

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Habituation reflects optimal exploration over noisy perceptual samples

Permalink

<https://escholarship.org/uc/item/4mh6103w>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Cao, Anjie
Raz, Gal
Saxe, Rebecca
et al.

Publication Date

2022

Peer reviewed

Habituation reflects optimal exploration over noisy perceptual samples

Anjie Cao¹ (anjiecao@stanford.edu), Gal Raz² (galraz@mit.edu),
Rebecca Saxe² (saxe@mit.edu), and Michael C. Frank¹ (mcfrank@stanford.edu)

¹Department of Psychology, Stanford University, ²Department of Brain and Cognitive Sciences, MIT

Abstract

From birth, humans constantly make decisions about what to look at and for how long. Yet the mechanism behind such decision-making remains poorly understood. Here we present the rational action, noisy choice for habituation (RANCH) model. RANCH is a rational learning model that takes noisy perceptual samples from stimuli and makes sampling decisions based on Expected Information Gain (EIG). The model captures key patterns of looking time documented in developmental research: habituation and dishabituation. We evaluated the model with adult looking time collected from a paradigm analogous to the infant habituation paradigm. We compared RANCH with baseline models (no learning model, no perceptual noise model) and models with alternative linking hypotheses (Surprisal, KL divergence). We showed that 1) learning and perceptual noise are critical assumptions of the model, and 2) Surprisal and KL are good proxies for EIG under the current learning context.

Keywords: decision making; learning; bayesian modeling; cognitive development

Introduction

From trying to find our way through a busy street to swiping through TikTok, people are constantly making the decision of whether to keep looking or to look at something else. Even the youngest infants decide whether to keep looking at what is in front of them or move on (Haith, 1980). How can we explain this decision-making process? Our goal in the current paper is to provide a model of the basic decision: whether to keep looking at a stimulus. To do so, we model looking as rational active selection of noisy perceptual samples for learning.

Developmental researchers have long capitalized on infants' ability to control their attention, making inferences about learning and mental representations from changes in looking duration (Baillargeon, Spelke, & Wasserman, 1985; Fantz, 1963). In a typical experiment, infants decrease their looking duration upon seeing the same stimulus repeatedly (habituation) but recover interest when seeing a novel stimulus (dishabituation) (Colombo & Mitchell, 2009). While these phenomena are well-documented, the mechanisms underlying them remain poorly understood, even though assumptions about habituation and dishabituation underpin many other claims about infants' cognitive repertoire (Aslin, 1991; Carey, 2009; Haith, 1998; Sim & Xu, 2019).

One classic theory of infant looking posits that infants look at stimuli in order to learn from them, so the dynamics of

looking time are driven by the dynamics of learning (Hunter & Ames, 1988). This theory describes looking duration as a function of exposure, stimulus complexity, and processing difficulty. The more an infant has already been exposed to a stimulus, the less they have yet to learn about it, but the more complicated the stimulus is, the more they have to learn overall, and older infants are assumed to learn faster than younger infants. Although this theory is influential, its predictions are qualitative, not quantitative, and have not been systematically tested (For exception, see Bergmann & Cristia, 2016; Hunter, Ames, & Koopman, 1983). Nevertheless, it is often invoked as a post-hoc explanation of observed patterns in infant data.

Recent work has attempted to overcome these limitations by describing infants' looking behaviors quantitatively through computational modeling. Kidd, Piantadosi, & Aslin (2012) developed a paradigm in which infants are shown sequences of events until they look away. A rational learning model computed the surprisal (negative log probability) of each event. Infants looked away most from events that were either too high or too low in surprisal (a "goldilocks" effect), suggesting that infants might be looking longer at stimuli with an optimal level of information content.

In this work, surprisal functioned as a quantitative linking hypothesis, connecting between a learning model and data. But other such linking hypotheses are possible. For example, research on information foraging postulates that human exploratory behaviors are driven by maximizing expected information gain (Hills et al., 2015; Pirolli & Card, 1999). In this formulation, agents focus on locations or examples where the amount to be learned is on average highest, a conclusion that is ratified by the emerging literature on curiosity in developmental robotics and reinforcement learning (Oudeyer, Kaplan, & Hafner, 2007). Indeed, EIG provides a good model of curiosity-driven learning in human children and adults (e.g., Liquin, Callaway, & Lombrozo, 2021).

