

The development of Bayesian integration in sensorimotor estimation

Claire Chambers

Department of Bioengineering,
University of Pennsylvania, Philadelphia, PA, USA
Department of Neuroscience,
University of Pennsylvania, Philadelphia, PA, USA



Taegh Sokhey

Sensory Motor Performance Program,
Shirley Ryan Abilitylab, Chicago, IL, USA
Department of Biological Sciences,
Northwestern University, Evanston, IL, USA



Deborah Gaebler-Spira

Sensory Motor Performance Program,
Shirley Ryan Abilitylab, Chicago, IL, USA
Department of Physical Medicine and Rehabilitation,
Feinberg School of Medicine, Northwestern University,
Chicago, IL, USA



Konrad Paul Kording

Department of Bioengineering,
University of Pennsylvania, Philadelphia, PA, USA
Department of Neuroscience,
University of Pennsylvania, Philadelphia, PA, USA



Examining development is important in addressing questions about whether Bayesian principles are hard coded in the brain. If the brain is inherently Bayesian, then behavior should show the signatures of Bayesian computation from an early stage in life. Children should integrate probabilistic information from prior and likelihood distributions to reach decisions and should be as statistically efficient as adults, when individual reliabilities are taken into account. To test this idea, we examined the integration of prior and likelihood information in a simple position-estimation task comparing children ages 6–11 years and adults. Some combination of prior and likelihood was present in the youngest sample tested (6–8 years old), and in most participants a Bayesian model fit the data better than simple baseline models. However, younger subjects tended to have parameters further from the optimal values, and all groups showed considerable biases. Our findings support some level of Bayesian integration in all age groups, with evidence that children use probabilistic quantities less efficiently than adults do during sensorimotor estimation.

Introduction

The behavior of adults under uncertainty is well described by Bayesian inference, in that adult humans weight different sources of information according to their relative uncertainty. Behavior is consistent with Bayesian computations in sensorimotor behavior (Berniker, Voss, & Kording, 2010; Kording & Wolpert, 2004), perception (Knill & Richards, 1996; Mamassian & Goutcher, 2001), cognition and reasoning tasks (Battaglia, Hamrick, & Tenenbaum, 2013; Tenenbaum & Griffiths, 2001), and cue combination across and within sensory modalities (Ernst & Banks, 2002; Hillis, Watt, Landy, & Banks, 2004). Adult humans seem to integrate information in a way that is well predicted by Bayesian inference (but see Bowers & Davis, 2012; Jones & Love, 2011; Rahnev & Denison, 2018).

These numerous findings of Bayesian behavior have led to the theory that the underlying neural computations are inherently Bayesian. For example, it has been argued that the activity of neural populations reflects probabilistic population codes that directly implement

Citation: Chambers, C., Sokhey, T., Gaebler-Spira, D., & Kording, K. P. (2018). The development of Bayesian integration in sensorimotor estimation. *Journal of Vision*, 18(12):8, 1–16, <https://doi.org/10.1167/18.12.8>.

<https://doi.org/10.1167/18.12.8>

Received January 11, 2018; published November 15, 2018

ISSN 1534-7362 Copyright 2018 The Authors



Bayesian computations (Beck et al., 2008; Ma, Beck, Latham, & Pouget, 2006; Ma, Beck, & Pouget, 2008; Zemel, Dayan, & Pouget, 1998). However, findings of Bayesian behavior are not sufficient to support this claim. Bayesian behavior simply represents optimal behavior under uncertainty, and there are ways of generating optimal behavior that do not explicitly implement Bayesian computation (Mandt, Hoffman, & Blei, 2017; Verstynen & Sabes, 2012; Weisswange, Rothkopf, Rodemann, & Triesch, 2011). Therefore, previous research has not fully established whether the neural code is inherently Bayesian.

If neural circuits are evolved to implement Bayesian computations, then behavior should always show Bayesian signatures, including during development. Therefore, children too should act in accordance with the rules of Bayesian integration. Specifically, they should weight information according to its relative uncertainty in simple tasks.

A good number of articles ask how Bayesian children are. Work on the development of cue combination across and within modalities has often shown that children do not integrate information but instead process it from each cue separately up to the age of approximately 9–11 years (Dekker et al., 2015; Gori, Del Viva, Sandini, & Burr, 2008; Nardini, Bedford, & Mareschal, 2010; Nardini, Jones, Bedford, & Braddick, 2008). Work on the integration of value-based information (reward/penalty) in a sensorimotor task has shown that children's behavior is suboptimal until late in development, at around 11 years (Dekker & Nardini, 2016). However, other findings suggest that whether children integrate information or not may be task dependent. In a hand-localization task, variance reduction consistent with multisensory integration has been reported in the responses of young children ages 4–6 years (Nardini, Begus, & Mareschal, 2013). Some work on looking times is consistent with optimal integration of information as early as infancy (Téglás, Tenenbaum, & Bonatti, 2011). It has been shown that children as young as 4 years old use probabilistic information to infer causality when performing actions (Gopnik & Wellman, 2013; Kushnir & Gopnik, 2007; Sobel, Tenenbaum, & Gopnik, 2004). Therefore, based on previous research, it is unclear whether the behavior of children is consistent with use of Bayesian inference.

Here, we investigate whether integration of information to perform sensorimotor estimation is present in young children or is acquired over the course of development. In our paradigm, we examined the use of probabilistic information in a simple sensorimotor estimation task, previously used in adults to examine integration of prior and likelihood information under uncertainty (Acuna, Berniker, Fernandes, & Kording, 2015; Berniker et al., 2010; Kording & Wolpert, 2004; Vilares, Howard, Fernandes, Gottfried, & Kording,

2012). The results of these studies show that adults learn experimentally imposed prior distributions and integrate prior statistics with uncertain sensory information efficiently and without instructions to do so.

Our task was designed to specifically probe the integration of prior and likelihood information during sensorimotor estimation. We isolated integration from the many other factors that contribute to movement execution under naturalistic conditions. In the real world, priors and likelihoods are complex and multi-dimensional. In order to discount the influence of the complexity of the distribution and task, we used simple unidimensional Gaussian distributions (Berniker et al., 2010). We also minimized the influence on our results of having to learn the prior distribution by presenting samples from the prior distribution on-screen. Movement in the real world has an associated loss function (Wolpert & Landy, 2012). In this task, we minimized the effect of motor effort: All possible responses could be made using minimal movement of a computer mouse. Our experiment thus isolated Bayesian integration from other processes.

In our experiment, visual targets were drawn from a prior distribution and participants were shown uncertain sensory information about each target. We measured reliance on sensory information and found that while young children ages 6–8 years used sensory information according to its precision, we did not find evidence that they used the uncertainty of the prior distribution, as did children ages 9–11 years or adults. In all age groups, a Bayesian model predicted estimation behavior better than models which used only one source of information (prior or likelihood) or switched between prior and likelihood (Laquitaine & Gardner, 2018).

Methods

Experimental details

We aimed to examine probabilistic inference during sensorimotor estimation in a child population. Our task was designed to examine use of probabilistic information during sensorimotor estimation (Acuna et al., 2015; Berniker et al., 2010; Vilares et al., 2012). Previous findings indicate that adults weight information according to its reliability and learn priors in a manner which resembles Bayesian integration during sensorimotor estimation. For the purposes of the current study, we adapted the experimental protocol for child participants by using a concept that was engaging to children, using simplified instructions, and reducing the number of trials.

