Emergence of certainty representations for guiding concept learning

Huiwen Alex Yang[1] (hw.yang@berkeley.edu)

Louis Martí[1] (lmarti13@gmail.com)

Carolyn Baer[2, 3] (cebaer@mail.ubc.ca)

Aaraam Granera[1] (ajj.granera@berkeley.edu)

Holly Palmeri[1] (holly.palmeri@berkeley.edu)

Celeste Kidd[1] (celestekidd@berkeley.edu)
1 Department of Psychology, University of California
2 Department of Psychology, University of British Columbia
3 Department of Psychology, Kwantlen Polytechnic University

Abstract

Previous research has shown that our subjective sense of certainty doesn't always accurately reflect the strength of the evidence that has been presented to us. We investigate several key factors that drive children's certainty using a Boolean concept learning task. We created an idealized learning model to predict children's accuracy and certainty during the experiment, given past evidence that they have seen in the task, and we compared its predictions with our behavioral results. Our results suggest that while predictors from the idealized learning model capture children's accuracy, behavioral predictors generated by the behavioral data can better predict children's certainty. We also show that younger children's certainty can be explained by the idealized learning model, while older children's certainty is primarily predicted by how well they observed themselves doing in the experiment.

Keywords: certainty; confidence; concept learning

Introduction

Certainty directly determines what agents learn and believe. Past research has shown that certainty guides attention and also moderates the way information is encoded (Wade & Kidd, 2019): if you think you know a lot, not only do you stop seeking out new information, but even when new information is presented to you, you do not encode the information the same way as if you were uncertain. In the context of learning, certainty serves as a general-purpose metacognitive signal that guides learning (Kidd & Hayden, 2015; Baer & Kidd, 2022). Human infants (Kidd, Piantadosi, & Aslin, 2012; Piantadosi, Kidd, & Aslin, 2014), children (Cubit, Canale, Handsman, Kidd, & Bennetto, 2021; Wang, Yang, Macias, & Bonawitz, 2021), and adults (Kang et al., 2009) prefer to direct their attention and engage in learning when they are moderately uncertain about the stimuli. Certainty also guides learning in non-human species. For instance, studies have shown that macaques exhibit a similar pattern: they show preferential attention to moderately complex events (Wu et al., 2022).

Research shows that subjective sense of certainty (how confident participants feel) and objective measures of certainty (the probability that participants are correct, given the strength of the evidence) can be two distinctive signals (Pouget, Drugowitsch, & Kepecs, 2016; Martí, Mollica, Piantadosi, & Kidd, 2018). Furthermore, a subjective sense of certainty can depend on many factors other than the objective measures of probability. For instance, research has found that subjective certainty, or confidence, can be influenced by participants' competence (Kruger & Dunning, 1999), the difficulties of tasks (Larrick, Burson, & Soll, 2005), and the amount of information available (Tsai, Klayman, & Hastie, 2008).

In the present study, we present an experiment to look at factors driving children's subjective certainty when they perform a Boolean concept-learning task. Concept learning is the process of learning categorization from examples. The Boolean concept-learning tasks we used involve learning concepts whose membership is determined by a combination of binary features (Shepard, Hovland, & Jenkins, 1961). They were widely used to study concept learning because the complexity of these Boolean concepts can be measured easily (Feldman, 2000). They also allow researchers to construct a learning model with a simplified hypothesis space (Goodman, Tenenbaum, Feldman, & Griffiths, 2008; Martí et al., 2018).

Our Boolean concept-learning task requires children to conduct rational inductive learning within that simplified hypothesis space. Inductive reasoning ability is important because it is required for many cognitive tasks, such as making abstractions and deriving rules (Perret, 2015). Previous research has shown that children, as young as 4 years old, can use their inductive reasoning abilities to infer causal structures from evidence (Gopnik & Schulz, 2004; Bright & Feeney, 2014; Sobel, Tenenbaum, & Gopnik, 2004). However, no experiment has measured and explained children's change of certainty in this inductive learning process. In this paper, we present experimental and modeling evidence that as children age, the factors that drive their certainty shift from

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.* ©2024 The Author(s). This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY).

