Policies or Knowledge: Priors differ between perceptual and sensorimotor tasks

Claire Chambers [1,2,3]

Hugo Fernandes [1,2,3]

Konrad Kording [1,2,3]

1. Sensory Motor Performance Program, Rehabilitation Institute of Chicago, Chicago, IL, USA
2. Biomedical Engineering, Robert R. McCormick School of Engineering and Applied Sciences, Northwestern University, Chicago, IL, USA
3. Department of Physical Medicine and Rehabilitation, Northwestern University, Chicago, IL, USA

If the brain abstractly represents probability distributions as knowledge, then the modality of a decision, e.g. movement vs perception, should not matter. If on the other hand, learned representations are policies, they may be specific to the task where learning takes place. Here, we test this by asking if a learned spatial prior generalizes from a sensorimotor estimation task to a two-alternative-forced choice (2-Afc) perceptual comparison task. A model and simulation-based analysis revealed that while subjects learn the experimentally-imposed prior distribution in the sensorimotor estimation task, measured priors are consistently broader than expected in the 2-Afc task. That the prior does not fully generalize suggests that sensorimotor priors strongly resemble policies. In disagreement with standard Bayesian thought, the modality of the decision has a strong influence on the implied prior distributions.

The acquisition of knowledge is thought to be at the core of the brain's function (Battaglia, Hamrick, & Tenenbaum, 2013; Tenenbaum, Griffiths, & Kemp, 2006; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). A behavioral signature of knowledge-use is strong generalization across situations. For instance, when a child learns a new word they can use it in many new situations, not just the sentence where the word was learned (Perfors, Tenenbaum, & Regier, 2011; Xu & Tenenbaum, 2007). However, the framing of learned representations as generalizable knowledge may not apply to all of the brain's functions equally. For example, generalization from movements of one arm to those of the other is not always complete (Criscimagna-hemminger et al., 2003; Shadmehr, 2004). Indeed, the reinforcement learning literature (Sutton & Barto, 1998) defines an alternative way of learning. Within this framework, learning is framed as policy-acquisition, i.e. mappings from states to actions (Daw & Doya, 2006; Haith & Krakauer, 2013). This definition implies that learning of policies is specific to the action for which it was learned and thus suggests limited generalization across tasks. We want to know if humans are policy animals, knowledge carriers, or something in between.

In sensorimotor estimation tasks, humans weigh prior knowledge with sensory information in a near-optimal way (Berniker, Voss, & Kording, 2010; Kording & Wolpert, 2004; Tassinari, Hudson, & Landy, 2006; Vilares, Howard, Fernandes, Gottfried, & Kording, 2012) and generalize learned prior statistics to new conditions (Fernandes, Stevenson, Vilares, & Kording, 2014). Thus, there is evidence for learning of sensorimotor priors. However, little is known about whether sensorimotor learning generalizes when the read-out modality of the decision changes. Therefore, we do not know if sensorimotor motor priors should be described as knowledge or policies. This is important because it has consequences for how neural representations should be conceptualized.

Here, we investigate if priors are always the same across modalities by examining whether priors generalize across two simple tasks. The tasks were computationally equivalent, in that both involved a judgment concerning spatial location and the stimuli were identical. Subjects learned a spatial prior in a sensorimotor estimation task, and we asked if they transferred the learned prior to a two-alternative-forced-choice (2-Afc) task, where subjects made a binary decision about object location. We inferred the variance of the learned prior and found that the learned sensorimotor prior does not generalize fully to the 2-Afc task. The prior variance measured from 2-Afc decisions was higher than the near-optimal variance measured from sensorimotor estimations. This shows that a learned prior does not generalize fully across sensorimotor and decisional modalities and suggests that sensorimotor priors are represented as policies.

Results

We asked if a learned prior distribution generalizes across tasks and thus consists of knowledge. To do so, we first had subjects learn a prior in a sensorimotor estimation task where subjects gave a continuous estimate of target location under uncertainty. We then quantified use of the prior in a 2-Afc task where instead subjects compared the locations of two hidden targets (Fig. 1). We used data from previous work (Acuna, Berniker, Fernandes, & Kording, 2015). Using model-based data analysis and simulations, we compared the performance of two models: a model with a common prior for the two tasks and a model with task-specific priors. We then inferred the parameters from the best-fitting model to quantify the prior used in each task.

In our tasks, subjects judged the location of visual targets on screen. Targets were samples from a Gaussian prior distribution, $N(\mu, \sigma_s^2)$, which were hidden from view. Instead, subjects were shown an uncertain visual cue in the form of *n* samples (*n*=4) from a Gaussian likelihood distribution, distributed around target location, $N(s, \sigma_l^2)$. When judging target location, the Bayesian optimal strategy is to combine the likelihood and the learned prior according to their relative precision. Therefore, to examine subjects' use of probabilistic information, we manipulated the variances of the prior and likelihood, σ_s^2 and σ_l^2 and quantified how much subjects rely on the likelihood or prior to reach a decision (see Methods, Fig. 2). Our paradigm allowed us to examine integration of probabilistic information and to infer subjects' learned priors in the estimation and 2-Afc tasks.