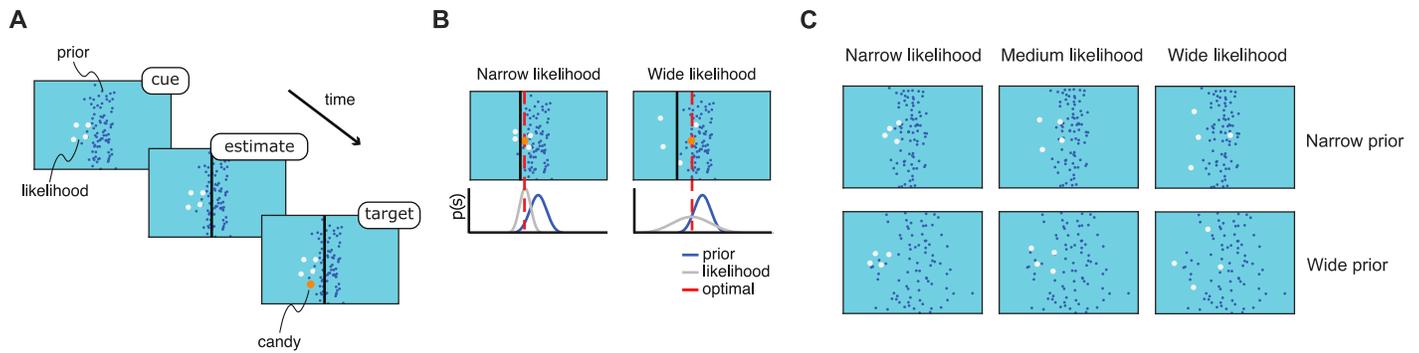


Figure 1. (A) Experimental protocol. Participants were shown a visual cue with experimentally controlled uncertainty or likelihood which was presented as a “splash” created by a hidden target or “piece of candy” drawn from a prior distribution. Participants were told that the splash was created by candy falling into a pond. They were prompted to place a vertical bar (“net”) where the hidden target fell and were then shown feedback on the target location. (B) Relying on the likelihood. A simple strategy would be to rely entirely on likelihood information by pointing at its centroid on each trial. This strategy is close to optimal when the likelihood is precise or narrow. The black bar or net overlaps with the target in the left panel. However, this strategy is less successful when the likelihood is wider, as samples from the likelihood become a less reliable indicator of target location and the optimal estimate shifts closer to the prior mean. The net is far from the target in the right panel. The optimal strategy involves weighting prior and likelihood information according to their relative uncertainties. (C) Experimental design. In order to quantify integration of the prior and likelihood, we measured reliance on the likelihood under different conditions of prior width and likelihood width. The prior could be narrow or wide, and the likelihood could be narrow, medium, or wide.

Participants were 16 children (eight boys, eight girls) ages 6–8 years ($M = 6.94$, $SD = 0.77$), 17 children (eight boys, nine girls) ages 9–11 years ($M = 10.06$, $SD = 0.75$), and 11 adults (five men, six women) over 18 years old ($M = 27.27$, $SD = 5.31$). The data of two participants were excluded due to their looking away from the screen during the experiment. Participants overlapped with the sample of a previous study, where some participants from the current study were included as age-matched controls and compared with a clinical population (Chambers, Sokhey, Gaebler-Spira, & Kording, 2017).

In a quiet room, participants sat at a comfortable distance (typically between 30 and 60 cm) from a computer monitor 52 cm wide and 32.5 cm high. Before starting the experiment, we presented participants with the instructions that someone behind them was throwing pieces of candy into a pond, represented by the screen, and that their aim was to estimate where the candy target landed and catch as many pieces of candy as possible over the course of the experiment. We told participants that one piece of candy was being thrown into the water per trial. We also told participants that the candy target was hidden and caused a splash on the surface of the water. We showed them “where the candy usually lands.” The participants’ goal was to hit these targets (“catch the candy”) by combining the information from the splash and where the candy usually lands. The instructions from the beginning of the experiment are shown in the Appendix (Figures A1 and A2).

Candy targets were drawn from a Gaussian prior distribution centered at the middle of the screen, $N(\mu, \sigma_s^2)$. The prior distribution was fixed for the duration of a block. On each trial, participants were presented with an uncertain “splash” stimulus for 1 s and were told that the splash was caused by a hidden candy target (Figure 1A). The splash was $n = 4$ samples (white dots on a blue background, diameter = 2% screen width) from a Gaussian likelihood distribution that was centered on the target location $N(s, \sigma_l^2)$. The stimulus-generation process was hierarchical: The target on each trial was sampled from a prior distribution that was fixed across trials; the likelihood distribution was centered on the target; and samples were then drawn from the likelihood distribution to form the splash stimulus that was displayed to participants. Participants provided an estimate of the candy target’s location on the horizontal axis using a vertical bar that extended from the top to the bottom of the screen (2% of screen width). In order to successfully catch a target, the center of the “net” or vertical bar had to be within 2% of screen width of the target center. The net appeared at the same time as the splash at a random location on screen. Participants had 6 s to respond. After providing a response, they were shown the true candy-target location.

The stimuli in this experiment were presented well above the threshold of perception (in terms of their size and contrast) and were clearly visible to those with normal or corrected-to-normal vision. Therefore, physical measurements of stimulus dimensions like the exact visual angle were not crucial in determining how

participants weighted information from the prior and likelihood distributions in this experiment.

One simple strategy for performing sensorimotor estimation under uncertainty is to consistently judge target location at the centroid c of the splash. This consists of full reliance on the likelihood and works well when the likelihood distribution is narrow, because the closely spaced points of the splash are an accurate indicator of target location (Figure 1B, left). However, full reliance on the likelihood would cause a participant to miss targets more frequently as the likelihood distribution widens (Figure 1B, right). When sensory information is unreliable, rather than relying on the likelihood completely we maximize performance by giving more weight to our prior belief on target location. More generally, the optimal strategy involves weighting sources of information according to their relative uncertainties.

Formally, weighting sources of information according to their relative precision corresponds to Bayesian inference. An optimal Bayesian observer combines noisy sensory information from the likelihood $N(s, \sigma_1^2/n)$ with the prior $N(\mu, \sigma_s^2)$, resulting in a posterior distribution over target location:

$$N\left(\left(\frac{\mu}{\sigma_s^2} + \frac{c}{\sigma_1^2/n}\right) / \left(\frac{1}{\sigma_s^2} + \frac{1}{\sigma_1^2/n}\right), 1 / \left(\frac{1}{\sigma_s^2} + \frac{1}{\sigma_1^2/n}\right)\right).$$

The mean of the posterior is a mean of the prior and sensory information weighted by their precisions. From this posterior distribution, an estimate is computed. Therefore, the optimal reliance on the likelihood is a function of prior and likelihood uncertainties, $\sigma_s^2 / (\sigma_s^2 + \sigma_1^2/n)$. We can manipulate the prior and likelihood variance and measure their influence on participants' reliance on the likelihood to investigate probabilistic information during sensorimotor estimation.

To investigate how children use probabilistic information during sensorimotor estimation, we manipulated the variances of prior distribution and likelihood distribution (Figure 1C). We used a Gaussian prior distribution with a mean at the center of the screen and standard deviation of 0.03 (Narrow Prior) or 0.1 (Wide Prior) in units of screen width. The likelihood distribution was centered on the target location and could have a standard deviation of 0.05 (Narrow Likelihood), 0.1 (Medium Likelihood), or 0.25 (Wide Likelihood) in units of screen width. There were six conditions: Narrow Prior/Narrow Likelihood, Narrow Prior/Medium Likelihood, Narrow Prior/Wide Likelihood, Wide Prior/Narrow Likelihood, Wide Prior/Medium Likelihood, and Wide Prior/Wide Likelihood.

The experiment consisted of four blocks, each lasting 120 trials, preceded by a practice block lasting 10 trials. Trials were blocked by prior condition, with all likelihood conditions presented in randomized order

within one block. The prior over target location switched from block to block with a randomly chosen starting condition for each participant (i.e., Narrow-Wide-Narrow-Wide or Wide-Narrow-Wide-Narrow).

We introduced a number of modifications to engage child participants in the task. Participants were shown how much candy they had won on-screen and won a “bonus” piece of candy for every 10 pieces of candy they caught. Sounds were presented to signal successfully catching a target and missed responses when they did not respond within the 6-s time window. Step-by-step instructions were shown to participants on-screen before the experiment, to ensure that all participants received the same instructions. Before the experiment, we told all participants that their payment was proportional to the number of pieces of candy they caught. Their target score was displayed on-screen. Then at the end of the experiment we paid a fixed amount, which was different for the parents of child participants (\$25) and for adults (\$10). These modifications ensured that child participants were not discouraged during the experiment.

Ethical approval was provided by the Northwestern University Institutional Review Board (#20142500001072). This study was performed in accordance with the Declaration of Helsinki. Participants signed a consent form before participation. For child participants, a parent provided consent for their child to take part and completed the Developmental Coordination Disorder questionnaire (Wilson et al., 2009), a modified Vanderbilt questionnaire to assess for attention deficit–hyperactivity disorder (Wolraich et al., 2003), and the Behavior Assessment System for Children parent rating scales (Reynolds & Kamphaus, 2004). After the participant completed the game, we administered the child Mini-Mental State Evaluation to obtain an approximate assessment of cognitive ability (Ouvrier, Goldsmith, Ouvrier, & Williams, 1993). No data were excluded on the basis of neuropsychological test results.

