towards unshown items

Subjects are more likely to predictively look at objects on their first appearance when the pop-up events are less surprising. The estimated GAM plot shows a decreasing trend between the probability of predictive-looking and the surprisal value, as in Figure 4. The pattern is supported by the unigram controlled regression model, which finds statistical significance in the linear trend ($\beta = -2.72, z = -2.65, p < 0.009$). Subjects are also more likely to predictively look at never shown objects if they appear earlier in a sequence ($\beta = -12.61, z = -2.13, p < 0.04$). These results show that subjects might be curious about unknown information, expecting that there is some change that will occur and, if it does, it will be informative. They also suggest that over time, as it is increasingly unlikely to see unopened box will ever open, they are less likely to allocate their attentional resources towards monitoring unopened box.

Discussion

Humans do not indiscriminately attempt to absorb any information they encounter. Instead, information-gathering is highly regulated and indeed strategic—we actively seek out information that is maximally useful (Oudeyer & Smith, 2016; Wang & Hayden, 2020; Cervera, Wang, & Hayden, 2020; Kidd & Hayden, 2015). One result of this strategic allocation of information-gathering effort is that we pay special attention to moderately surprising events. Unsurprising events, which provide no additional information beyond what we already understand, do not need constant monitoring to verify their unsurprisingness because sparse sampling is sufficient if there is no change in the underlying statistics. Overly surprising events are also disfavored, likely because they exceed our learning capacity. As a consequence, an inverse-U-shaped preference function is a key signature of the strategic allocation of attention in the service of information gathering. Here we show that this pattern, previously only observed in humans (Kidd et al., 2012, 2014; Piantadosi et al., 2014; Cubit et al., 2021), is also observed in rhesus macaques, a primate species that diverged from humans roughly 25 million years ago.

The presence of this pattern in rhesus macaques suggests that the capacity to adaptively seek maximally useful information is not uniquely human, but instead reflects longstanding evolutionary pressures that have been present since at least the time of our last common ancestor. This is important because a good deal of theorizing highlights the uniqueness of human curiosity, with the implication that curiosity is a factor that has driven human divergence (Berlyne, 1957). Our results, then, suggest an alternative hypothesis that humans and animals share a broad suite of cognitive adaptations, and that humans differ in quantity but not in quality from non-human animals. We found that unigram statistics were a more robust predictor of monkey learners' behaviors than the transitional statistics—in contrast to infants, for whom transitional models outperformed unigram models for both visual (Kidd et al., 2012) and auditory (Kidd et al., 2014) stimuli. This difference may suggest a species-level difference in attentional preferences; however, it is important to note that this conclusion is premature in light of the fact that the macaque participants tested here were in the habit of tasks that require tracking unigram statistics (e.g., k-arm bandit tasks). The apparent species-level difference, thus, could instead be the result of experiential differences across our particular participant pools.

Finally, these results highlight the importance of executive control in curiosity. We do not simply amble around the world gleaning whatever information is available to us; instead we act in concert with our environment, following cues, taking hints, and moving towards stimuli and settings that likely offer us the best chance of getting information. Indeed, we search for information in an optimizing manner. Just as we forage for food, we deliberately seek a rich and balanced diet of information that can drive maximally efficient learning and, ultimately, adaptive fitness. It would be interesting to know how the brain processes underlying these two forms of foraging relate (Hayden, Pearson, & Platt, 2011; Blanchard et al., 2015).

References

Aslin, R. N. (2007). What's in a look? *Developmental Science*, *10*(1), 48–53.

Azab, H., & Hayden, B. Y. (2018). Correlates of economic decisions in the dorsal and subgenual anterior cingulate cortices. *European Journal of Neuroscience*, *47*(8), 979–993.

Berlyne, D. E. (1957). Conflict and information-theory variables as determinants of human perceptual curiosity. *Journal of Experimental Psychology*, *53*(6), 399.

Berlyne, D. E. (1960). Conflict, arousal, and curiosity.

Blanchard, T. C., Hayden, B. Y., & Bromberg-Martin, E. S. (2015). Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron*, *85*(3), 602–614.

Blanchard, T. C., Wolfe, L. S., Vlaev, I., Winston, J. S., & Hayden, B. Y. (2014). Biases in preferences for sequences of outcomes in monkeys. *Cognition*, *130*(3), 289–299.

Bonawitz, E. B., van Schijndel, T. J., Friel, D., & Schulz, L. (2012). Children balance theories and evidence in exploration, explanation, and learning. *Cognitive Psychology*, *64*(4), 215–234.

Cervera, R. L., Wang, M. Z., & Hayden, B. (2020). Curiosity from the perspective of systems neuroscience. `https://doi.org/10.31234/osf.io/znrbf`.

Civan, A., Teller, D. Y., & Palmer, J. (2005). Relations among spontaneous preferences, familiarized preferences, and novelty effects: Measurements with forced-choice techniques. *Infancy*, *7*(2), 111–142.

