8-month-old infants meta-learn from sparse evidence

Francesco Poli^{1*} (ORCID : 0000-0001-9841-0056)

Tommaso Ghilardi¹ (ORCID : 0000-0003-0855-2814)

Rogier B. Mars^{1,2} (ORCID : 0000-0001-6302-8631)

Max Hinne¹ (ORCID : 0000-0002-9279-6725)

Sabine Hunnius¹ (ORCID : 0000-0001-5510-4032)

¹ Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, Netherlands.

² Wellcome Centre for Integrative Neuroimaging, Centre for Functional MRI of the Brain (FMRIB), Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, University of Oxford, Oxford, UK.

* Corresponding author (email: francesco.poli@donders.ru.nl)

Abstract

Infants learn to navigate the complexity of the physical and social world at an outstanding pace, but how they accomplish this learning is still largely unknown. Recent advances in human and artificial intelligence research propose that a key feature to achieving quick and efficient learning is meta-learning, the ability to make use of prior experiences to optimize how information is acquired in the future. Here we show that 8-month-old infants successfully engage in meta-learning within very short timespans after being exposed to a new learning environment. We developed a Bayesian model that captures how infants attribute informativity to incoming events, and how this process is optimized by the meta-parameters of their hierarchical models over the task structure. We fitted the model with infants' gaze behaviour during a learning task. Our results reveal how infants actively use past experiences to generate new inductive biases that allow future learning to proceed faster.

Keywords: Meta-learning, Infants, Computational modelling, Eye-tracking.

Introduction

Infants use their learning abilities to build sophisticated models of the environment and constantly update and refine them to better navigate the complexity of the physical and social world (Hunnius, 2022; Köster et al., 2020). These abilities are advanced from early in life (Emberson et al., 2015; Kidd et al., 2012; Poli et al., 2020) and possibly from birth (Bulf et al., 2011; Craighero et al., 2020). A key aspect of infant learning is the ability to exploit newly learned content to improve further learning (Dewar & Xu, 2010; Thiessen & Saffran, 2007). However, the cognitive mechanisms that support this ability are still unknown.

This ability to learn how to learn is often referred to as meta-learning, and has been placed at the core of recent theories of human (Baram et al., 2021; Wang et al., 2018) and artificial intelligence (Lake et al., 2017). When agents meta-learn, they do not simply accumulate experiences, but actively use them to generate new inductive biases and knowledge. This in turn makes learning proceed faster and more efficiently in the future (Kemp et al., 2007; Wang et al., 2018).

A classical study by Harlow (1949) illustrates the key aspects of meta-learning. Macaque monkeys were given a choice between two objects. One of the objects predicted the location of some food with perfect accuracy, while the other object was always paired with an empty well. Every six trials the two objects (i.e., the learning set) changed, but the underlying rule remained the same: one object led to a reward, the other did not. This setup meant that in theory it was always possible to know which object was rewarded after a single trial: the chosen object if a reward was received, or the non-chosen one if no reward was received. Across hundreds of trials, the macaques learned the structure of the task and became able to predict after one single presentation of a new object set where the food was. Harlow also tested preschool children with a similar task, obtaining a comparable performance in even shorter timespans (Harlow, 1949). In contrast, squirrel monkeys and marmosets take almost a thousand sets and never reach the same level of performance; rodents never score above chance (Passingham & Wise, 2012).

The study by Harlow (1949) on learning sets offers a simple but elegant example of learning to learn, a sophisticated ability that requires the agent to form meta-representations over the task space, allowing past experience with similar situations to support the learning in new situations (Wang et al., 2018). Recent advances in cognitive computational modelling capture how adults acquire structured priors that support their learning over multiple scales (Gershman & Niv, 2010; Lake et al., 2015; Vossel et al., 2014). Similarly, research in artificial intelligence focuses on how complex structures can emerge from statistical regularities, and what prior knowledge is required to support such processes (Alet et al., 2020; Grant et al., 2018), if any (Piantadosi, 2021). Such questions are extremely relevant within developmental science (Rule et al., 2020; Xu, 2019; Yuan et al., 2020), and Bayesian theories of brain functioning (Tenenbaum et al., 2011) hold promise to answer them. This would help overcome rigid

dichotomies such as “empiricism versus nativism” or “domain-specific versus domain-general”, moving towards a better understanding of the nature of the cognitive mechanisms that underlie learning in infancy and of the priors that might support acquisition of knowledge from birth.