Data analysis

We were interested in the integration of probabilistic information from a prior distribution and sensory information from the likelihood during sensorimotor estimation. To investigate this, we examined whether samples from the likelihood distribution, $X = \{x_1, x_2, x_3, x_4\}$, were combined with information about the prior distribution $N(\mu, \sigma_s^2)$ in producing an estimate of target location. We quantified this for each condition using the extent to which participants relied on the likelihood, given by the linear relationship between the centroid of the splash $c = \sum_i^N x_i/n$ and their estimate on each trial. We performed a linear

regression with the estimate (placement of the net) as the dependent variable and the likelihood centroid c as the independent variable. The slope provides an estimate of the reliance on the likelihood, which we term the *estimation slope*. If participants relied only on the likelihood to generate their estimate, then they should point close to the centroid of the splash c on all trials, leading to an estimation slope ≈ 1 . If instead participants ignore the likelihood and use only their representation of the prior, then their estimates should not depend on c , leading to an estimation slope ≈ 0 . The intercept/ $(1 - \text{estimation slope})$ computed from the fitted function reflects the subjective prior mean used to provide estimates. Therefore, from participants' estimates we obtain a measure of their reliance on the likelihood and the subjective prior mean.

For a Bayesian observer, sources of information are weighted according to their relative reliabilities, leading to an estimation slope of $\sigma_s^2 / (\sigma_s^2 + (\sigma_1 + \Delta_1)^2 / n)$. To account for the fact that participants may not have learned the experimentally imposed likelihood variance σ_1^2 , we added an extra parameter Δ_1 , a constant source of variance added to the likelihood variance. We use this equation to model the estimation slopes and to infer probabilistic variables used by participants (σ_s^2 , Δ_1).

We also investigated a simple switching model as an alternative to full Bayesian integration, where participants randomly switched between using the prior and likelihood across trials. The proportion of trials where participants used the likelihood to form their estimate, $p(\text{likelihood})$, was a free parameter. In this model, participants ignored the uncertainty of prior and likelihood information and used a fixed $p(\text{likelihood})$ for all conditions.

We quantified task performance using the proportion of correct responses, $p(\text{correct})$, and the root mean square error (RMSE) with respect to the regression line in each condition. We note that the RMSE contains contributions from both additional noise Δ_1 added to the likelihood and motor errors.

Motor errors, which occur after the combination of task-relevant information, influence estimates. Linear regression provides an unbiased estimate of the estimation slope for realistic amounts of motor noise added to estimates with standard deviations of up to approximately 10%–15% of screen width. Larger amounts of motor noise lead to many estimates at screen limits and estimation slopes that are closer to zero than the true values predicted by prior and likelihood parameters. Here, our analysis of estimation slopes assumes that sensorimotor estimates are not corrupted by large amounts of motor noise. However, such large motor errors are unlikely given that the RMSE of the most variable subject was less than 15% of screen width (see Figure 2A, right panel).

We examined whether children's estimation behavior resembled Bayesian inference using model selection and by estimating the parameters of the best-fitting model. We first compared the performance of a Bayesian model with three alternative models: a model where participants alternated between the likelihood and prior (Switch model), a model where they fully relied on the prior (Prior-only model), and a model where they fully relied on the likelihood (Likelihood-only model). We computed estimation slopes from each participant's data, then fitted the estimation slopes to the Bayesian model and the Switch model. The Bayesian model contained two prior variance parameters and the variance parameter added to the likelihood ($\sigma_{s\text{NP}}^2$, $\sigma_{s\text{WP}}^2$, and Δ_1). We minimized the mean squared error (MSE) of the objective function $\sigma_s^2 / (\sigma_s^2 + (\sigma_1 + \Delta_1)^2 / n)$ relative to the data, using bounds of 0.02 and 0.4 for the prior variance parameters. For the Switch model, we minimized the MSE between estimation slopes and the $p(\text{likelihood})$ parameter, using bounds of 0.1 and 0.9. We ran the optimizer five times with randomly selected initial parameters. We performed leave-one-out cross validation on the data of each participant by fitting model parameters to five conditions and computing the MSE for the left-out condition. For the Prior-only model, the MSE was computed by comparing estimation slopes to 0; for the Likelihood-only model, it was computed by comparing estimation slopes to 1. We selected the model with the lowest MSE summed across left-out conditions. We estimated the parameters of the Bayesian and Switch models by fitting the models to 1,000 data sets resampled with replacement from each participant's data. We used estimation slopes measured from participants' data for model selection and parameter estimation.

For validation purposes, we performed model selection on 1,000 data sets simulated from each model. Each simulated participant generated estimates using a maximum a posteriori decision rule with added motor noise ($SD = 5\%$ in screen units). If estimates fell outside screen limits, they were set to the screen limit. Prior-only and Likelihood-only participants were generated by excluding the likelihood and prior terms from the posterior distribution, respectively. Switch subjects alternated between Prior-only and Likelihood-only strategies. In order to illustrate that we can infer model parameters with reasonable accuracy, we performed parameter estimation for 1,000 simulated Bayesian subjects whose subjective prior variance was at intermediate values between the theoretical variance parameters. We simulated cases where subjective prior variances were undifferentiated ($\sigma_{s\text{NP}}^2 = 0.06$, $\sigma_{s\text{WP}}^2 = 0.06$), where they were partly differentiated ($\sigma_{s\text{NP}}^2 = 0.048$, $\sigma_{s\text{WP}}^2 = 0.083$), and where subjects used the experimentally imposed prior ($\sigma_{s\text{NP}}^2 = 0.03$; $\sigma_{s\text{WP}}^2 = 0.1$). We simulated conditions where different amounts of