Unfortunately, EIG can be difficult to compute. Because EIG is a measure of expected information gain, its computation requires combinatorial search over all future possibilities. As a result, EIG is often approximated by what is termed "learning progress": the amount of information the agent just learned (Haber, Mrowca, Fei-Fei, & Yamins, 2018). Following this intuition, Poli, Serino, Mars, & Hunnius (2020) formalized learning progress as the Kullback-Leibler (KL) divergence between the model's knowledge before and after

each stimulus and found that a higher KL (i.e. more learning progress) predicted longer looking time in a look-away paradigm.

Thus, existing models take important steps towards a quantitative model of infant attention, but they have three key limitations. First, conceptually, these previous models are not models of choice. The models retrospectively fit infants' overall pattern of attention, without modeling the decision that infants must make in each moment, whether to keep looking at the current stimulus. Second, and relatedly, these models assume that infants acquire a perfect representation of a given stimulus upon exposure. These prior models do not accommodate the noisy nature of perception, and thus cannot explain why infants would have more information to gain from longer looking at the same, already perceived stimulus (Callaway, Rangel, & Griffiths, 2021; Kersten, Mamasian, & Yuille, 2004). Third, because of the first two limitations, the prior models could only directly predict one specific infant looking behavior given an unusual learning problem: looking away from an ongoing stream of stimuli while learning about event probabilities. Substantial further innovation is required to use these models to generate quantitative predictions for infant looking behavior in more standard habituation-dishabituation experimental designs.

Here we attempt to overcome these limitations by providing a model of looking behaviors as arising from optimal decision-making over noisy perceptual representations (Bitzer, Park, Blankenburg, & Kiebel, 2014; Callaway et al., 2021). We present the rational action, noisy choice for habituation (RANCH) model. RANCH works by accumulating noisy samples and choosing at each moment whether to continue to look at the current stimulus or to look away to the rest of the environment. Critically, RANCH allows us to explore a learning problem closer to the problem faced by infants in a standard habituation experiment: instead of assuming a learner is estimating the probability of events, the model learns a category based on the exemplars that are presented during habituation (Oakes, 2010). Furthermore, the architecture allows us to investigate different information-theoretic linking hypotheses as informing choice, including EIG, surprisal, and KL. We make a preliminary evaluation of the RANCH model using adult looking time data collected from a self-paced habituation paradigm that captures habituation, dishabituation, and how these phenomena are modified by stimulus complexity. We begin by presenting our experiment, since it frames the learning task for our model.

Experiment

To reproduce the key looking time patterns from infant habituation experiments in adult participants, we chose a learning context in which participants learn about the stimuli as they look at visually presented exemplars for as long as they like, with no explicit task. The time participants spent exploring the exemplars served as the adult proxy for looking time. This experimental setup resembles the classic infant habituation-

dishabituation paradigm, rather than the look-away paradigm where infants were assumed to learn about event probabilities (Kidd et al., 2012; Poli et al., 2020).

Our initial data come from adults for two reasons. First, adult data are suitable for establishing quantitative links between models and human behaviors, since infants' looking time data tend to have small sample sizes and are therefore limited in their quantitative details (Frank et al., 2017). Second, adult data allow us to test the hypothesis that similar rational choice processes underlie infant and adult behavior under similar learning contexts.

Methods

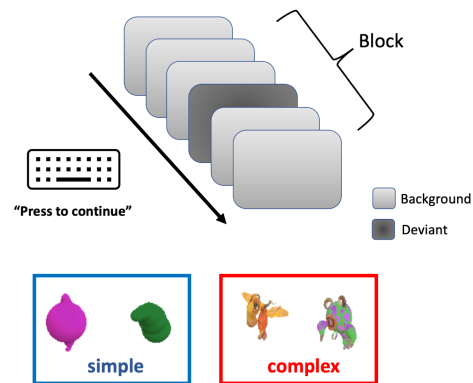


Figure 1: Experimental design and examples of simple and complex stimuli. In each block, a deviant could appear on the second, fourth (as depicted here) or sixth trial or not at all. Stimuli within a block were either all simple or all complex.

Stimuli We created the animated creatures using Spore (a game developed by Maxis in 2008). There were forty creatures in total, half of which had low perceptual complexity and half of which have high perceptual complexity (see Fig. 1 for examples). We used the “animated avatar” function in Spore to capture the creatures in motion.