From the early studies on learning sets (Harlow, 1949; Koch & Meyer, 1959), developmental science has come a long way. Infants’ learning skills (Emberson et al., 2015; Trainor, 2012), their flexibility (Kayhan et al., 2019; Tummeltshammer & Kirkham, 2013) and complexity (Werchan et al., 2015) have been widely demonstrated. However, current research lacks a mechanistic explanation of how meta-learning occurs in the infant mind. In this paper, we show how we can leverage onto hierarchical Bayesian models to bring novel insights into the cognitive mechanism that allow infants to exploit prior knowledge to optimize how new stimuli are learned.

We presented 8-month-old infants with multiple probabilistic sequences of stimuli while monitoring their eye movements via eye-tracking, and we quantified the information content of each stimulus using Bayesian updating and information theory. We propose that when infants meta-learn, they gain the ability to strategically weight the value of incoming information. Namely, they might up-weight the information content of stimuli that are expected to carry more information, or down-weight the information content of stimuli that are expected to be irrelevant. Similar to Harlow’s task (1949), our task provided a general structure that infants could extract over multiple sequences. In each sequence, the most relevant information was in the first trials, and later trials were less relevant to learn to predict the following stimuli. If infants understood this structure, they could make use of it to change how they learned: By up-weighting evidence acquired early in the sequence and down-weighting evidence acquired later on, they could optimize their learning.

Methods

Participants

Ninety 8-month-old infants ($M = 8.05$ months, $SD = 11.56$ days, 42 females) were recruited for this study from a database of volunteer families. Infants that carried out less than 20 trials were excluded from the analysis ($N = 17$). The final sample consisted of 73 infants ($M = 8.02$ months, $SD = 11.37$ days, 34 females). For this study, a previously published dataset (Poli et al., 2020) was reanalysed ($N = 50$), and new data were collected using the same task ($N = 40$).

Procedure

Infants were tested in a silent room with fixed levels of light. They were placed in a baby seat, which was held by the parents on their lap. The eyes of the infants were approximately 65 cm far from the screen. Parents were instructed not to interact with their child, unless infants sought their attention and,

even in that case, not to try to bring infants' attention back to the screen. During the experiment, the infants' looking behaviour was monitored using both a Tobii X300 eye-tracker and a video camera. We extracted visual fixations using I2MC (Hessels et al., 2017) with the default parameters of the algorithm. We then extracted saccadic latency, looking time and look-away from the pre-processed data.

Experimental Paradigm

Infants were presented with a visual learning eye-tracking task in which 16 sequences of cue-target pairs were shown. In each sequence, the cue consisted of a shape appearing in the middle of the screen. The target was the same shape appearing in one of four screen quadrants around the cue location (Figure 1). The shape was the same across all trials of the sequence but changed across sequences. Four of the sixteen sequences were deterministic, as the target always appeared in the same location. The other twelve sequences were probabilistic: the target could appear in any location, but it appeared in one specific location more often than the others (either 60% or 80% of the times, see Figure 2). The sequences were shown one after the other. When the infant looked away from the screen for one second or more, the sequence was stopped. When the infant looked back to the screen, the following sequence was played. The experiment lasted until the infant had watched all sixteen sequences or became fussy.

Figure 1. An example trial of the visual learning task. Each sequence was composed of a maximum of 15 trials and terminated when infants looked away for 1 second. ITI = Inter-Trial Interval, ISI = Inter-Stimulus Interval.

Computational model

We developed a cognitive Bayesian model of infants' looking behaviour that exploits the richness of eye-tracking data to infer the values of latent parameters capturing how infants processed the information content of the stimuli over time. To do so, we engaged in a 'reverse-engineering' procedure: given the data, we wanted to infer information about infants' cognitive processes (i.e., the latent parameters of the model). A consequence of this approach is that we were not constrained to the use of only one dependent variable, as is often the case in statistical models. Instead, we exploited multiple dependent variables to improve the model estimates of the latent parameters (e.g., Lee & Wagenmakers, 2014). Specifically, we collected three variables from the infants' looking behaviour that have been shown to relate to information processing (Kidd, Piantadosi, & Aslin, 2012; O'Reilly et al., 2013; Poli et al., 2020): 1) *Look-Away*. For each trial, we recorded whether infants kept looking at the screen or looked away; 2) *Saccadic Latency*. We measured how quickly infants moved their eyes from the cue to the target location, from the moment the target appeared. Negative times (i.e., anticipations to the target location) were also possible; 3) *Looking Time*. We measured how long infants looked at the target location, from the moment it appeared to 750ms after its disappearance.