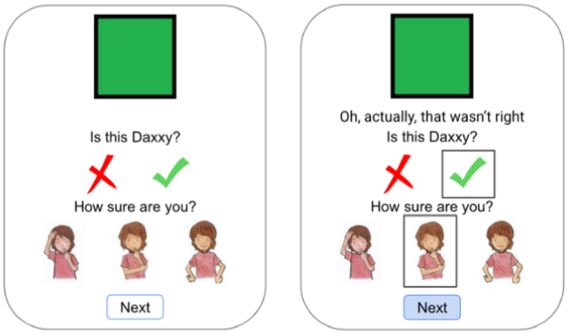

Figure 1: Experimental design. Left: participants see 12 trials that measure their accuracy and certainty. Right: feedback was given after each trial.

Table 1: Concepts and their corresponding rules.

Concept	Rule
1	red ∧ square
2	red ∨ square
3	red
4	((red ∧ triangle) ∨ (green ∧ square))
5	((red ∧ square) ∨ (green ∧ triangle))

the strength of the observed evidence to how well they observe themselves doing in the course of learning.

Methods

Participants

76 children between the ages of 4 to 8 years old participated in the experiment either in person or on Zoom. Data from 5 participants were removed due to missing information on their precise age at the testing date. The analysis is done on the remaining 71 participants (Mean age = 6.36, SD = 1.35).

Design

We asked children to perform a standard Boolean concept-learning task, in which we measured their knowledge of a hidden concept and their certainty throughout the learning process. In this experiment, participants were shown positive and negative examples of a target concept "daxxy", whose meaning was determined by a latent rule on two feature dimensions (the category and the color of a shape), following experimental work by Shepard et al. (1961) and Feldman (2000). The latent rules participants were required to learn varied across various logical forms. We tested each child on just one of 5 possible latent rules, chosen at random. Thus, each child had to infer only one rule for a set of objects. The conceptual rules are shown in Table 1. Concept difficulties are calculated by the complexities of the Boolean logical operations that define the concept (Feldman, 2000). Some concepts are simpler because they contain fewer logical operations and fewer shape parameters. For instance, concept 3 is the simplest because its rule contains no logical operations and only one shape parameter. Some concepts are harder because they contain more logical operations and more shape parameters, like concepts 4 and 5.

The experiment contains 12 trials. In each trial, participants see one of the four possible shapes (red square, red triangle, green square, green triangle) and provide their guess on whether the shape belongs to the concept "daxxy" or not. After that, they also provided their certainty-how confident they are that a shape is or isn't "daxxy". In our experiment, the confidence level was measured using a 3-point scale accompanied by pictures of a child expressing high, medium, and low confidence in their answer (Hembacher & Ghetti, 2014). Participants were trained to point to the low-confidence option if they were "not so sure", the medium-confidence option if they were "a little sure", and the high-confidence option if they were "really sure".

Before they did these 12 test trials, the researchers did pre-task training with participants using a different training concept that was very easy to learn so that they became familiar with Boolean concept-learning tasks. Participants also received training and practice using the confidence scale to make sure that they felt comfortable expressing their certainty with this scale and that they understood what each option meant.

Analysis

For the analyses, we proposed several factors that could explain participants' uncertainty and compared them to determine the best ones. These predictors can be classified into two broad categories: model-based predictors, and behavioral predictors. Model-based predictors were calculated using our idealized learning model, and behavioral predictors were calculated using the behavioral data. Since logarithmic transformations are common in psychophysics, many of our predictors were considered in their standard form, as well as under a logarithmic transformation. Examples of top-performing model-based and behavioral predictors and their definitions are shown in Table 2.

We created a model to determine how confident a learner should be, given past evidence they have seen in the experiment. We call the model an idealized learning model because it acts as a rational learner that considers the strength of all the past information and weighs each new piece of evidence when determining which of all possible rules in a hypothesis space is more likely to be the concept. Our implementation was developed using the Language Of Thought library, LOTlib (Piantadosi, 2014).