Procedure The experiment was a web-based, self-paced visual presentation task. Participants were instructed to look at a sequence of animated creatures at their own pace and answer some questions throughout. On each trial, an animated creature showed up on the screen. Participants could press the down arrow to go to the next trial whenever they wanted to, after a minimum viewing time of 500 ms.

Each block consisted of six trials. Unbeknownst to the participants, each trial within the block was either a background trial or a deviant trial. One creature was assigned to be the ‘background’ for each block, and was presented five or six times. If the block contained a deviant trial, then a new, unique, creature was presented on that trial. The deviant trial could appear at either the second, the fourth, or the sixth trial in the block, or not at all. The creatures presented in the deviant trials and background trials were matched for complexity. Each participant saw eight blocks in total, four with

simple creatures and four with complex creatures, in random order across participants.

To test whether behavior was related to task demands, participants were randomly assigned to one of three attention check conditions, differing in the questions asked following each block: Curiosity, Memory, and Math. In the Curiosity condition, participants were asked to rate “How curious are you about the creature?” on a 5-point Likert scale. In the Memory condition, a forced-choice recognition question followed each block (“Have you seen this creature before?” showing either a creature presented in the preceding block or a novel creature matched in complexity). In the Math condition, the participants were asked a simple arithmetic question (“What is $5 + 7$?”) in a multiple-choice format.

To check if our complexity manipulation was successful, at the end of the eight blocks, participants were asked to rate the complexity of creatures they encountered on a 7-point Likert scale.

Participants We recruited 449 participants (Age $M = 30.49$; $SD = 9.74$) on Prolific. They were randomly assigned to one of the three conditions of the experiment. Participants were excluded if they showed irregular reaction times or their responses in the filler tasks indicated low engagement with the experiment. All exclusion criteria were pre-registered. The final sample included 380 participants (Curiosity: $N = 143$; Memory: $N = 98$; Math: $N = 139$).

Results

The sample size, methods, and main analyses were all pre-registered and are available at https://aspredicted.org/3CR_VDR. Data and analysis scripts are available at https://github.com/anjiecao/pokebaby_CogSci2022

We first checked whether the basic complexity manipulations were successful. Complex animated creatures were rated as more perceptually complex ($M = 4.63$; $SD = 1.08$) than the simple animated creatures ($M = 1.06$; $SD = 1.06$; $p < 0.001$).

Next, we tested whether the task (Curiosity, Memory, or Math) affected reaction times in self-paced viewing (our measure of interest). There were no task effects so we averaged all results across three conditions.

We were interested in whether our paradigm successfully captured the characteristic looking time patterns observed in infant literature: habituation (the decrease in looking time for a stimulus with repeated presentations), dishabituation (the increase in looking time to a new stimulus after habituated to one stimulus), and complexity effects (longer looking time for perceptually more complex stimuli). The visualization of our results suggests that we reproduce the phenomena qualitatively (Fig. 3, row 1). To evaluate the phenomena quantitatively, we ran a linear mixed effects model with maximal random effect structure. The predictors included in the model were a three-way interaction term between the trial number (modeled as an exponential decay; Kail, 1991), the type of trial (background vs. deviant) and the complexity of the stim-

uli (simple vs. complex). The model failed to converge, so we pruned the model following the pre-registered procedure. The final model included per-subject random intercepts. All predictors except for the three-way interaction were significant in the model (all $p < .001$), providing a quantitative confirmation that our paradigm successfully captured the key looking time patterns: habituation (trial number), dishabituation (the deviant effect), and complexity (the stimulus complexity effect). We next tested whether we could capture these behavioral results using the RANCH model.

Model

RANCH treats the learning problem that participants face in our experiment as a form of Bayesian concept learning (Goodman, Tenenbaum, Feldman, & Griffiths, 2008; Tenenbaum, 1999). In this setting, multiple noisy samples inform the learner’s hypothesis about a probabilistic concept represented by a set of binary features (Figure 2). Like our participants, the model needs to decide at every step whether to keep looking at the current stimulus or terminate the trial by “looking away.”