Bayes-optimal learning. In every trial t of a sequence s , a stimulus is shown in the target location $x_{s,t} \in \{1, \dots, K\}$, where $K = 4$. Starting from the initial uniform prior $\gamma_s = [1, 1, 1, 1]$, which assumes that the target is equally likely to appear in any of the four locations (i.e., 25%), the probability $P(X)_{s,t}$ of seeing the stimulus in any given location is updated in light of the new evidence $x_{s,t}$:

$$P(X)_{s,t} = \frac{X_{s,t} + \gamma_s}{t + K}$$

where $X_{s,t}$ indicates all the evidence that has been observed up until trial t . From these probabilities, we quantified how much information the stimulus was conveying using the Kullback-Leibler (KL) Divergence (Kullback & Leibler, 1951), or D_{KL} :

$$D_{KL} = \sum P(X)_{s,t} \log \left(\frac{P(X)_{s,t}}{P(X)_{s,t-1}} \right)$$

KL-Divergence captures how much the new stimulus changed the probability distribution of the events or, in other words, how much information is gained from the new stimulus to understand the underlying (hidden) probability distribution of the events. These estimates are reported in Figure 2.

Figure 2. Estimates of information gain for all the sequences of the visual learning task. These estimates were used as input in the model, and the data from infants (looking times, saccadic latencies, and look-aways from the screen) informed the model about whether information was up- or -down weighted across time. Colours indicate the four target locations (pseudo-randomized across participants).

To test our hypotheses on up- and down-weighting of information, we introduced an exponential decay of information gain over trials (see Figure 3, in green). The decay could vary across sequences, thus allowing information on early sequences to be processed differently from information acquired later in the task. The decay was regulated by four additional parameters λ_α^0 , λ_α^1 , β_α^0 , and β_α^1 , such that:

$$IG_{s,t} = \underbrace{(\lambda_\alpha^0 + \lambda_\alpha^1 s)}_{\text{Early up-weight}} \underbrace{\exp(-t(\beta_\alpha^0 + \beta_\alpha^1 s))}_{\text{Late down-weight}} D_{KL_{s,t}}$$

where λ_α^0 and λ_α^1 regulate the up-weighting across sequences of the information acquired in trials early in the sequence, while β_α^0 and β_α^1 regulate the down-weighting across sequences of the information acquired in trials late in the sequence. The parameters λ_α^0 and β_α^0 are simple intercepts, while λ_α^1 and β_α^1 are the regression coefficients of interest, and our focus will thus be on the latter. If meta-learning occurs, we would expect that as sequences are observed, infants start to extract the underlying informational structure. As a consequence, the more sequences they are exposed to, the more they would up-weight the information acquired in the first trials or down-weight the information of late trials of each sequence. These two hypotheses are not mutually exclusive: Infants might both up-weight early evidence and down-weight late evidence. This would translate in the model returning positive values for both λ_α^1 and

β_{α}^1 . Finally, the parameter λ_s controls for changes in baseline attention to the task across sequences (see Figure 3, in red). This allows us to disentangle whether changes in attention across sequences are in fact due to meta-learning or simply related to fatigue.

We modelled how the information content in the stimuli was related linearly to the different observed variables, Saccadic Latency (SL), Looking Time (LT), and Look-Away (LA), in a probabilistic way. When estimating the relationship between information gain and the dependent variables, the regression coefficients (see Figure 3, in yellow) are estimated for each participant, thus taking into account individual differences. Further specifications of the priors and likelihoods of the model are available in the online supplementary materials.

Inference. The model was fitted in Python using PyMC3 (Salvatier et al., 2016). Two chains of 400,000 samples each were run using No-U-Turn-Sampler (NUTS), a common gradient-based MCMC algorithm. 390,000 samples were discarded as burn-in, and the last 10,000 samples were kept for further analysis. All the parameters reached convergence, as indicated by $\hat{R} < 1.05$ for all parameters (see Gelman & Rubin, 1992).

Figure 3. The generative model. Squares indicate discrete variables and circles indicate continuous variables; Grey boxes indicate observed variables; In green, parameters regulating the up- and down-weighting of information; In yellow, parameters accounting for individual differences across infants; In red, the parameter controlling for infants' changing attention across sequences.

Results

Model comparison. We compared the full model to reduced models that did not include some (or all) of the parameters modulating the up and down-regulation of information. The up-weight model assumed that, as sequences accumulated, information early in the sequences was progressively up-weighted. The down-weight model assumed that, as sequences accumulated, information late in the sequences was progressively down-weighted. The full model assumed concurrent up- and down-weighting of information, while the no-weight model assumed no change across sequences in the way information gain decayed over trials. The negative expected log predictive densities (ELPD) of the models were compared with Leave-One-Out (LOO) and WAIC methods.