To use this modeling technique, we first define a probabilistic context-free grammar (PCFG) with a set of primitives that contains logical primitives (and, or, not), category primitives (triangle, square), and color primitives (red, green). We can combine these primitives to form latent rules such as those

Table 2: Key behavioral predictors and their definitions (shaded) as well as key model-based predictors and their definitions (unshaded)

Predictor	Description
Local Accuracy	Performance on previous N trials (N = 2, 3, 4, 5)
Local Accuracy Current	Performance on previous N trials (N = 2, 3, 4, 5) and a guess on the current trial
Total Accuracy	Total number of correct trials so far
Trial	Number of trials participants have seen so far
Entropy	Model uncertainty over hypotheses regarding what the concept is
Domain Entropy	Model uncertainty over which objects belong to the concept
MAP	The probability of the best hypothesis
MAPNoPrior	The probability of the best hypothesis ignoring the prior probability

used in concepts shown in Table 1. Each rule that this grammar can generate serves as a hypothesis of what the concept could be. Together, these rules span an infinite hypothesis space. We also defined the probability distribution such that there is a uniform prior over each basic rule in the grammar. As a result of the multiplication of compositional rules, simpler rules are favored over more complex ones.

To establish a tractable hypothesis space, the model drew 10000 samples from the posterior distribution of hypotheses using tree-regeneration Metropolis-Hastings and stored the best 100 hypotheses based on the simplicity and the fit of hypotheses. The model also includes a parameter for the noise in the data (alpha) and a parameter for the power law memory decay on the likelihood of previous data (beta). We set the value of alpha to 0.40 and beta to 0 as a result of doing a grid search and selecting the values that yielded the best fit for our data.

Results

Accuracy and certainty improve throughout the experiment

Figure 2 shows participants' performance over trials of the experiment, separated by concepts. The red circles represent the mean accuracy of all participants. Participants become more accurate throughout the experiment, which indicates that they gradually learn the concept with more trials. We can also see that they didn't learn every concept equally well. In the first three concepts, participants reached high accuracy, demonstrating that they learned the concepts. However, the accuracy of concepts 4 and 5 is significantly lower than that of other concepts. Therefore, it's not clear whether children learned the concepts in these two concepts. This result is consistent with what Shepard et al. (1961) and Feldman (2000, 2003) found, that concepts referring to more dimensions with more logical operators were harder to learn.

In Figure 2, the blue triangles represent the mean certainty of all participants. In our analysis, certainty is coded as a binary variable that can have the value 0 or 1. Certainty was coded as 0 if the participant pointed to either the low-confidence option or the medium-confidence option. It was coded as 1 if the participant pointed to the high-confidence option. We coded both the low-confidence option and the medium-confidence option as certainty being 0 because we found that children tend to either only use a combination of the low-confidence option and the high-confidence option, or use a combination of the medium-confidence option and the high-confidence option to signal their change in confidence level.

Table 3: Top three behavioral predictors and their AIC (shaded) and top three model-based predictors and their AIC (unshaded).

Predictor	AIC
LocalAccuracy3BackCurrent	965.3644
LocalAccuracy1BackCurrent	966.4860
LocalAccuracy2BackCurrent	966.9417
Entropy	987.9356
MAP	988.0496
Domain Entropy	991.0423

Figure 2 shows that, in general, participants also become more confident with their answers through the experiment, regardless of the difficulty of the concept. However, they only achieve high certainty when their accuracy is also high, in the first three concepts.

Idealized learning model predicts accuracy

We used a generalized logistic mixed-effect model with random subject and concept effects to determine how well our model can predict behavioral accuracy. Figure 3 shows our result. It shows that behavioral accuracy can be well-predicted by model accuracy ($\beta = .748$, $z = 30.423$, $p < .001$). This is an indication that the model can capture children's learning behaviors and outcomes reasonably well.