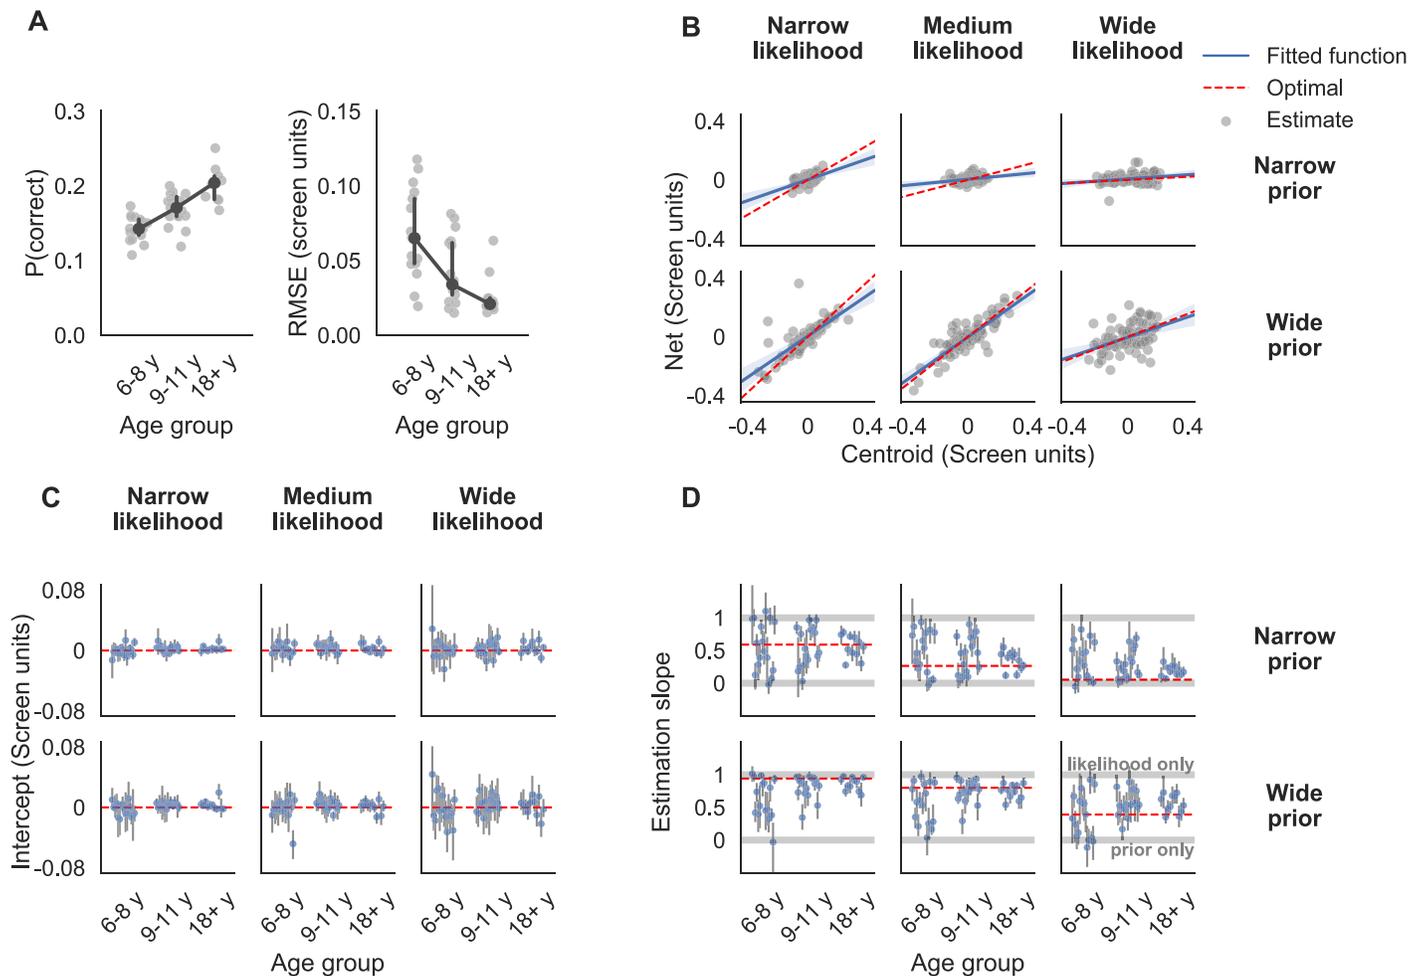


Figure 2. Task performance and estimation data. (A) Left: The proportion of candy targets caught as a function of age group (median, error bars = 95% confidence intervals [CIs]). Right: The root mean square error relative to the regression line in the Wide Prior/ Narrow Likelihood condition gives an indication of how noisy participants were and is shown as a function of age group (median, error bars = 95% CIs). (B) Estimation data overlaid with linear fit for a representative participant age 11 years. The net position as a function of the centroid of the likelihood is shown for each trial (points). The fitted (blue) and optimal (red) functions are displayed. Note that optimal values here and in (C–D) and Figure 3B assume that participants use the experimentally imposed likelihoods and priors. Each panel displays estimation data for one condition, as defined by prior and likelihood width. (C) The median bootstrapped intercept of individual participants is shown as a function of age group (error bars = 95% CI). The optimal intercept at zero is shown (red). (D) The median bootstrapped estimation slope of individual participants is shown as a function of age group (error bars = 95% CI). The optimal estimation slope values are shown (red).

variance were added to the experimentally imposed likelihood ($\Delta_l = 0, 0.05, 0.1$). We also inferred the $p(\text{likelihood})$ parameter of the Switch model (0.2, 0.4, 0.6, 0.8) from the data of 1,000 simulated subjects per condition. Simulations allowed us to ensure that our model-selection and parameter-estimation procedures produced unbiased results.

Results

We wanted to investigate the development of Bayesian integration. To do so, we examined whether

children ages 6–11 years and adults could learn to use uncertainty of different sources of information (prior and likelihood) during sensorimotor estimation (Figure 1). We first quantified participants' task performance. We then examined how their estimates depended on the prior and the likelihood. We then performed model selection on the data of each participant to assess whether their behavior was more consistent with Bayesian inference, full reliance on the prior, full reliance on the likelihood, or alternation between prior and likelihood. Finally, we devised estimated parameters of the Bayesian model for each age group.

It was first important to establish that the all age groups understood and carried out the task. We

therefore examined the proportion $p(\text{correct})$ of candy caught (Figure 2A, left panel). Performance increased significantly with age (one-way analysis of variance [ANOVA]), $F(2, 41) = 27.44$, $p < 0.0001$, with significant differences between age groups—6–8 years versus 9–11 years: $t(31) = 3.86$, $p < 0.001$; 6–8 years versus 18+ years: $t(25) = 7.72$, $p < 0.0005$; 9–11 years versus 18+ years: $t(26) = 3.83$, $p < 0.001$ (corrected $\alpha = 0.0167$). We compared the proportion of candy targets caught, $p(\text{correct})$, to chance level for each age group (Figure 2A). The performance of all age groups far exceeded chance level, as tested with one-sample t tests—6–8 years: $t(15) = 33.77$, $p < 0.0001$; 9–11 years: $t(16) = 32.26$, $p < 0.0001$; 18+ years: $t(10) = 29.31$, $p < 0.0001$ (corrected $\alpha = 0.0167$). In addition, we examined the RMSE of estimates relative to the regression line in the Narrow Likelihood/Wide Prior condition, which gave an indication of the variability of sensorimotor estimates (Figure 2A, right panel). We found a significant effect of age group on RMSE (one-way ANOVA), $F(2, 41) = 10.76$, $p < 0.0005$, and significant differences between age groups—6–8 years versus 9–11 years: $t(31) = 2.80$, $p < 0.01$; 6–8 years versus 18+ years: $t(25) = 4.31$, $p < 0.001$; 9–11 years versus 18+ years: $t(26) = 2.19$, $p = 0.0378$ (corrected $\alpha = 0.0167$). The RMSE (in units of screen width) was within an acceptably low range for all age groups—6–8 years: M (SD) = 0.07 (0.03); 9–11 years: 0.04 (0.02); 18+ years: 0.03 (0.01). This shows that all age groups understood the candy-catching task and carried out the task above chance level. Although we do observe differences between age groups, these differences between cannot be attributed to a lack of understanding of the task.

Having found differences in performance between groups, we examined how weighting of prior and likelihood information changes over the course of development. In order to quantify the nature of integration between the likelihood and the prior, we used the parameters of the linear fit between estimates and the likelihood centroid. The mean of the prior used by participants is calculated using the parameters of the linear fit: $\text{intercept}/(1 - \text{estimation slope})$ (Berniker et al., 2010). Full reliance on the likelihood indicates a close relationship between estimates and the likelihood and results in estimation slope ≈ 1 . Full reliance on the prior indicates a lack of relationship between estimates and the likelihood and results in estimation slope ≈ 0 . We display the fitted estimation slope, intercept, and estimation data for an 11-year old participant (Figure 2B), and the reliability of the slope and intercept estimates for each individual participant (Figure 2C and 2D). The linear fit explains a reasonable amount of variance in raw estimates in each age group—6–8 years: mean R^2 (SD) = 0.32 (0.22); 9–11 years: 0.54 (0.20); 18+ years: 0.70 (0.12). This procedure allowed us to

quantify the nature of prior and likelihood integration for each participant.

We were interested in whether sensorimotor integration of prior and likelihood improved during development. To examine this, we first tested whether participants' estimates were consistent with simple prior statistics. We analyzed the prior mean used by participants (Figure 3A). We note that this information was provided to participants on-screen, since they were shown samples from the prior distribution. However, since we were examining a child population, it was necessary to make sure that they used visually available information. We examined the influence of the prior width, likelihood width, and age group on the prior mean (Intercept/[1 – estimation slope]). The analysis of the prior mean did not reveal a significant influence of prior width or age group (repeated-measures ANOVA)—prior width: $F(1, 42) = 0.37$, $p = 0.55$; age group: $F(1, 42) = 2.53$, $p = 0.12$ —nor their interaction, $F(1, 42) = 0.76$, $p = 0.39$. However, there was a significant effect of likelihood width, $F(2, 84) = 3.94$, $p = 0.03$, and a significant Likelihood width \times Age interaction, $F(2, 84) = 3.69$, $p = 0.04$. Post hoc t tests comparing levels of the likelihood-width variable and the interaction between likelihood width and age group did not show significant differences between conditions (see Appendix, Tables A1 and A2). Therefore, age group and stimulus factors did not play a strong role in influencing the prior mean used by participants.

We next examined whether the weighting of prior and likelihood information changed during development by analyzing the estimation slope (Figure 3B). We examined the influence of the prior width, likelihood width, and age group on the estimation slope. A repeated-measures ANOVA applied to the estimation slope revealed significant main effects of prior width, $F(1, 42) = 5.66$, $p < 0.05$, and likelihood width, $F(2, 84) = 53.78$, $p < 0.0001$, as well as a nonsignificant main effect of age group, $F(1, 42) = 0.42$, $p = 0.52$. Therefore, there is evidence that the sample as a whole integrated both the prior and the likelihood into their judgments.