The results of model comparison are reported in Table 1. As expected, LOO and WAIC methods returned similar estimates for all of the dependent variables. When comparing the total scores of the different models, the down-weight model and the full model performed better than the null model and the up-weight model, as indexed by lower negative ELPD scores (Kruschke, 2011). The difference in performance between the down-weight model and the full model was negligible. The up-weight parameter did not affect the results in either of the two models: it was absent in the down-weight model, and it was not different from zero for the full model (see below and in Figure 5). Hence, both models support the conclusion that down-weighting is present, while up-weighting is not.

Information gain predicted infants' looking behaviour. Before analysing the parameters of interest (i.e., the up- and down-weight parameters), we examined whether the dependent variables were significantly related to the information gain estimates. This step is crucial as we assumed that the dependent variables can inform latent parameters related to how information is processed. We found that the beta coefficients relating the dependent variables to information gain were different from zero for look-aways (mean = -4.98, SD = 0.61, 94%HDI = [-6.21, -3.87]), looking time (mean = -0.57, SD = 0.20, 94%HDI = [-0.99, -0.20]) and saccadic latency (mean = 0.69, SD = 0.19, 94%HDI = [0.34, 1.05]), and none of the credible intervals included zero. Hence, the “reverse-engineering” procedure was successful, as the dependent variables were all contributing to inform the levels of information content and their change over time.

To check whether our model was a good fit to the data, we performed a posterior predictive check. In other words, the values of the posterior distributions were used to predict new data, thus testing whether they could produce reliable estimates. The real data and the predicted distribution of the data matched well (Figure 4).

Table

1

Model comparison between the full model and different reduced models, for each dependent variable and as a whole.

Model	Dependent Variable	-ELPD LOO	Δ LOO	-ELPD WAIC	Δ WAIC
Null model	Looking Time	4686		4685	
	Saccadic Latency	4555		4555	
	Look-Away	1201		1201	
	Total	10 442	0	10 441	0
Up-weight	Looking Time	4683		4683	
	Saccadic Latency	4552		4552	
	Look-Away	1205		1205	
	Total	10 440	2	10 440	1
Down-weight	Looking Time	4675		4674*	
	Saccadic Latency	4549		4549*	
	Look-Away	1194		1194	
	Total	10 418	24	10 418*	23
Full Model	Looking Time	4675		4675	
	Saccadic Latency	4549		4549	
	Look-Away	1192		1192	
	Total	10 416*	26*	10 416*	25*

Note: Delta values indicate the difference from the null model. Asterisks indicate the best models.

Figure 4. Results for the coefficients relating information gain to looking behaviour. A-D. The peak of the distributions indicate the most likely value for the parameter, and the rest of the curve indicates other possible values and their relative likelihood. E-F. The posterior predictive distributions for looking time and saccadic latency.

Infants meta-learn by down-weighting irrelevant information. Table values of the parameters λ_{α}^1 and β_{α}^1 accounted for the up- and down-regulation of information across sequences. The up-weight parameter (λ_{α}^1) was not different from zero (mean = 0.004, SD = 0.003, 94%HDI = [0, 0.009]), confirming that there was not a significant relationship between up-weighting of information found early in a sequence and sequence number. The parameter β_{α}^1 was different from zero (mean = 0.065, SD = 0.019, 94%HDI = [0.031, 0.102]), confirming that there was a significant relationship between down-weighting of information found late in a sequence and sequence number (Figure 5).

Figure 5. The posterior distributions for λ_{α}^1 and β_{α}^1 . The peaks of the distributions indicate the most likely values for the parameters, and the rest of the curve indicates other possible values and their relative probability. The parameters λ_{α}^1 and β_{α}^1 regulate the up- and down-weighting of information gain, respectively.

To better understand how the parameters capture infants' attribution of informativity to incoming stimuli, we used the mean values of the parameters estimated by the model that was fit with infants' gaze data (as reported in table 2) to generate the perceived values of information gain across the task. The results are reported in Figure 6. Comparing an early sequence to a late sequence, the information gain of the first trials did not change as the sequences progressed, while information gain of trials later on in the sequence was down-weighted for late sequences. This effect was present while controlling for changes in baseline attention across sequences (see λ_s in Figure 4).