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The eyelink toolbox: eye tracking with matlab and the psychophysics toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613–617.

Cubit, L. S., Canale, R., Handsman, R., Kidd, C., & Bennetto, L. (2021). Visual attention preference for intermediate predictability in young children. *Child Development*.

Dember, W. N., & Earl, R. W. (1957). Analysis of exploratory, manipulatory, and curiosity behaviors. *Psychological Review*, *64*(2), 91.

Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, *127*(3), 455.

Ebitz, R. B., Tu, J. C., & Hayden, B. Y. (2020). Rules warp feature encoding in decision-making circuits. *PLoS Biology*, *18*(11), e3000951.

Fantz, R. L. (1964). Visual experience in infants: Decreased attention to familiar patterns relative to novel ones. *Science*, *146*(3644), 668–670.

Gruber, M. J., Valji, A., & Ranganath, C. (2019). Curiosity and learning: a neuroscientific perspective. `https://doi.org/10.31234/osf.io/s36e5`.

Hayden, B. Y., Pearson, J. M., & Platt, M. L. (2011). Neuronal basis of sequential foraging decisions in a patchy environment. *Nature Neuroscience*, *14*(7), 933.

Hunter, M. A., & Ames, E. W. (1988). A multifactor model of infant preferences for novel and familiar stimuli. *Advances in Infancy Research*.

Kaldy, Z., Blaser, E. A., & Leslie, A. M. (2006). A new method for calibrating perceptual salience across dimensions in infants: the case of color vs. luminance. *Developmental Science*, *9*(5), 482–489.

Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T.-y., & Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science*, *20*(8), 963–973.

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, *88*(3), 449–460.

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The goldilocks effect: Human infants allocate attention to visual sequences that are neither too simple nor too complex. *PloS One*, *7*(5), e36399.

Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2014). The goldilocks effect in infant auditory attention. *Child Development*, *85*(5), 1795–1804.

Kinney, D. K., & Kagan, J. (1976). Infant attention to auditory discrepancy. *Child Development*, 155–164.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin*, *116*(1), 75.

Oudeyer, P.-Y., & Smith, L. B. (2016). How evolution may work through curiosity-driven developmental process. *Topics in Cognitive Science*, *8*(2), 492–502.

Pearson, J. M., Hayden, B. Y., Raghavachari, S., & Platt, M. L. (2009). Neurons in posterior cingulate cortex signal exploratory decisions in a dynamic multioption choice task. *Current Biology*, *19*(18), 1532–1537.

Pelz, M., & Kidd, C. (2020). The elaboration of exploratory play. *Philosophical Transactions of the Royal Society B*, *375*(1803), 20190503.

Piaget, J. (1970). Structuralism harper & row. *New York*.

Piantadosi, S. T., & Kidd, C. (2016). Extraordinary intelligence and the care of infants. *Proceedings of the National Academy of Sciences*, *113*(25), 6874–6879.

Piantadosi, S. T., Kidd, C., & Aslin, R. (2014). Rich analysis and rational models: Inferring individual behavior from infant looking data. *Developmental Science*, *17*(3), 321–337.

Roder, B. J., Bushnell, E. W., & Sasseville, A. M. (2000). Infants' preferences for familiarity and novelty during the course of visual processing. *Infancy*, *1*(4), 491–507.

Rose, S. A., Gottfried, A. W., Melloy-Carminar, P., & Bridger, W. H. (1982). Familiarity and novelty preferences in infant recognition memory: Implications for information processing. *Developmental Psychology*, *18*(5), 704–713.

Schulz, L. E., & Bonawitz, E. B. (2007). Serious fun: preschoolers engage in more exploratory play when evidence is confounded. *Developmental Psychology*, *43*(4), 1045.

Sokolov, E. N. (1963). Perception and the conditioned reflex.

Strait, C. E., Sleezer, B. J., & Hayden, B. Y. (2015). Signatures of value comparison in ventral striatum neurons. *PLoS Biology*, *13*(6), e1002173.

Wade, S., & Kidd, C. (2018). Cross-cultural differences in the influence of peers on exploration during play. *Cognitive Science*, *42*(8), 3050–3070.

Wade, S., & Kidd, C. (2019). The role of prior knowledge and curiosity in learning. *Psychonomic Bulletin & Review*, *26*(4), 1377–1387.

Wagner, S. H., & Sakovits, L. J. (1986). A process analysis of infant visual and cross-modal recognition memory: Implications for an amodal code. *Advances in Infancy Research*.

Wang, M. Z., & Hayden, B. Y. (2019a). Beyond "incentive hope": Information sampling and learning under reward uncertainty. *Behavioral and Brain Sciences*, *42*.

Wang, M. Z., & Hayden, B. Y. (2019b). Monkeys are curious about counterfactual outcomes. *Cognition*, *189*, 1–10.

Wang, M. Z., & Hayden, B. Y. (2020). Latent learning, cognitive maps, and curiosity. *Current Opinion in Behavioral Sciences*, *38*, 1–7.