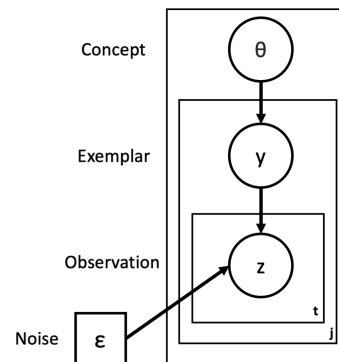


Figure 2: Graphical representation of RANCH. Circles indicate random variables. The square indicates fixed model parameters.

The formulation of the learner as taking noisy samples from a stimulus allows us to do two things. First, we can explicitly model the learner’s decision about when to stop sampling by asking the model to decide, after every sample, whether it wants to continue sampling from the same stimulus or not. This aspect of RANCH contrasts with previous models, which correlate information-theoretic measures to looking data overall (Kidd et al., 2012; Poli et al., 2020) but do not provide a mechanism for how these measures could control moment-to-moment sampling decisions. Second, a consequence of making a decision at every time step is that we can study the behavior of another information-theoretic measure: the model’s expected information gain (EIG). EIG is commonly used in rational analyses of information-seeking behavior to assess whether information-seeking is optimal with respect to the learning task (Markant & Gureckis, 2012; Oaksford & Chater, 1994).

Model definition

In our setting, the goal is to learn a concept θ , which is a set of probabilities over independent binary features $\theta_{1,2,\dots,n}$, where n is the number of features. θ in turn generates exemplars y : instantiations of $\bar{\theta}$, where each feature $y_{1,2,\dots,n}$ is present or absent. The weights on each feature θ_i are sampled from a Beta prior, and individual exemplars y_i are distributed as a binomial with parameter θ_i , forming a conjugate Beta-Bernoulli distribution. Since the features are independent, this relationship holds for the entire concept θ .

To model the timecourse of attention, RANCH does not observe exemplars directly. Instead, it can observe repeated noisy samples \bar{z} from each exemplar. For any sample z from an exemplar y there is a small probability ϵ that the observation is flipped and the feature is seen to be present when it was actually absent or vice versa. ϵ is assumed to be unknown but to have a Beta prior; in practice, we integrate over all possible values of ϵ . Therefore, by making noisy observations \bar{z} , RANCH obtains information about the true identity of the exemplar y , and by extension, about the concept $\bar{\theta}$. By Bayes' rule:

$$P(\theta|\bar{z}) = p(\bar{z}|y)p(y|\theta)p(\theta)/p(\bar{z}) \quad (1)$$

To compute approximate posterior probability distributions during inference, we used a discrete grid approximation with a step size of .001 over both θ and ϵ .

Upon observing a sample, RANCH then decides whether to keep sampling or not. We chose EIG from the next sample as the main linking hypothesis between the learned posterior and sampling choice.

RANCH computes EIG by iterating through each possible next observation and weighing the information gain from each observation by its posterior predictive probability $p(z|\theta)$. We defined information gain as the KL between the hypothetical posterior after observing a future sample z_{t+1} and the current posterior (Baldi & Itti, 2010):

$$EIG(z_{t+1}) = \sum_{z_{t+1} \in [0,1]} p(z_{t+1}|\theta_t) * D_{KL}(\theta_{t+1} || p(\theta_t)) \quad (2)$$

Finally, to get actual sampling behavior from the model, it has to convert EIG into a binary decision about whether to continue looking at the current sample, or to advance to the next trial. The model does so via a Luce choice between the EIG from the next sample and a constant "environmental EIG" that is assumed to be the amount of information to be gained via looking away from the stimulus.

$$p(\text{lookaway}) = \frac{EIG(env)}{EIG(z_{t+1}) + EIG(env)} \quad (3)$$

The basic structure of the model can be described in the following pseudocode:

RANCH model

```

for each exemplar  $y$ 
   $sample \leftarrow T$ 
  while  $sample$  take another sample  $z$ 
    update posterior  $P(\theta|z)$ 
    compute EIG of next sample  $z_{t+1}$ 
    flip  $coin$  with  $p(\text{lookaway}) = \frac{EIG(env)}{EIG(z_{t+1}) + EIG(env)}$ 
    if  $coin = T$ 
       $sample \leftarrow F$ 

```

Simulations

To model the behavioral experiment, we first represented the stimuli as binary-valued vectors indicating the presence (1) or absence (0) of each feature. All stimulus vectors were chosen to be length 6 to provide sufficient representational flexibility. Complex stimuli were represented as having three 1s and simple stimuli were represented as having one 1, with the rest of the features set to 0. Individual stimuli were then assembled into sequences to reflect the stimuli sequences in the behavioral experiment. For a particular sequence, we constructed the deviant stimulus based on the background stimulus to make sure that they were always maximally different and had the same number of features present.