In our analysis of estimation slopes, we then turned to statistical interactions which reveal age-specific effects. There was a significant Age group \times Prior width interaction, $F(1, 42) = 26.34$, $p < 0.0001$, and no significant Age group \times Likelihood width interaction, $F(2, 84) = 3.39$, $p = 0.06$. The use of likelihood may not change considerably over the course of development. However, there is evidence for an influence of age group on use of prior statistics. We therefore examined the influence of the prior width on the estimation slope for different age groups, using paired t tests—6–8 years: $t(30) = 0.54$, $p = 0.59$; 9–11 years: $t(32) = 3.76$, $p < 0.001$; 18+ years: $t(20) = 7.08$, $p < 0.0001$ (corrected $\alpha = 0.0167$). At a group level, 6- to 8-year-olds do not distinguish between Narrow and Wide Prior condi-

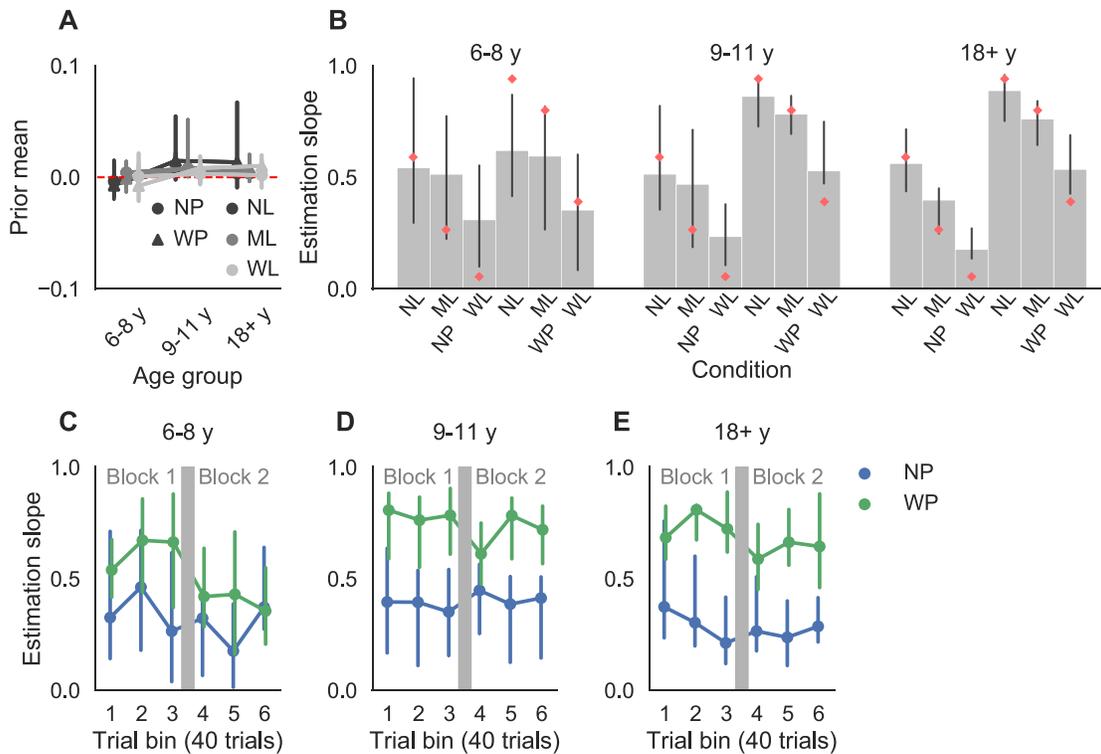


Figure 3. Prior mean, estimation slope, and estimation slope as a function of trial bin. (A) Median prior mean as a function of prior width (NP = narrow prior, WP = wide prior), likelihood width (NL = narrow likelihood, ML = medium likelihood, WL = wide likelihood), and age group (error bars = 95% confidence interval). The optimal value is shown (red). (B) Median estimation slope as a function of prior width, likelihood width, and age group (error bars = 95% confidence interval). Optimal values are shown (red). (C–E) Estimation slopes were computed for separate blocks and bins of 40 consecutive trials, then averaged across likelihood conditions. Median estimation slopes (error bars = 95% confidence) are shown for the three age groups: (C) 6–8 years, (D) 9–11 years, and (E) 18+ years.

tions, with this difference becoming significant at 9–11 years. Therefore, young children ages 6–8 years show the ability to incorporate likelihood into their judgments as adults do, but there is no evidence that they make use of the prior width until 9–11 years.

Children ages 6–8 years did not appear to use the prior distribution when making sensorimotor estimates. It could have been the case that they learned to do so over the course of the experiment. We therefore examined how the influence of the prior width on the estimation slope changed during the experiment. We computed slopes for bins of 40 consecutive trials and averaged estimation slopes across likelihood conditions (Figure 3C–3E). As shown by a repeated-measures ANOVA, we did not find significant main effects of age group, $F(1, 41) = 0.01$, $p = 0.94$; prior width, $F(1, 41) = 1.96$, $p = 0.17$; or trial bin, $F(2, 82) = 0.60$, $p = 0.55$; but we did find a significant main effect of experimental block, $F(1, 41) = 7.79$, $p < 0.01$. As before, we found a significant Prior width \times Age group interaction, $F(1, 41) = 14.58$, $p < 0.001$. However, interactions between age group, prior width, trial bin and experimental block were not significant. The significant effect of experimental block revealed a subtle tendency to rely more on the likelihood in Block 1— M (SD) = 0.55 (0.25)—than

in Block 2—0.46 (0.22). However, we did not find evidence for a change over the course of the experiment in how participants distinguished between conditions based on uncertainty. Therefore, we did not find evidence that young children learned the prior variance over the course of the experiment.

We formally tested whether estimation behavior is consistent with Bayesian inference by comparing the performance of a model which integrated prior and likelihood information (Bayesian model) with baseline models that used only prior information (Prior-only model), used only likelihood information (Likelihood-only model), and alternated between prior and likelihood across trials (Switch model). We first show that we can reliably infer the correct model from simulated estimation data (Figure 4A). When applied to estimation data, on average, we found a lower MSE for the Bayesian model compared to the baseline models (Figure 4B). Most individual participants' data were best fit by the Bayesian model: 11 out of 16 participants ages 6–8 years, 15 out of 17 ages 9–11 years, and 11 out of 11 adults (Figure 4C). On average, the Bayesian model accounts for a reasonable amount of variance of estimation slopes—6–8 years: mean R^2 (SD) = 0.74 (0.26); 9–11 years: 0.94 (0.07); 18+ years: 0.96

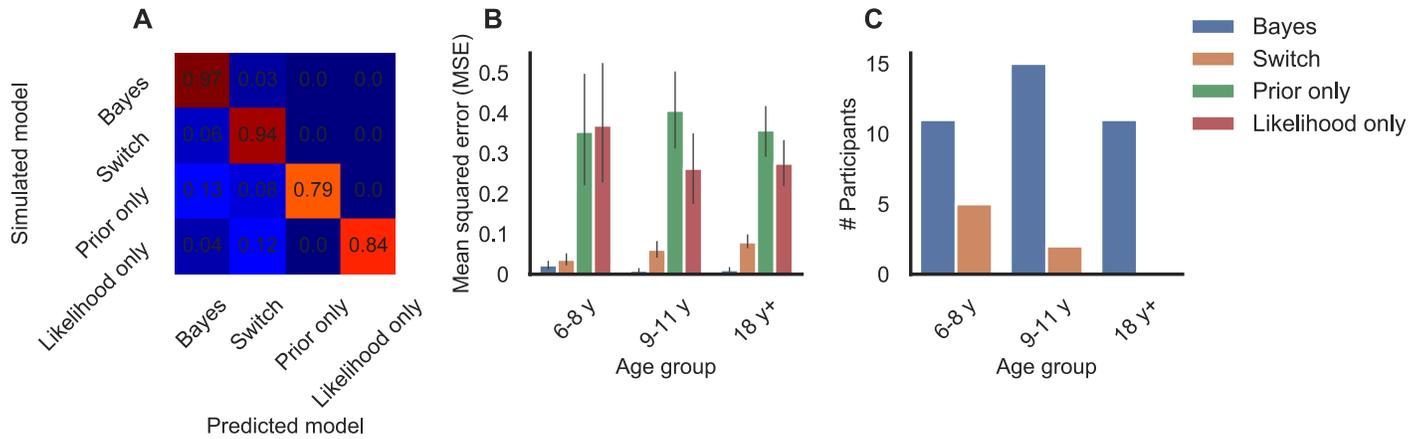


Figure 4. Model selection. (A) Confusion matrix showing the proportion of cases where each model was selected, computed from the data of 1,000 participants per simulated model. We can infer the correct model from simulated data with reasonable accuracy. (B) Median mean squared error for each model as a function of age group (error bars = 95% confidence interval). (C) The number of participants for whom each model was selected. The Bayesian model provides an improved fit for most participants (11 out of 16 ages 6–8 years, 15 out of 17 ages 9–11 years, 11 out of 11 adults).