Figure 6. Information gain values generated from the model fit to infant data. The information gain that the infants attributed to the stimuli in the 1st and 5th sequence is reported. On average, infants watched 7 sequences. Compared to the blue line, the orange line is not higher at the start of the sequence (information is not up-weighted) but it is lower in following trials (information is down-weighted).

Infants’ meta-learning improves learning efficiency. Once we identified the mechanisms underlying infants’ meta-learning, it remains to be determined whether meta-learning led to more efficient and faster learning. To test this, we analysed saccadic latencies for predictable trials. When a target repeatedly appears in the same predictable location, saccadic latencies to the target should decrease. Moreover, in later sequences, if infants are exploiting what they meta-learned from previous sequences, saccadic latencies should be even shorter from earlier on. We analysed saccadic latencies as a function of sequence number and (log-transformed) trial number. The model was run on R with brms (Bürkner, 2017) with two chains of 8000 samples, and the first 4000 samples were discarded as burn-in. We found a main effect of trial number ($\beta = -0.43$, $SD = 0.05$, $94\%HDI = [-0.53, -0.33]$), indicating that they were learning the predictable target locations successfully. More importantly, we found an interaction between sequence number and trial number ($\beta = 0.04$, $SD = 0.01$, $94\%HDI = [0.02, 0.06]$), indicating that learning was faster for later sequences, just as predicted by our meta-learning results (Figure 7).

Figure 7. Saccadic latencies as a function of sequence and trial number. Saccadic latencies for predictable trials decreased across time, indicating a learning effect. This effect changed across sequences, as later sequences presented faster saccadic latencies from early on, suggesting that meta-learning led to more efficient learning.

Discussion

Meta-learning is a complex ability that entails building and tuning meta-parameters over a task to shape future expectations. Whereas it was known that infants' learning is incremental and that infants build new knowledge relying on what they have learned in the past (Thiessen & Saffran, 2007), the cognitive mechanisms underlying meta-learning were still unknown. In the current study, we show that infants can meta-learn from sparse evidence on a relatively small timescale (i.e., within minutes). More importantly, computational modelling demonstrates *how* they achieve this meta-learning: Infants identify where in the environment information is, and they use this knowledge to down-weight the informativity of irrelevant stimuli. As a result of meta-learning, infants needed less information and their learning efficiency increased, as indexed by faster saccadic latencies. These results were obtained while controlling for the effect of fatigue and for the passage of time, and thus cannot be explained by a simple decrease in interest in the stimuli over time.

The idea that infants learn inductive biases that change their information processing over time is in line with the predictive processing framework. Predictive processing holds that the brain is a hierarchically-structured prediction engine, with prediction travelling from higher to lower levels in the hierarchy, and prediction errors moving in the opposite direction, from lower to higher levels (Clark, 2013). In predictive processing terms, prediction errors are precision weighted, that is, their importance is modulated depending on their relevance. The down-weighting of information content as we find it is

assimilable to the predictive processing concept of precision-weighting of the prediction error. Thus, the current study contributes to recent literature on infant development suggesting that the architecture of a predictive brain is present from early on in life (Köster et al., 2020).

The ability to meta-learn is crucial for learning optimization and the acquisition of new abilities (Yoon et al., 2018). Its presence in the first year of life is in line with the idea that complex functions can emerge dynamically from the interaction between powerful learning mechanisms (Chen & Westermann, 2018; Poli et al., 2020; Silverstein et al., 2021), fundamental attentional biases (Di Giorgio et al., 2017), and early accumulation of evidence (Craighero et al., 2020; Ossmy & Adolph, 2020). Recent theoretical and empirical work favours the idea that high-level cognitive systems can be generated via sub-symbolic dynamics (Piantadosi, 2021; Sheya & Smith, 2019). For example, Yuan and colleagues (2020) investigated how preschoolers learn to map multi-digit number names onto their written forms. They presented children and a deep learning neural network with minimal training material and showed that both children and machines could reach systematic generalizations from sparse evidence. Hence, complex structures can emerge from simple mechanisms, and meta-learning might greatly ease this process from a very young age.

Finally, the current work offers a new methodological approach to infant research. Usually, statistical models try to link one or more independent variables to a single dependent variable by fitting a number of parameters (Lee, 2018). Conversely, in the computational model that we designed, the values of latent parameters were informed by multiple sources of data. In our case, saccadic latency, looking time to the targets, and looking away from the screen contributed to informing the relevant latent parameters. This approach is especially useful in infant research, which is often characterized by smaller datasets with lower signal-to-noise ratios compared to adult research (Havron et al., 2020). The limitations inherent to infant research can thus be compensated using richer measurements and dedicated computational models.