Since the model makes stochastic choices about how many samples to take from each stimulus, behavior varies substantially across runs. Thus, we conducted 500 runs for each stimulus sequence and parameter value to obtain a reasonably precise estimate of the model's behavior.

Parameter estimation

We performed an iterative grid search in parameter space. We a priori constrained our parameter space on the prior beta distribution to have shape parameters $\alpha_\theta > \beta_\theta$, which describe the prior beliefs as "more likely to see the absence of a feature than the presence of a feature." We then searched for the priors over the concept (θ), the noise parameter that decides how likely a feature would be misperceived (ϵ), and the constant EIG from the environment ($EIG(env)$). The prior over the noise parameter was fixed for all searches ($\alpha_\epsilon = 1; \beta_\epsilon = 10$). We selected the parameters that achieved the highest correlation with the behavioral data averaged across participants and blocks ($\alpha_\theta = 1, \beta_\theta = 4, \epsilon = 0.065, EIG(env) = 0.01$). No parameter regimes showed qualitatively different patterns, though the magnitude of dishabituation was strongly dependent on the priors over θ ; a test of the generality of these specific parameter values is left for future work.

Results

RANCH exhibited the main phenomena of interest, showing habituation, dishabituation, and complexity effects (Fig. 3, row 2). We also quantitatively explored the model by fitting the model results to the behavioral data (See Table 1, row 1). Overall RANCH achieved a good fit, though it did show a slightly more gradual habituation process than the behavioral