(0.02). Our findings suggest that some combination of prior and likelihood was present in children as young as 6 years.

Having found that the Bayesian model provides a better fit for the majority of participants across age groups, we estimated the parameters of the Bayesian model. We first show that our parameter-estimation procedure leads to unbiased estimates using simulated data (Figure 5A–5C). We simulated optimal participants with different amounts of noise added to the likelihood ($\Delta_1 = 0.01, 0.05, 0.1$), who either used a prior whose standard deviation was the mean of the experimentally imposed prior standard deviations ($\sigma_{s\ NP}, \sigma_{s\ WP} = 0.065$), used partly differentiated priors ($\sigma_{s\ NP} = 0.048, \sigma_{s\ WP} = 0.083$), or used the experimentally imposed priors ($\sigma_{s\ NP} = 0.03, \sigma_{s\ WP} = 0.1$). The parameters inferred from the model agreed with the simulated parameters (Figure 5A–5C). We then estimated the prior width and noise added to the likelihood ($\sigma_{s\ NP}, \sigma_{s\ WP}$, and Δ_1) from the data of the 37 out of 44 participants whose data were better fit by the Bayesian model. We inferred the prior variance from data sets sampled with replacement 1,000 times from each participant's data. While parameters estimated from the model are corrupted by noise, there is a trend for prior variance parameters to be unaffected by the task (experimental prior width) in young children (Figure 5D) and to shift in the direction of experimentally imposed prior variances during development (Figure 5D–5F). We observe that, in young children ages 6–8 years, sensorimotor estimates are unaffected by prior width both here and in our analysis of estimation slopes. While the priors used by all age groups were suboptimal, this bias appears to decrease with age.

The behavior of a small number of younger participants was better fit by the switching model

(Figure 5H and 5I). We can accurately infer the proportion $p(\text{likelihood})$ of trials where participants use the likelihood from simulated data (Figure 5H). The $p(\text{likelihood})$ is variable across participants who use this strategy (Figure 5I). This simple strategy provides an alternative to Bayesian integration but is not prevalent in our sample.

Discussion

We investigated the development of Bayesian integration during sensorimotor estimation. Statistically efficient estimation required integrating information from a prior distribution with sensory information from a likelihood distribution. We found that participants' estimates reflect the experimentally imposed prior mean, which was expected, since samples of the prior were displayed on screen. Our analysis of estimation slopes showed that all age groups relied on the likelihood according to its uncertainty. While the older child group (9–11 years) and adults distinguished between conditions based on prior width, we did not find evidence for this in the youngest group (6–8 years). We did not find learning related to the prior over the course of the experiment, but we did find slightly less reliance on the likelihood in later blocks. In all age groups, a Bayesian model that integrated likelihood and prior information performed better than simple models that used only one or the other source of information and a model that alternated between them. This finding is consistent with the significant main effect of the likelihood width on estimation slopes. We also found that younger children exhibit larger deviations from optimality in fitted parameters.

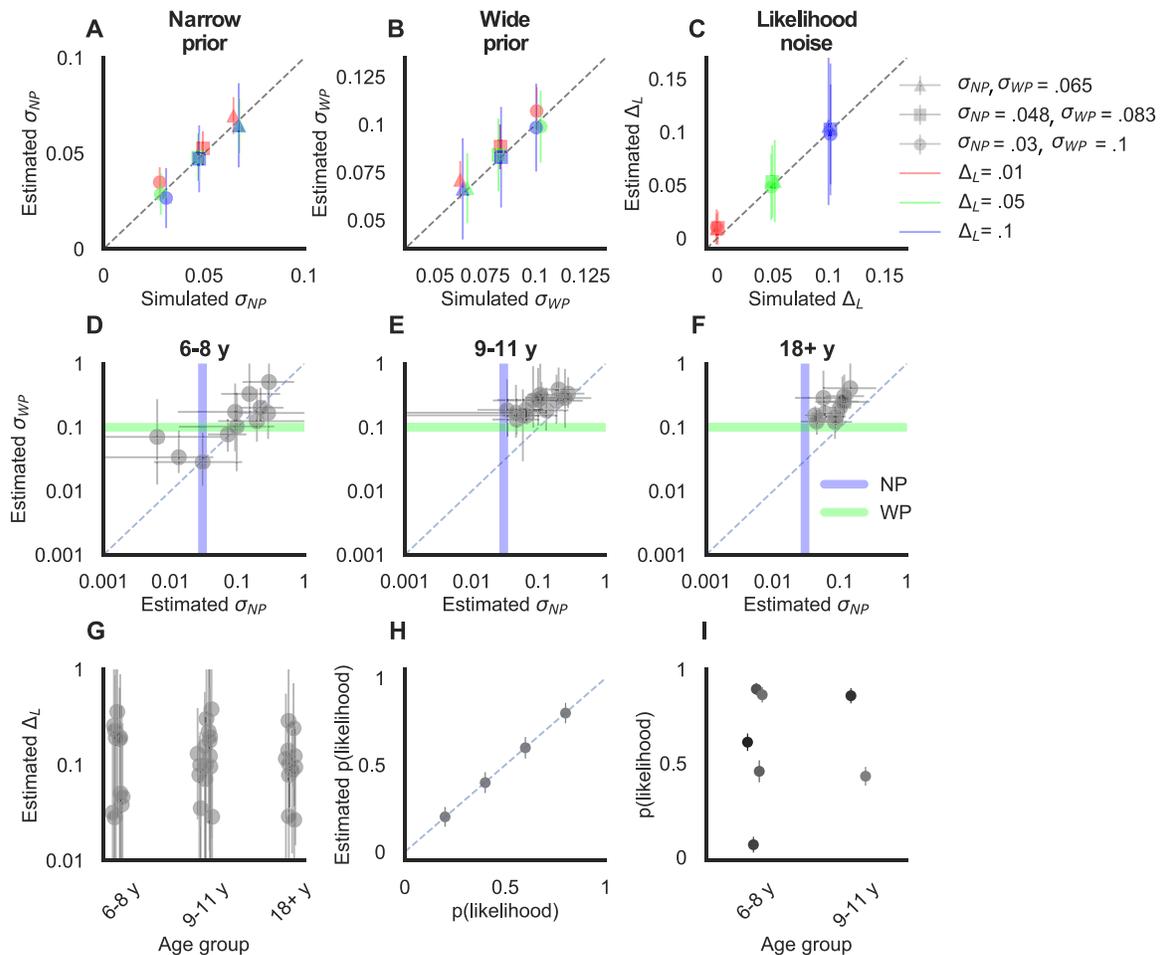


Figure 5. Estimates of model parameters. (A–C) Estimates of model parameters from 1,000 simulated Bayesian participants. (A) Estimates of the Narrow Prior standard deviation ($\sigma_{s, NP}$). (B) Estimates of the Wide Prior standard deviation ($\sigma_{s, WP}$). (C) Estimates of the standard deviation added to the likelihood (Δ_L). Simulated participants used the same prior in both conditions ($\sigma_{s, NP}, \sigma_{s, WP} = 0.065$), used partly differentiated priors ($\sigma_{s, NP} = 0.048, \sigma_{s, WP} = 0.083$), or used the experimentally imposed prior ($\sigma_{s, NP} = 0.03, \sigma_{s, WP} = 0.1$). We also varied the amount of noise added to the likelihood ($\Delta_L = 0.01, 0.05, 0.1$). The median (error bars = 95% confidence interval) is shown in all panels. (D–F) Wide Prior variance as a function of the Narrow Prior variance inferred from the estimation data of participants whose data was best fit by the Bayesian model: (D) ages 6–8 years, (E) ages 9–11 years, and (F) adults. Blue and green lines show the experimentally imposed prior variance in the Narrow and Wide Prior conditions, respectively. (G) Shows the variance added to the likelihood inferred from the estimation data. (H–I) Switch model. (H) Shows the $p(\text{likelihood})$ inferred from the Switch models for 1,000 simulated subjects per $p(\text{likelihood})$ condition. (I) Shows the $p(\text{likelihood})$ for participants whose data was best fit by the Switch model.