Acknowledgements

Funding: Radboud University, Nijmegen supported this work. Additional funding was from Donders Centre for Cognition internal grant to S.H. and R.B.M. (“Here’s looking at you, kid.” A model-based approach to interindividual differences in infants’ looking behaviour and their relationship with cognitive performance and IQ; award/start date: 15 March 2018), BBSRC David Phillips Fellowship to R.B.M. (“The comparative connectome”; award/start date: 1 September 2016; serial number: BB/N019814/1), Netherland Organization for Scientific Research NWO to R.B.M. (“Levels of social inference: Investigating the origin of human uniqueness”; award/start date: 1 January 2015; serial number: 452-13-015), Wellcome Trust center grant to R.B.M. (“Wellcome Centre for Integrative

Neuroimaging”; award/start date: 30 October 2016; serial number: 203139/Z/16/Z), EPA Cephalosporin Fund and Biotechnology and Biological Sciences Research Council (BB/N019814/1) grants to R.B.M. The Wellcome Centre for Integrative Neuroimaging is supported by core funding from the Wellcome Trust [203139/Z/16/Z]; Dutch Research Council NWO to S.H. (“Loving to learn - How curiosity drives cognitive development in young children”; start date: 1 February 2021; serial number: VI.C.191.022).
Competing interests: The authors declare that they have no competing interests. **Data availability statement:** Experimental data and computational models can be found on OSF: https://osf.io/a93qr/?view_only=864b5f415f5a40de82b578392ff59219

References

- Alet, F., Schneider, M. F., Lozano-Perez, T., & Kaelbling, L. P. (2020). Meta-learning curiosity algorithms. In *arXiv [cs.LG]*. arXiv. <http://arxiv.org/abs/2003.05325>
- Baram, A. B., Muller, T. H., Nili, H., Garvert, M. M., & Behrens, T. E. J. (2021). Entorhinal and ventromedial prefrontal cortices abstract and generalize the structure of reinforcement learning problems. *Neuron*, 109(4), 713-723.e7.
- Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, 121(1), 127–132.
- Chen, Y.-C., & Westermann, G. (2018). Different novelties revealed by infants’ pupillary responses. *Scientific Reports*, 8(1), 9533.
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, 36(3), 181–204.
- Craighero, L., Ghirardi, V., Lunghi, M., Panin, F., & Simion, F. (2020). Two-day-old newborns learn to discriminate accelerated-decelerated biological kinematics from constant velocity motion. *Cognition*, 195, 104126.
- Dewar, K. M., & Xu, F. (2010). Induction, overhypothesis, and the origin of abstract knowledge: Evidence from 9-month-old infants. *Psychological science*, 21(12), 1871-1877.
- Di Giorgio, E., Lunghi, M., Simion, F., & Vallortigara, G. (2017). Visual cues of motion that trigger animacy perception at birth: the case of self-propulsion. *Developmental Science*, 20(4). <https://doi.org/10.1111/desc.12394>

- Emberson, L. L., Richards, J. E., & Aslin, R. N. (2015). Top-down modulation in the infant brain: Learning-induced expectations rapidly affect the sensory cortex at 6 months. *Proceedings of the National Academy of Sciences of the United States of America*, 112(31), 9585–9590.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 7, 457–511.
- Gershman, S. J., & Niv, Y. (2010). Learning latent structure: carving nature at its joints. *Current Opinion in Neurobiology*, 20(2), 251–256.
- Grant, E., Finn, C., Levine, S., Darrell, T., & Griffiths, T. (2018). Recasting Gradient-Based Meta-Learning as Hierarchical Bayes. In *arXiv [cs.LG]*. arXiv. <http://arxiv.org/abs/1801.08930>
- Harlow, H. F. (1949). The formation of learning sets. *Psychological Review*, 56(1), 51–65.
- Havron, N., Bergmann, C., & Tsuji, S. (2020). Preregistration in infant research-A primer. *Infancy: The Official Journal of the International Society on Infant Studies*, 25(5), 734–754.
- Hessels, R. S., Niehorster, D. C., Kemner, C., & Hooge, I. T. C. (2017). Noise-robust fixation detection in eye movement data: Identification by two-means clustering (I2MC). *Behavior Research Methods*, 49(5), 1802–1823.
- Hunnius, S. (2022). Early cognitive development: Five lessons from infant learning. *Oxford Research Encyclopedia of Psychology*.
- Ibrahim, J. G., Chen, M. H., Sinha, D. (2001). *Bayesian survival analysis (Vol. 2)*. New York: Springer.
- Kayhan, E., Hunnius, S., O'Reilly, J. X., & Bekkering, H. (2019). Infants differentially update their internal models of a dynamic environment. *Cognition*, 186, 139–146.
- Kemp, C., Perfors, A., & Tenenbaum, J. B. (2007). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*, 10(3), 307–321.
- 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), 1–8.
- Koch, M. B., & Meyer, D. R. (1959). A relationship of mental age to learning-set formation in the preschool child. In *Journal of Comparative and Physiological Psychology* (Vol. 52, Issue 4, pp. 387–389). <https://doi.org/10.1037/h0043510>