Local accuracy predicts certainty

Next, we aimed to see if our idealized learning model can also predict participants' certainty during the experiment. We ran a generalized logistic mixed-effect model with random

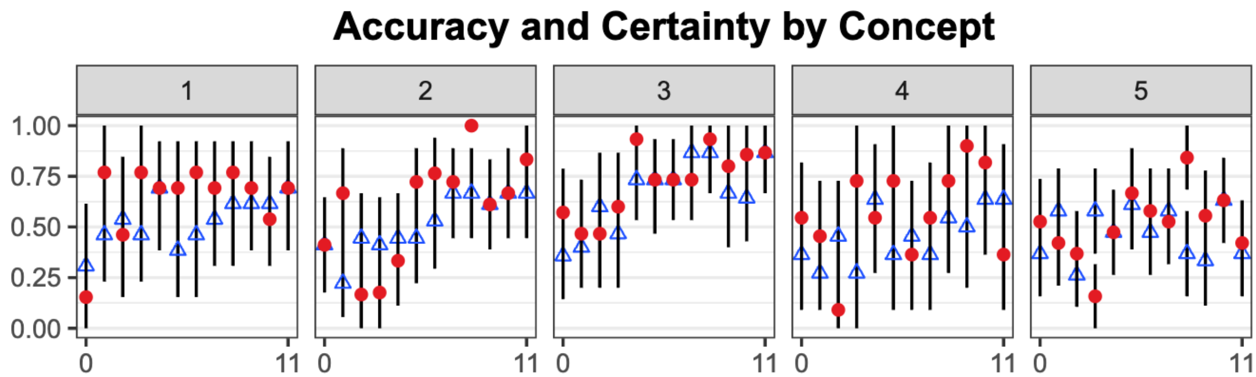

Figure 2: Mean accuracy (red circles) and mean certainty (blue triangles) throughout the experiment, separated by concepts. The upward trend of both red circles and blue triangles indicates that participants become more accurate and more certain over the experiment. However, the relatively low accuracy and certainty in concepts 4 and 5 indicate that concepts with more complex rules are harder to learn and participants only achieve high certainty when the accuracy is also high.

Table 4: Regression for best predictors with behavioral predictors (shaded) and model-based predictors (unshaded).

Predictor	Beta	Standard Error	z Value	Pr(>|z|)
Intercept	0.14	0.07	1.93	0.054
LocalAccuracy3BackCurrent	0.60	0.14	4.14	< 0.001
log(totalCorrect)	-0.02	0.22	-0.07	0.943
log(Trial)	-0.11	0.15	-0.75	0.455
Entropy	-0.73	0.47	-1.55	0.120
MAP	0.05	0.23	0.22	0.822
Domain Entropy	0.52	0.38	1.46	0.174

subject and concept effects for every model-based predictor and behavioral predictor. We ranked these models by their Akaike information criterion (AIC), which quantifies the fit of each model while penalizing its number of free parameters (a model with a lower AIC is considered to be better). Table 3 shows the top 3 behavioral and model-based predictors and their AIC. We found that the best group of predictors is LocalAccuracyCurrent. This means that the performance of the previous few trials and a guess of the current trial can best capture the changes in participants' certainty. Among this group of predictors, LocalAccuracy3BackCurrent, the accuracy averaged over the past 3 trials and a guess on the current trial has the best performance. This indicates that participants report higher certainty when they performed more accurately over the past 3 trials.

We also found that, in general, behavioral predictors outperform model-based predictors. As Table 3 shows, the AIC of our top behavioral predictors is significantly lower than that of our top model-based predictors.

We created a single regression model using the top three behavioral predictors and the top three model predictors to determine the unique contributions of each predictor when multiple predictors are allowed in the same model. We used LocalAccuracy3BackCurrent as the representative predictor of all predictors from LocalAccuracy and LocalAccuracyCurrent group, and we added predictors totalCorrect and Trial to the regression model since they yielded the next best individual model in our previous analysis. Table 4 shows our regression results. As it shows, LocalAccuracy3BackCurrent significantly outperforms all other predictors in its contribution to certainty. All other predictors were not significant when controlling for LocalAccuracy. This result aligns with the result from individual generalized logistic mixed-effect models that LocalAccuracy predicts certainty.

Our results mirror Martí et al. (2018)'s study, which shows that adults' certainty is also primarily driven by local feedback. Martí et al. (2018) employed a similar but more complex Boolean concept-learning task with adults. The results showed that adults' certainty can also be predicted by LocalAccuracy as opposed to model-based predictors from the idealized learning model. These results, together, suggest that while an idealized learning model that only considers the strength of past evidence can capture participants' overall performance, it does not predict participants' certainty. Instead, their certainty is predicted by how well they perceive themselves doing during the experiment based on their feedback.