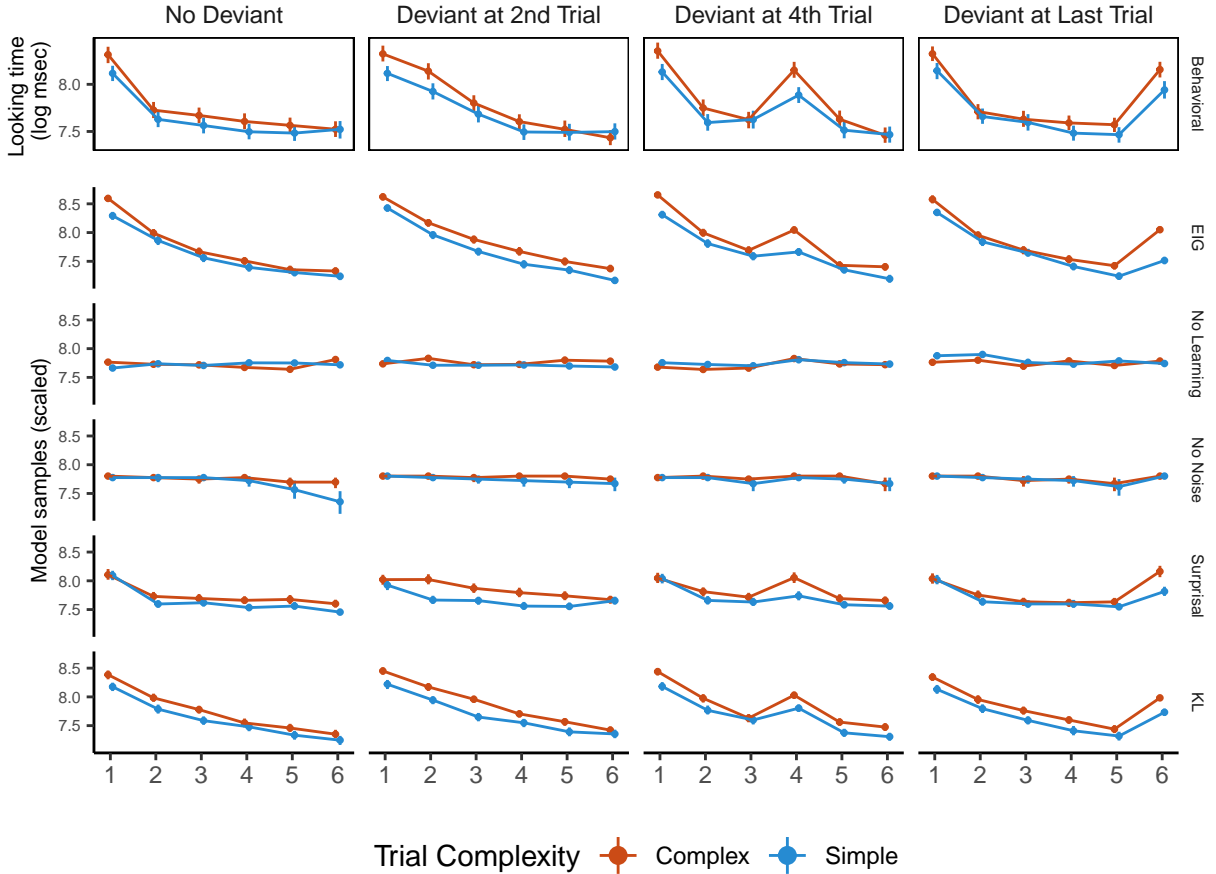


Figure 3: The first row shows behavioral data. All models results were adjusted to match behavioral data scale and intercepts for easier comparisons. All results were log-transformed. Red lines indicate results for complex stimuli, and blue lines indicate results for simple stimuli.

data.

Model Type (Linking Hypothesis)	Pearson's r	RMSE
RANCH (EIG)	0.92 [0.84, 0.96]	0.19 [0.16, 0.24]
Baseline: No Learning	0.21 [-0.09, 0.46]	0.27 [0.23, 0.34]
Baseline: No Noise	0.5 [0.36, 0.65]	0.25 [0.21, 0.31]
RANCH (Surprisal)	0.92 [0.85, 0.95]	0.13 [0.11, 0.16]
RANCH (KL-divergence)	0.93 [0.88, 0.96]	0.12 [0.1, 0.15]

Table 1: This table shows the correlations between the log-transformed model results and the log-transformed looking time data. The values in square brackets are 95% confidence intervals. RANCH model implemented with the three different linking hypotheses showed similar performance with slight numerical differences and outperformed the baseline models.

Alternative Models

Baseline models We next wanted to test what aspects of the model are necessary to produce the phenomena. We focused on two assumptions: 1) the model makes decisions based on learning and 2) perception is noisy. We implemented lesioned baseline models corresponding to each assumption.

The first baseline model (No Learning) made random sampling decisions by drawing $p(\textit{lookaway})$ from a uniform distribution between 0 and 1 at every time step. The second baseline model (No Noise) omitted the noisy sampling aspect of RANCH. We assumed that learning is free from perceptual noise, i.e. that learners can observe the exemplars y directly. To do so, we set ϵ to 0 and replaced the learner's beliefs about the true value of ϵ with the assumption that perception was noiseless (for numerical stability we set the value to 0.000001 instead of 0). The baseline models used the parameters obtained from fitting the EIG model to the behavioral data.

The baseline models fit the data poorly (Table 1, row 2-3; Fig 3, row 3-4), suggesting that both learning and noisy perception are critical for modeling the phenomena of interest. **Alternative linking hypotheses** We also studied the behavior of RANCH using two other linking hypotheses, surprisal and Kullback-Leibler (KL) divergence. Both have been used in previous attempts to model infant looking behavior (Kidd et al., 2012; Poli et al., 2020) and to approximate EIG in the reinforcement learning literature (Kim, Sano, De Freitas, Haber, & Yamins, 2020).

We implemented these by replacing $EIG(z_{t+1})$ in Equa-

tion 2. Surprisal, formally described as $-\log(p(z|\theta))$, intuitively refers to how surprising an observation z is given the model’s beliefs about θ – the intuition that surprising events should result in longer looking times has served as a foundational assumption in developmental psychology (Sim & Xu, 2019). KL, formally described as $\sum_{x \in X} p(\theta = x|z) \log \frac{p(\theta=x|z)}{p(\theta=x)}$, measures how much the model changed to accommodate the most recent observation z . The intuition behind using KL as a linking hypothesis is that, if one observation causes a large change, the next one might too, so continuing to sample is likely to be informative. We re-fit the free parameters (prior, noise and the environmental EIG) for these linking hypotheses to ensure a fair comparison.