We examined the possibility that instead of integrating information from the prior and likelihood, subjects relied on only one source of information on each trial and randomly switched between the two sources across trials. The form of switching that we investigated was simple, in that the probability that subjects relied on the likelihood (the parameter of this model) was fixed across conditions. However, other forms of switching behavior that resemble Bayesian integration are possible and have been explored in previous work (Landy & Kojima, 2001; Laquitaine & Gardner, 2018). For example, participants may weight the prior and likelihood proportionally to their

uncertainties, which would lead to the same set of estimation slopes in our experiment but differences in the variance of estimates. Our experiment did not allow us to accurately distinguish between Bayesian integration and this more complex variant of switching behavior. Switching rather than full integration during sensorimotor estimation remains a possible explanation for sensorimotor estimation in children to be addressed in future work.

Our finding that younger children exhibit larger deviations from optimality than adults suggests that efficient statistical inference is at least partly learned during development, but leaves open questions on the

sources of this suboptimality. Inaccurate likelihoods, priors, cost functions, decisions rules, and approximations to Bayesian inference can all lead to suboptimality (Rahnev & Denison, 2018). It will be interesting to tease apart their contributions in future work.

In a real-world context, priors must be learned from our interactions with the world (Berniker et al., 2010). In future investigations, it may be interesting to examine learning of prior statistics without providing visual cues on screen. Although this was outside the scope of the current work, our findings provide certain predictions on the changes that might occur during development. We found that young children can weight information by its relative uncertainty and that their judgments reflect the mean of the prior but not its variance. This could be further tested by varying the parameters of an unseen prior. It may be the case that children can learn simple statistics of distributions such as the mean but have difficulty with higher order statistics, and do not show signs of representing full distributions as adults do until late in development (Acerbi, Vijayakumar, & Wolpert, 2014; Kording & Wolpert, 2004).

Previous work on Bayesian cue combination has found that children's behavior suggests late integration in childhood, often finding that children's behavior up to the ages of 9–12 years is more consistent with reliance on separate sources of available information and switching between sources of information (Adams, 2016; Gori et al., 2008; Nardini et al., 2008). However, in our study, most of the children in the youngest group tested (6–8 years) integrated prior and likelihood information as shown by the improved fit of a Bayesian model to the data compared with baseline models that use one source of information only or alternate between sources of information.

Why do we observe early signs of integration, which are typically not observed in the case of cue combination? It may be that we learn to integrate information early to perform functions that are crucial to survival such as movement, as in the present work, or basic inferences on the behavior of visual objects (Téglás et al., 2011). It could be that we integrate at an earlier stage of development when the nature of the information to be integrated is the same, like in the current work. Efficient cue combination, on the other hand, requires bringing together signals that are processed separately, whether this be across or within the senses (Gori et al., 2008; Nardini et al., 2010), and it could be that maturation of the circuits which process these cues is necessary before information can be combined (Dekker et al., 2015).

Apparent inefficiencies in the behavior of young children might be in some sense efficient given children's goals and constraints. For example, in the case of cue combination it may be beneficial to avoid fusion during development in order to process cues separately,

so that the senses may calibrate themselves (Gori et al., 2008; Nardini et al., 2008). It will be important for future research to define the factors that lead to integration of information under certain conditions and tasks and not under others (Ernst, 2008). “Bayesian brain” theories must be developed further to account for how Bayesian inference emerges in the developing brain, and in particular must account for developmental trajectories of different behaviors.

If Bayesian computation is at the core of the neural code (Beck et al., 2008; Zemel et al., 1998), behavior should show the signatures of Bayesian inference under all conditions, including during development. Our findings support some level of Bayesian integration at all age groups, but they also show that children do not use available probabilistic information to the same extent as adults. It may be that during development the brain learns to approximate Bayesian principles by means other than explicitly implementing Bayesian computations in neural circuits (Mandt et al., 2017; Weisswange et al., 2011). Our findings fit with ideas suggested by Piaget (1954) on the role of constructivism in child development: that abilities are acquired through experience by building on more basic forms of knowledge. In this sense, learning statistics may be seen as a very basic form of knowledge. While we may be born with a general learning architecture, it seems that statistics should be seen not as core knowledge (Spelke & Kinzler, 2007) but as an acquired skill.

Keywords: development, Bayesian integration, sensorimotor

Acknowledgments

We would like to acknowledge Joshua Glaser for useful comments on this manuscript. This work was funded by NIH grant 5R01NS063399-08 awarded to KPK.

Commercial relationships: none.
Corresponding author: Claire Chambers.
Email: clairenc@seas.upenn.edu.
Address: Department of Bioengineering and Department of Neuroscience, University of Pennsylvania, Philadelphia, PA, USA.

References

- Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the origins of suboptimality in human probabilistic inference. *PLoS Computational Biology*,

- 10(6), e1003661, <https://doi.org/10.1371/journal.pcbi.1003661>.
- Acuna, D. E., Berniker, M., Fernandes, H. L., & Kording, K. P. (2015). Using psychophysics to ask if the brain samples or maximizes. *Journal of Vision*, 15(3):7, 1–16, <https://doi.org/10.1167/15.3.7>. [PubMed] [Article]
- Adams, W. J. (2016). The development of audio-visual integration for temporal judgements. *PLoS Computational Biology*, 12(4):e1004865, <https://doi.org/10.1371/journal.pcbi.1004865>.
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences, USA*, 110(45), 18327–18332, <https://doi.org/10.1073/pnas.1306572110>.
- Beck, J. M., Ma, W. J., Kiani, R., Hanks, T., Churchland, A. K., Roitman, J., ... Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. *Neuron*, 60(6), 1142–1152, <https://doi.org/10.1016/j.neuron.2008.09.021>.
- Berniker, M., Voss, M., & Kording, K. (2010). Learning priors for Bayesian computations in the nervous system. *PLoS One*, 5(9):e12686, <https://doi.org/10.1371/journal.pone.0012686>.
- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389–414, <https://doi.org/10.1037/a0026450>.
- Chambers, C., Sokhey, T., Gaebler-Spira, D., & Kording, K. P. (2017). The integration of probabilistic information during sensorimotor estimation is unimpaired in children with cerebral palsy. *PLoS One*, 12(11), e0188741.
- Dekker, T. M., Ban, H., Van Der Velde, B., Sereno, M. I., Welchman, A. E., & Nardini, M. (2015). Late development of cue integration is linked to sensory fusion in cortex. *Current Biology*, 25(21), 2856–2861, <https://doi.org/10.1016/j.cub.2015.09.043>.
- Dekker, T. M., & Nardini, M. (2016). Risky visuo-motor choices during rapid reaching in childhood. *Developmental Science*, 19(3), 427–439, <https://doi.org/10.1111/desc.12322>.
- Ernst, M. O. (2008). Multisensory integration: A late bloomer. *Current Biology*, 18(12), 519–521, <https://doi.org/10.1016/j.cub.2008.05.003>.
- Ernst, M. O., & Banks, M. S. (2002, January 24). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429–433, <https://doi.org/10.1038/415429a>.
- Gopnik, A., & Wellman, H. M. (2013). Reconstructing constructivism: Causal models, Bayesian learning mechanisms and the theory theory. *Psychological Bulletin*, 138(6), 1085–1108, <https://doi.org/10.1037/a0028044>.
- Gori, M., Del Viva, M., Sandini, G., & Burr, D. C. (2008). Young children do not integrate visual and haptic form information. *Current Biology*, 18(9), 694–698, <https://doi.org/10.1016/j.cub.2008.04.036>.
- Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, 4(12):1, 967–992, <https://doi.org/10.1167/4.12.1>. [PubMed] [Article]
- Jones, M., & Love, B. C. (2011). Bayesian fundamentalism or enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. *Behavioral and Brain Sciences*, 34, 169–231, <https://doi.org/10.1017/S0140525X10003134>.
- Knill, D. C., & Richards, W. (1996). *Perception as Bayesian inference*. New York, NY: Cambridge University Press.
- Kording, K. P., & Wolpert, D. M. (2004, January 15). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247, <https://doi.org/10.1038/nature02169>.
- Kushnir, T., & Gopnik, A. (2007). Conditional probability versus spatial contiguity in causal learning: Preschoolers use new contingency evidence to overcome prior spatial assumptions. *Developmental Psychology*, 43(1), 186–196, <https://doi.org/10.1037/0012-1649.43.1.186>.
- Landy, M. S., & Kojima, H. (2001). Ideal cue combination for localizing texture-defined edges. *Journal of the Optical Society of America*, 18(9), 2307–2320, <https://doi.org/10.1364/JOSAA.18.002307>.
- Laquitaine, S., & Gardner, J. L. (2018). A switching observer for human perceptual estimation. *Neuron*, 97(2), 462–474. e6, <https://doi.org/10.1016/j.neuron.2017.12.011>.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, 9(11), 1432–1438, <https://doi.org/10.1038/nn1790>.
- Ma, W. J., Beck, J. M., & Pouget, A. (2008). Spiking networks for Bayesian inference and choice. *Current Opinion in Neurobiology*, 18(2), 217–222, <https://doi.org/10.1016/j.conb.2008.07.004>.
- Mamassian, P., & Goutcher, R. (2001). Prior knowledge on the illumination position. *Cognition*, 81(1), B1–B9. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11525484>.
- Mandt, S., Hoffman, M. D., & Blei, D. M. (2017).