- Köster, M., Kayhan, E., Langeloh, M., & Hoehl, S. (2020). Making Sense of the World: Infant Learning From a Predictive Processing Perspective. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 15(3), 562–571.
- Kruschke, J. K. (2011). *Doing Bayesian Data Analysis: A Tutorial with R and BUGS*. Academic Press.
- Kullback, S., & Leibler, R. A. (1951). On Information and Sufficiency. *Annals of Mathematical Statistics*, 22(1), 79–86.
- Lake, B. M., Salakhutdinov, R., & Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332–1338.
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *The Behavioral and Brain Sciences*, 40, e253.
- Lee, M. D. (2018). Bayesian methods in cognitive modeling. *The Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*, 5, 37–84.
- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian Cognitive Modeling: A Practical Course*. Cambridge University Press.
- Ossmy, O., & Adolph, K. E. (2020). Real-Time Assembly of Coordination Patterns in Human Infants. *Current Biology: CB*, 30(23), 4553–4562.e4.
- O'Reilly, J. X., Schüffelgen, U., Cuell, S. F., Behrens, T. E., Mars, R. B., & Rushworth, M. F. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. *Proceedings of the National Academy of Sciences*, 110(38), E3660–E3669.
- Passingham, R. E., & Wise, S. P. (2012). *The Neurobiology of the Prefrontal Cortex: Anatomy, Evolution, and the Origin of Insight*. OUP Oxford.
- Piantadosi, S. T. (2021). The computational origin of representation. *Minds and Machines*, 31, 1–58.
- Poli, F., Serino, G., Mars, R. B., & Hunnius, S. (2020). Infants tailor their attention to maximize learning. *Science Advances*, 6(39). <https://doi.org/10.1126/sciadv.abb5053>
- Rodríguez, G. (2008). Multilevel Generalized Linear Models. In J. de Leeuw & E. Meijer (Eds.), *Handbook of Multilevel Analysis* (pp. 335–376). Springer New York.
- Rule, J. S., Tenenbaum, J. B., & Piantadosi, S. T. (2020). The Child as Hacker. *Trends in Cognitive Sciences*, 24(11), 900–915.

- Salvatier, J., Wiecki, T. V., & Fonnesebeck, C. (2016). Probabilistic programming in Python using PyMC3. *PeerJ. Computer Science*, 2(e55), e55.
- Sheya, A., & Smith, L. (2019). Development weaves brains, bodies and environments into cognition. *Language, Cognition and Neuroscience*, 34(10), 1266–1273.
- Silverstein, P., Feng, J., Westermann, G., Parise, E., & Twomey, K. E. (2021). Infants learn to follow gaze in stages: Evidence confirming a robotic prediction. *Open Mind*, 1–15.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science*, 331(6022), 1279–1285.
- Thiessen, E. D., & Saffran, J. R. (2007). Learning to learn: Infants’ acquisition of stress-based strategies for word segmentation. *Language learning and development*, 3(1), 73-100.
- Trainor, L. J. (2012). Predictive information processing is a fundamental learning mechanism present in early development: evidence from infants. *International Journal of Psychophysiology: Official Journal of the International Organization of Psychophysiology*, 83(2), 256–258.
- Tummeltshammer, K. S., & Kirkham, N. Z. (2013). Learning to look: probabilistic variation and noise guide infants’ eye movements. *Developmental Science*, 16(5), 760–771.
- Vossel, S., Mathys, C., Daunizeau, J., Bauer, M., Driver, J., Friston, K. J., & Stephan, K. E. (2014). Spatial attention, precision, and Bayesian inference: a study of saccadic response speed. *Cerebral Cortex*, 24(6), 1436–1450.
- Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., Hassabis, D., & Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6), 860–868.
- Werchan, D. M., & Amso, D. (2020). Top-down knowledge rapidly acquired through abstract rule learning biases subsequent visual attention in 9-month-old infants. *Developmental Cognitive Neuroscience*, 42, 100761.
- Werchan, Denise M., Collins, A. G. E., Frank, M. J., & Amso, D. (2015). 8-month-old infants spontaneously learn and generalize hierarchical rules. *Psychological Science*, 26(6), 805–815.
- Xu, F. (2019). Towards a rational constructivist theory of cognitive development. *Psychological Review*, 126(6), 841–864.