In our experiment, the performance of surprisal and KL matched that of EIG (Table 1, row 4-5, Fig 3, row 5-6). To calculate EIG, the model needs to consider all combinations of possible features for the next observation and how informative they would be, a computation that can be intractable in richer environments. The similarity of model fits between EIG, surprisal and KL suggests that easier-to-compute metrics could be viable heuristics for choice behavior, at least in the current learning context.

General discussion

The current work aims to provide a computational model that can explain key phenomena observed in typical infant looking time paradigms: habituation, dishabituation, and how these are modified by stimulus complexity. RANCH assumes a rational learner that takes noisy perceptual samples from stimuli and makes sampling decisions based on EIG. We evaluated the model with adult looking time data collected from a paradigm that mirrors classic infant looking time paradigms, in which participants are learning about multi-feature concepts, and found that RANCH could successfully reproduce the patterns observed in behavioral data. By contrasting the model results with our baseline models, we showed that habituation, dishabituation, and complexity effects only arise in a learning model that takes into account the noisy nature of perception. Moreover, we found that, in the current learning context, other information theoretic quantities (surprisal and KL) are good proxies for the optimal linking hypothesis, EIG.

RANCH constitutes a significant step forward in the modeling of looking time in that it models the moment-to-moment decision making process of whether to keep sampling or look away. Previous approaches incremented time in steps of whole stimuli and therefore correlated information-theoretic variability in the stimulus sequence to look-away probability and looking time, rather than producing these behaviors endogenously. Our account of the sampling process depends on assuming that perception is noisy, which makes it necessary to take multiple samples from a stimulus until the information content of the stimulus has been learned sufficiently.

The similarity between model fits among models with different linking hypotheses highlights the significance of learning contexts. Our results should not be interpreted as evi-

dence showing that the three linking hypotheses are indistinguishable across all learning contexts. Previous work has shown that adopting surprisal as learning policy can lead to undesirable behaviors in artificial agents (e.g. “the white noise problem,” Oudeyer et al., 2007). Moreover, the two alternative linking hypotheses are backward-looking metrics that utilize heuristics about the past to make decisions. This characteristic could constrain their application to situations in which the environment is stable and the cost of sampling is low. Since adult exploration is sensitive to environmental complexity, a forward-looking metric like EIG might be particularly suitable to predict behaviors in a more dynamic learning context (Dubey & Griffiths, 2020; Vogelstein et al., 2022).

There are several limitations to our work. For our behavioral data, one concern is that adult looking time might not be driven by intrinsic interest to the same degree as infant looking time. Rather, they might be driven by task-preparation. However, across the three conditions with different cover tasks, we found no differences in looking time patterns. In regards to the model, a few concerns can be raised. First, the current stimulus representation is oversimplified, using an unweighted collection of binary features. Future research could apply RANCH to stimulus representations generated from a perceptual model. Second, RANCH assumes that the EIG from the environment is a constant throughout the experiment, but one can argue that environmental EIG might increase as the experiment progresses (e.g. the longer you have not attended to the things in your surroundings, the more they may have changed in the meantime). While implementing more sophisticated assumptions could potentially explain additional variance in the data, our current work suggests that even a simple rational learner that takes noisy samples from a set of independent binary features is capable of explaining key phenomena.

Our ultimate goal is to provide a rational learner model that can account for information-seeking behaviors reflected in infants’ looking time. Here we have shown that a simple model of learning from sampling can reproduce habituation, dishabituation, and complexity effects. Moving forward, we aim to capture and explain more contentious phenomena documented in the infant looking time literature such as familiarity preferences and age effects (Hunter & Ames, 1988). Our ongoing work with infants will eventually enable us to evaluate our model with developmental data. When combined with adult results, the data and model will provide insights into the general mechanisms through which learners decide what to look at, and when to stop looking.