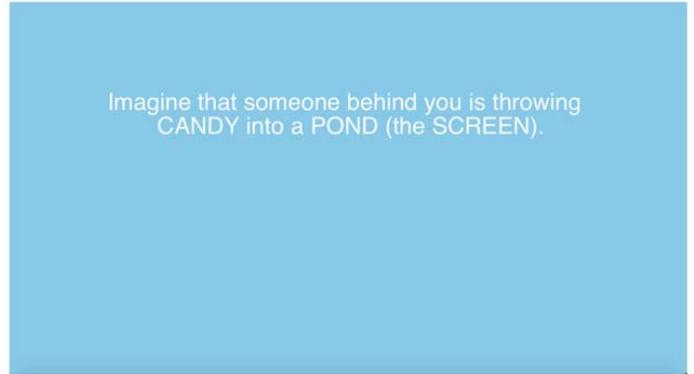
- Stochastic gradient descent as approximate Bayesian inference. *ArXiv*, <https://doi.org/1704.04289v1>.
- Nardini, M., Bedford, R., & Mareschal, D. (2010). Fusion of visual cues is not mandatory in children. *Proceedings of the National Academy of Sciences, USA*, *107*(39), 17041–17046, <https://doi.org/10.1073/pnas.1001699107>.
- Nardini, M., Begus, K., & Mareschal, D. (2013). Multisensory uncertainty reduction for hand localization in children and adults. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(3), 773–787, <https://doi.org/10.1037/a0030719>.
- Nardini, M., Jones, P., Bedford, R., & Braddick, O. (2008). Development of cue integration in human navigation. *Current Biology*, *18*(9), 689–693, <https://doi.org/10.1016/j.cub.2008.04.021>.
- Ouvrier, R., Goldsmith, R. F., Ouvrier, S., & Williams, I. C. (1993). The value of the Mini-Mental State Examination in childhood: A preliminary study. *Journal of Child Neurology*, *8*(2), 145–148, <https://doi.org/10.1177/088307389300800206>.
- Piaget, J. (1954). *The construction of reality in the child*. New York, NY: Basic Books.
- Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, February, 1–107, 1–18, <https://doi.org/10.1017/S0140525X18000936>.
- Reynolds, C. R., & Kamphaus, R. (2004). *Behavior Assessment System for Children, (BASC-2) Handbook*. AGS Publishing, 4201, 55014-1796.
- 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, <https://doi.org/10.1016/j.cogsci.2003.11.001>.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, *10*(1), 89–96, <https://doi.org/10.1111/j.1467-7687.2007.00569.x>.
- Téglás, E., Tenenbaum, J. B., & Bonatti, L. L. (2011, May 27). Pure reasoning in 12-month-old infants as probabilistic inference. *Science*, *332*(6033), 1054–1059, <https://doi.org/10.1126/science.1196404>.
- Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity and Bayesian inference. *Behavioral and Brain Sciences*, *24*(4), 629–640, <https://doi.org/10.1017/S0140525X01000061>.
- Verstynen, T., & Sabes, P. N. (2012). How each movement changes the next: An experimental and theoretical study of fast adaptive priors in reaching. *The Journal of Neuroscience*, *31*(27), 10050–10059, <https://doi.org/10.1523/JNEUROSCI.6525-10.2011>.
- Vilares, I., Howard, J. D., Fernandes, H. L., Gottfried, J. A., & Kording, K. P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. *Current Biology*, *22*(18), 1641–1648, <https://doi.org/10.1016/j.cub.2012.07.010>.
- Weisswange, T. H., Rothkopf, C. A., Rodemann, T., & Triesch, J. (2011). Bayesian cue integration as a developmental outcome of reward mediated learning. *PLoS One*, *6*(7):e21575, <https://doi.org/10.1371/journal.pone.0021575>.
- Wilson, B. N., Crawford, S. G., Green, D., Roberts, G., Aylott, A., & Kaplan, B. J. (2009). Psychometric properties of the revised Developmental Coordination Disorder Questionnaire. *Physical & Occupational Therapy in Pediatrics*, *29*(2), 182–202.
- Wolpert, D. M., & Landy, M. S. (2012). Motor control is decision-making. *Current Opinion in Neurobiology*, *22*(6), 996–1003, <https://doi.org/10.1016/j.conb.2012.05.003>.
- Wolraich, M. L., Lambert, W., Doffing, M. A., Bickman, L., Simmons, T., & Worley, K. (2003). Psychometric properties of the Vanderbilt ADHD diagnostic parent rating scale in a referred population. *Journal of Pediatric Psychology*, *28*(8), 559–568.
- Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, *10*(2), 403–430, <https://doi.org/10.1162/089976698300017818>.

Appendix: Instructions during experiment, analysis of prior mean

Screen 1



Screen 2



Screen 3



Screen 4



Screen 5



Screen 6

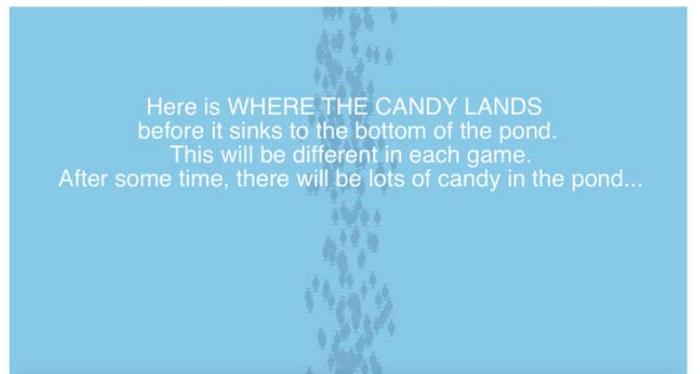


Figure A1. Instructions Part 1. Screens 1 and 2. We presented participants with the instructions that someone behind them was throwing candy into a pond, represented by the screen. Screens 3 to 6. We showed participants 200 samples from the prior distribution (“Where the candy lands”). The first 10 samples (two shown here) were shown “falling” into the pond one by one.

Screen 7



Screen 8



Screen 9



Screen 10



Screen 11



Screen 12



Figure A2. Instructions Part 2. Screens 7 to 10. We showed participants an example trial, where they were given noisy feedback on target location in the form of $n=4$ samples from the likelihood (Splash). We showed participants a vertical bar and asked them to use it to catch the candy by moving the bar from left to right. After they provided their estimate, we showed them the candy's true location. Screens 11 and 12. We gave participants information on “bonuses” and the duration of the experiment.

Comparison	Age group	<i>t</i> (<i>df</i>)	<i>p</i>
Narrow versus medium likelihood	6–8 years	1.64 (30)	0.1123
	9–11 years	0.60 (32)	0.5537
	18+ years	0.96 (86)	0.3395
Medium versus wide likelihood	6–8 years	1.23 (30)	0.2296
	9–11 years	1.07 (32)	0.2921
	18+ years	1.50 (86)	0.1365
Narrow versus wide likelihood	6–8 years	1.01 (30)	0.3197
	9–11 years	1.77 (32)	0.0862
	18+ years	1.56 (86)	0.8767

Table A1. Paired *t* tests comparing prior mean across likelihood width and age conditions (corrected $\alpha = 0.0056$).

Comparison	<i>t</i> (<i>df</i>)	<i>p</i>
Narrow versus medium likelihood	0.96 (86)	0.3395
Medium versus wide likelihood	1.50 (86)	0.1365
Narrow versus wide likelihood	0.16 (86)	0.8767

Table A2. Paired *t* tests comparing prior mean across likelihood width (corrected $\alpha = 0.0167$).