Table 5: Top three behavioral predictors and their AIC (shaded) and top three model-based predictors and their AIC (unshaded) for younger, older, and all participants.

All participants		Younger participants		Older participants	
Predictors	AIC	Predictors	AIC	Predictors	AIC
LocalAccu3BackCurrent	965.36	log(MAP)	527.77	LocalAccu1BackCurrent	450.68
LocalAccu1BackCurrent	966.49	log(MAPNoPrior)	528.26	LocalAccu2BackCurrent	452.82
LocalAccu2BackCurrent	966.94	Entropy	528.36	LocalAccu3BackCurrent	453.33
Entropy	987.94	LocalAccu4BackCurrent	528.50	Entropy	480.18
MAP	988.05	LocalAccu3BackCurrent	529.31	Domain Entropy	483.35
Domain Entropy	991.04	log(LocalAccu3BackCurrent)	530.28	MAP	485.80

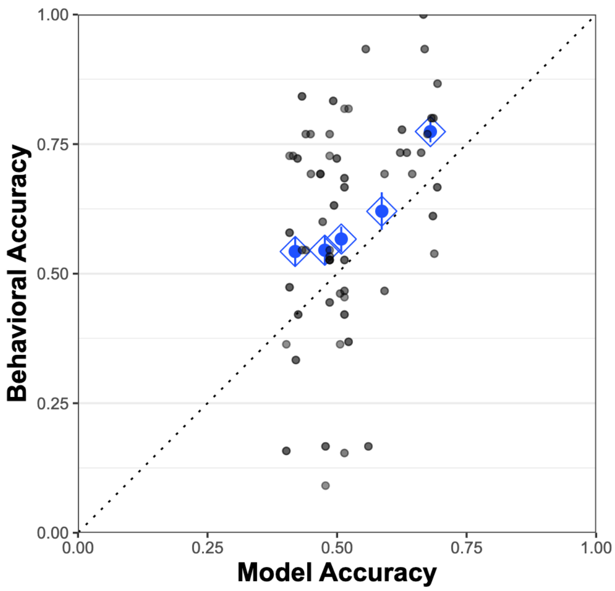

Figure 3: Behavioral accuracy as a function of accuracy predicted by the idealized learning model.

Younger children's certainty is predicted by the idealized learning model and older children's certainty is predicted by local feedback

We were interested in whether the primary factors that drive children's uncertainty change throughout development. Therefore, we split the dataset by age in order to compare trials from younger participants (N = 35, Age = 4.05 - 6.29, mean = 5.18, SD = 0.66) to those from older participants (N = 36, Age = 6.4 - 8.98, mean = 7.51, SD = 0.69). We also ran a generalized logistic mixed-effect model with random subject and concept effects for all model-based and behavioral predictors separately, as we have run before, on these two datasets. Table 5 shows our modeling results for all participants, younger participants, and older participants. As it shows, while older participants' certainty is primarily driven by predictors that are related to local feedback, younger participants' certainty can be well-predicted by our model-based predictors.

For both age groups, we created a regression model using the top three behavioral predictors and the top three model-based predictors, and we compared the results with those of all participants. Figure 4 shows certainty as a function of the best behavioral and model-based predictors for younger, older, and all participants. Our results are consistent with the previous results from the generalized logistic mixed-effect model. We confirmed that for younger participants, when controlling for the effect of top model-based predictors, behavioral predictors do not make a significant contribution to predicting certainty. For older participants, when controlling for the effect of local feedback, model-based predictors don't make a significant contribution to predicting certainty.

Conclusions and Discussion

Our experiment demonstrated that an idealized learning model can capture participants' accuracy in a Boolean concept-learning task. It, however, cannot predict participants' subjective sense of certainty during the course of learning. Instead, we found that predictors derived from local feedback have the most predictive power of their certainty. Our age-analysis results suggest that as opposed to older children and adults (Martí et al., 2018), younger children's certainty is best predicted by the model-based predictors that evaluate the strength of the accumulated evidence. This is interesting because it suggests that either maturation or experience triggers a shift in what determines human certainty—and that older humans are more susceptible to distortion in how certain they should be derived from feedback as a heuristic.