References

- Aslin, R. N. (1991). Development of processing speed in childhood and adolescence. *Advances in Child Development and Behavior*, 23, 151–185.
- Baillargeon, R., Spelke, E. S., & Wasserman, S. (1985). Object permanence in five-month-old infants. *Cognition*,

- 20(3), 191–208.
- Baldi, P., & Itti, L. (2010). Of bits and wows: A bayesian theory of surprise with applications to attention. *Neural Networks*, 23(5), 649–666.
- Bergmann, C., & Cristia, A. (2016). Development of infants' segmentation of words from native speech: A meta-analytic approach. *Developmental Science*, 19(6), 901–917.
- Bitzer, S., Park, H., Blankenburg, F., & Kiebel, S. J. (2014). Perceptual decision making: Drift-diffusion model is equivalent to a bayesian model. *Frontiers in Human Neuroscience*, 8, 102.
- Callaway, F., Rangel, A., & Griffiths, T. L. (2021). Fixation patterns in simple choice reflect optimal information sampling. *PLoS Computational Biology*, 17(3), e1008863.
- Carey, S. (2009). *The origin of concepts*. Oxford University Press.
- Colombo, J., & Mitchell, D. W. (2009). Infant visual habituation. *Neurobiology of Learning and Memory*, 92(2), 225–234.
- Dubey, R., & Griffiths, T. L. (2020). Reconciling novelty and complexity through a rational analysis of curiosity. *Psychological Review*, 127(3), 455.
- Fantz, R. L. (1963). Pattern vision in newborn infants. *Science*, 140(3564), 296–297.
- Frank, M. C., Bergelson, E., Bergmann, C., Cristia, A., Floccia, C., Gervain, J., . . . others. (2017). A collaborative approach to infant research: Promoting reproducibility, best practices, and theory-building. *Infancy*, 22(4), 421–435.
- Goodman, N. D., Tenenbaum, J. B., Feldman, J., & Griffiths, T. L. (2008). A rational analysis of rule-based concept learning. *Cognitive Science*, 32(1), 108–154.
- Haber, N., Mrowca, D., Fei-Fei, L., & Yamins, D. L. (2018). Learning to play with intrinsically-motivated self-aware agents. *arXiv Preprint arXiv:1802.07442*.
- Haith, M. M. (1980). *Rules that babies look by: The organization of newborn visual activity*. Lawrence Erlbaum Associates.
- Haith, M. M. (1998). Who put the cog in infant cognition? Is rich interpretation too costly? *Infant Behavior and Development*, 21(2), 167–179.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., Group, C. S. R. others. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19(1), 46–54.
- Hunter, M. A., & Ames, E. W. (1988). A multifactor model of infant preferences for novel and familiar stimuli. *Advances in Infancy Research*.
- Hunter, M. A., Ames, E. W., & Koopman, R. (1983). Effects of stimulus complexity and familiarization time on infant preferences for novel and familiar stimuli. *Developmental Psychology*, 19(3), 338.
- Kail, R. (1991). Development of processing speed in childhood and adolescence. *Advances in Child Development and Behavior*, 23, 151–185.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as bayesian inference. *Annu. Rev. Psychol.*, 55, 271–304.
- 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.
- Kim, K., Sano, M., De Freitas, J., Haber, N., & Yamins, D. (2020). Active world model learning with progress curiosity. In *International conference on machine learning* (pp. 5306–5315). PMLR.
- Liquin, E. G., Callaway, F., & Lombrozo, T. (2021). Developmental change in what elicits curiosity. In *Proceedings of the annual meeting of the cognitive science society* (Vol. 43).
- Markant, D., & Gureckis, T. (2012). Does the utility of information influence sampling behavior? In *Proceedings of the annual meeting of the cognitive science society* (Vol. 34).
- Oakes, L. M. (2010). Using habituation of looking time to assess mental processes in infancy. *Journal of Cognition and Development*, 11(3), 255–268.
- Oaksford, M., & Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4), 608.
- Oudeyer, P.-Y., Kaplan, F., & Hafner, V. V. (2007). Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(2), 265–286.
- Pirolli, P., & Card, S. (1999). Information foraging. *Psychological Review*, 106(4), 643.
- Poli, F., Serino, G., Mars, R., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science Advances*, 6(39), eabb5053.
- Sim, Z. L., & Xu, F. (2019). Another look at looking time: Surprise as rational statistical inference. *Topics in Cognitive Science*, 11(1), 154–163.
- Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. *Advances in Neural Information Processing Systems*, 59–68.
- Vogelstein, J. T., Verstynen, T., Kording, K. P., Isik, L., Krakauer, J. W., Etienne-Cummings, R., . . . others. (2022). Prospective learning: Back to the future. *arXiv Preprint arXiv:2201.07372*.