- Yoon, J., Kim, T., Dia, O., Kim, S., Bengio, Y., & Ahn, S. (2018). Bayesian model-agnostic meta-learning. *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 7343–7353.
- Yuan, L., Xiang, V., Crandall, D., & Smith, L. (2020). Learning the generative principles of a symbol system from limited examples. *Cognition*, 200, 104243.

Supplementary Materials

Model priors. We report the prior distributions of all the parameters in the model. It must be noted that we always preferred vague priors (i.e., high variance), allowing the model to be predominantly informed by the data. The parameters λ_α^0 , λ_α^1 , β_α^0 , and β_α^1 are assumed to follow a half-Normal distribution, which constrains them to take only values equal or greater than zero:

$$\begin{aligned}\lambda_\alpha^j &\sim \text{HalfNormal}(0,1) & j &\in \{0,1\} \\ \beta_\alpha^j &\sim \text{HalfNormal}(0,1) & j &\in \{0,1\}\end{aligned}$$

This is a constrain due to the fact that KL-Divergence is a non-negative score, and thus it must scale with nonnegative terms. The half-Normal distribution was chosen to ease the interpretation of these parameters as regression coefficients. Similarly, hazard (which is proportional to a probability) cannot be negative, and thus the parameter λ_s has a Gamma distribution, which can only take positive values:

$$\lambda_s \sim \text{Gamma}(\kappa, \theta)$$

Where values of 0 would indicate that infants never look away from the screen, and the likelihood of looking away increases as λ values increase.. A different value of λ was estimated for each sequence s , and the general values across sequences were captured by the hyperpriors κ and θ :

$$\begin{aligned}\kappa &\sim \text{Gamma}(1,1) \\ \theta &\sim \text{Gamma}(1,1)\end{aligned}$$

We chose Gamma distributions following Ibrahim, Chen, and Sinha (2001). This prior requires the partitioning of the time range in question into intervals, and it thus fits well with the trial-by-trial structure of our design. All β_* for the regressions follow normal distributions, and they are thus free to detect both positive and negative relationships between the independent variable (i.e., information gain) and the data:

$$\begin{aligned}\beta_{LAi} &\sim \text{Normal}(\mu_{LA}, 1) \\ \beta_{SLi}^j &\sim \text{Normal}(\mu_{SL}, 1) & j &\in \{0,1\} \\ \beta_{LTi}^j &\sim \text{Normal}(\mu_{LT}, 1) & j &\in \{0,1\}\end{aligned}$$

The values of β_* were estimated for each participants, starting from the hyperparameters μ_* , that also followed a normal distribution:

$$\mu_* \sim \text{Normal}(0,1)$$

Model likelihood. The priors were linked to the data (i.e., the likelihoods) using a Student-t distribution for Saccadic latency (SL), and a normal distribution for Looking Time (LT); Look-Away (LA) is

modelled with a piecewise-constant proportional hazard model with a Poisson likelihood (Rodríguez, 2008):

$$\begin{aligned}
LA_{i,s,t} &\sim \text{Poisson}(\lambda_s \exp(\beta_{LA_i} IG_{i,s,t})) \\
SL_{i,s,t} &\sim \text{Norm}(\beta_{SL_i}^0 + \beta_{SL_i}^1 IG_{i,s,t}, \sigma_{SL}) \\
LT_{i,s,t} &\sim \text{T-Student}(\beta_{LT_i}^0 + \beta_{LT_i}^1 IG_{i,s,t}, \sigma_{LT}, v = 15)
\end{aligned}$$

Saccadic latency and looking time were standardized across participants and at the group level. The standard deviation terms σ_{SL} and σ_{LT} have half-Cauchy distributions, which again constrains their values to be nonnegative, as standard deviations cannot take negative values:

$$\begin{aligned}
\sigma_{SL} &\sim \text{halfCauchy}(0,1) \\
\sigma_{LT} &\sim \text{halfCauchy}(0,1)
\end{aligned}$$

We chose half-Cauchy distributions because they are recommended by Gelman (2006) for modelling standard deviations, given the advantage of being approximately uniform in the tail and weak near zero.