Past research findings suggest that young children are capable of understanding and learning from the statistical structure of information in a way that is consistent with using a Bayesian framework (Xu & Tenenbaum, 2007; Gopnik & Schulz, 2004; Griffiths, Sobel, Tenenbaum, & Gopnik, 2011), much like the idealized learning model we created to capture the performance of children's learning. It could be that children use those inferred statistical structures as an indication of how confident they should feel, given past evidence. However, as children grow—and perhaps also as they develop metacognitive and social awareness—they increasingly rely on feedback to inform their certainty. The implications of this age-related difference are many, and include that younger

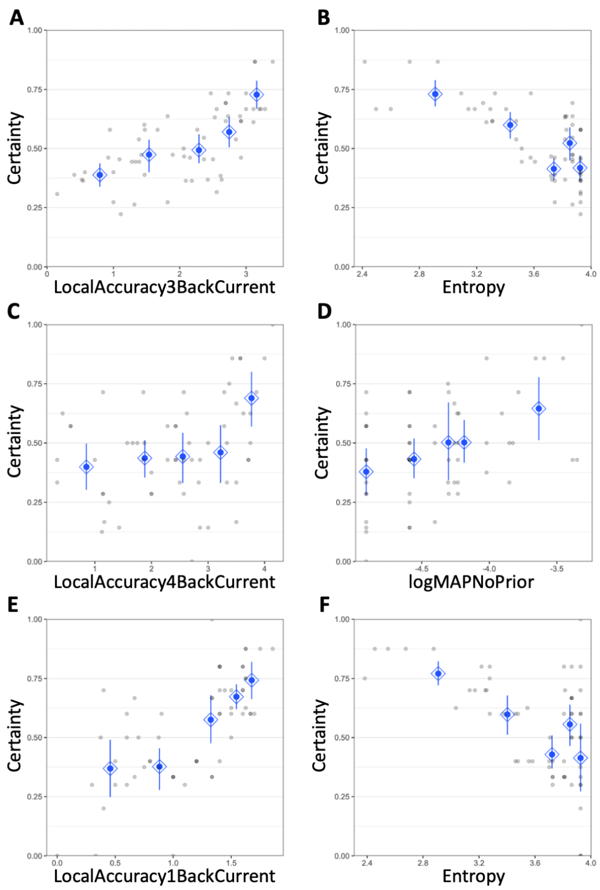

Figure 4: Best behavioral and model-based fits for all, younger and older children. (**A**) For all participants, localAccuracy3BackCurrent is the best behavioral predictor. The more positive the local feedback, the more certainty participants report on the current trial. (**B**) For all participants, entropy is the best model-based predictor. Participants report higher certainty when the data has low entropy. (**C**) For younger participants, localAccuracy4BackCurrent is the best behavioral predictor. It is, however, not a strong predictor. (**D**) For younger participants, the log of likelihood is the best model-based predictor. (**E**) For older participants, localAccuracy1BackCurrent is the best behavioral predictor. The more positive the local feedback, the more certainty participants report on the current trial.(**F**) For older participants, entropy is the best model-based predictor. Participants report higher certainty when the data has low entropy.

children may be less susceptible to social feedback-related influences on their beliefs (Orticio, Martí, & Kidd, 2022).

We also found age-group differences in children's sensitivity to certainty and the criteria they use to classify different levels of certainty (Baer & Kidd, 2022). For instance, children are usually overconfident in their performances (Van Loon, De Bruin, Leppink, & Roebers, 2017), however, overconfidence tends to decrease as children age (Schneider, Visé, Lockl, & Nelson, 2000). In our analysis, we introduced the random subject and concept effects into our idealized learning model to address these differences, while maintaining our focus on the contributions of the key factors that drive children's certainty during concept learning. Further work is needed to disentangle the mechanisms responsible for these age-related changes, as they could depend upon maturation, experience, or both.

The observed shift—from younger children's reliance on the strength of evidence towards older children's reliance on feedback for determining their certainty—makes specific predictions with profound implications on how and why younger children's curiosity and learning may operate differently. Key implications include that their information-seeking and informational encoding—both of which are tied directly to their curiosity (Wade & Kidd, 2019)—are more closely aligned with what they truly do not know than only what they believe they do not know. We will further investigate the mechanisms underlying this shift and its implications in subsequent work.

References

Baer, C., & Kidd, C. (2022). Learning with certainty in childhood. *Trends in Cognitive Sciences*, *26*(10), 887–896.

Bright, A. K., & Feeney, A. (2014). Causal knowledge and the development of inductive reasoning. *Journal of Experimental Child Psychology*, *122*, 48–61.

Cubit, L. S., Canale, R., Handsman, R., Kidd, C., & Bennetto, L. (2021). Visual Attention Preference for Intermediate Predictability in Young Children. *Child Development*, *92*(2), 691–703.

Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature*, *407*(6804), 630–633.

Feldman, J. (2003). The Simplicity Principle in Human Concept Learning. *Current Directions in Psychological Science*, *12*(6), 227–232.

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.

Gopnik, A., & Schulz, L. (2004). Mechanisms of theory formation in young children. *Trends in Cognitive Sciences*, *8*(8), 371–377.

Griffiths, T. L., Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2011). Bayes and Blickets: Effects of Knowledge on Causal Induction in Children and Adults: Cognitive Science. *Cognitive Science*, *35*(8), 1407–1455.

Hembacher, E., & Ghetti, S. (2014). Don't Look at My Answer: Subjective Uncertainty Underlies Preschoolers' Exclusion of Their Least Accurate Memories. *Psychological Science*, *25*(9), 1768–1776.

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. (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.

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it: How difficulties in recognizing one's own incompetence lead to inflated self-assessments. *Journal of Personality and Social Psychology*, *77*(6), 1121–1134.

Larrick, R. P., Burson, K. A., & Soll, J. B. (2005). Social Comparison and Confidence: When Thinking You're Better than Average Predicts Overconfidence. *SSRN Electronic Journal*.

Martí, L., Mollica, F., Piantadosi, S., & Kidd, C. (2018). Certainty Is Primarily Determined by Past Performance During Concept Learning. *Open Mind*, *2*(2), 47–60.

Orticio, E., Martí, L., & Kidd, C. (2022). Social Prevalence Is Rationally Integrated in Belief Updating. *Open Mind*, *6*, 77–87.

Perret, P. (2015). Children's Inductive Reasoning: Developmental and Educational Perspectives. *Journal of Cognitive Education and Psychology*, *14*(3), 389–408.

Piantadosi, S. T. (2014). *LOTlib: Learning and Inference in the Language of Thought.*

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.

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, *19*(3), 366–374.

Schneider, W., Visé, M., Lockl, K., & Nelson, T. O. (2000). Developmental trends in children's memory monitoring. *Cognitive Development*, *15*(2), 115–134.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs: General and Applied*, *75*(13), 1–42.

Sobel, D. M., Tenenbaum, J. B., & Gopnik, A. (2004). Children's causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science*, *28*(3), 303–333.

Tsai, C. I., Klayman, J., & Hastie, R. (2008). Effects of amount of information on judgment accuracy and confidence. *Organizational Behavior and Human Decision Processes*, *107*(2), 97–105.

Van Loon, M., De Bruin, A., Leppink, J., & Roebers, C. (2017). Why are children overconfident? Developmental differences in the implementation of accessibility cues when judging concept learning. *Journal of Experimental Child Psychology*, *158*, 77–94.

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

Wang, J. J., Yang, Y., Macias, C., & Bonawitz, E. (2021). Children With More Uncertainty in Their Intuitive Theories Seek Domain-Relevant Information. *Psychological Science*, *32*(7), 1147–1156.

Wu, S., Blanchard, T., Meschke, E., Aslin, R. N., Hayden, B. Y., & Kidd, C. (2022). Macaques preferentially attend to intermediately surprising information. *Biology Letters*, *18*(7), 20220144.

Xu, F., & Tenenbaum, J. B. (2007). Word learning as Bayesian inference. *Psychological Review*, *114*(2), 245–272.