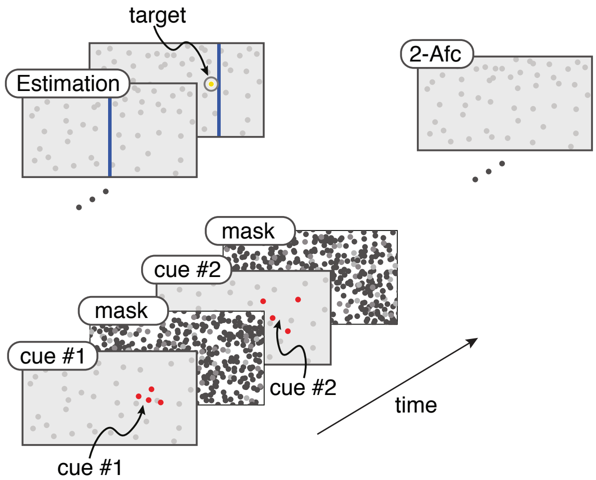

Figure 1. Experimental protocol. Subjects were shown two splashes in succession, created by hidden targets (described as coins in the task) falling into a pond interleaved with white noise masks. Subjects were then presented with one of two possible tasks. In the estimation task, subjects were prompted to place a net where the second hidden target fell. In the 2-Afc task, subjects reported which hidden target landed farther to the right.

To examine use of probabilistic information in making estimations, we first examined influence of the experimentally-imposed prior variance (σ_s^2) and likelihood variance (σ_l^2) on subjects' reliance on the likelihood or prior (*Reliance on the likelihood*, see Fig. 2 and Methods for details). We found that both σ_s^2 and σ_l^2 influence the *Reliance on the likelihood* (main effect of σ_s^2, p<.001, Scheirer-Ray-Hare test, F(1, 28) = 67.33; main effect of σ_l^2, p<.001, Scheirer-Ray-Hare test, F(1, 28) = 90.18). Therefore, subjects use the prior and likelihood variances to judge target location. Second, we compared the *Reliance on the likelihood* to optimal values and found agreement between *Reliance on the likelihood* and optimal predictions (Fig. 2a,b). The *Reliance on the likelihood* averaged across conditions is unbiased relative to the optimal (p = .67, paired, 2-sided Wilcoxon signed-rank test, Z=.42), suggesting that behavior does not systematically deviate from optimal predictions. Thus, in the estimation task, subjects use probabilistic information to perform the task and behavior is well accounted for by Bayesian optimal predictions. Therefore, it makes sense to describe subjects' sensorimotor estimations as Bayesian and to quantify the prior used during the task.

We then examined usage of probabilistic information in the 2-Afc task, with a measure of reliance on prior or likelihood in decision data, which we termed the *PSE slope* (see Fig. 2c,d and Methods for details). There was a significant effect of σ_s^2 on *PSE slope* (p = .049, paired, 2-sided Wilcoxon signed-rank test, Z = 1.97). Thus, subjects have a greater reliance on the prior in the Narrow-Prior condition, which is consistent with Bayesian computation. However, comparison of the *PSE slope* with optimal values shows that the *PSE slope* is positively biased relative to the optimal ($p = .01$, paired, 2-sided Wilcoxon signed-rank test, $Z = 2.52$). Thus, there is evidence for use of probabilistic information in the 2-Afc task, but the positive bias in *PSE slope* suggests that subjects' priors in this task are broader than the theoretical prior.

Behavior in the two tasks is in accordance with the use of probabilistic information. This was shown by an influence of the uncertainty of the prior and likelihood on judgments in both tasks. However, the priors used in the estimation and 2-Afc tasks appear to be different. The *Reliance on the likelihood* computed from estimation data is in agreement with the statistical optimal, unlike the *PSE slope* fit to the 2-Afc data. This result is in violation of standard Bayesian thought where the prior representation is considered as knowledge and hence, domain general and available for use across tasks.

We then asked if the data supports a model with task-dependent prior representations. For a full description of the model-based analysis, see Methods. We compared the performance of a model with a common prior for both tasks, 1-Prior model, with the performance of a model with task-specific priors, 2-Prior model. The 1-Prior model contained a variance parameter for the Narrow Prior (NP) condition and the Wide Prior (WP) condition ($\sigma^2_{s\,NP}$, $\sigma^2_{s\,WP}$) and a lapse parameter to account for attentional 2-Afc task errors (λ). The 2-Prior model contained task-specific variance parameters for each prior condition ($\sigma^2_{s\,NP\,Est}$, $\sigma^2_{s\,WP\,Est}$, $\sigma^2_{s\,NP\,Afc}$, $\sigma^2_{s\,WP\,Afc}$) and a lapse parameter (λ). To compare model performance, we used the log of the likelihood ratio between the 2-Prior and 1-Prior model. A positive log-likelihood ratio indicates better performance of the 2-Prior model relative to the 1-Prior model. The log-likelihood ratio was significantly above 0 ($p = .01$, paired, 2-sided Wilcoxon signed-rank test, $Z = 2.52$, Fig. 3a). Therefore, subjects use different priors in the different tasks.

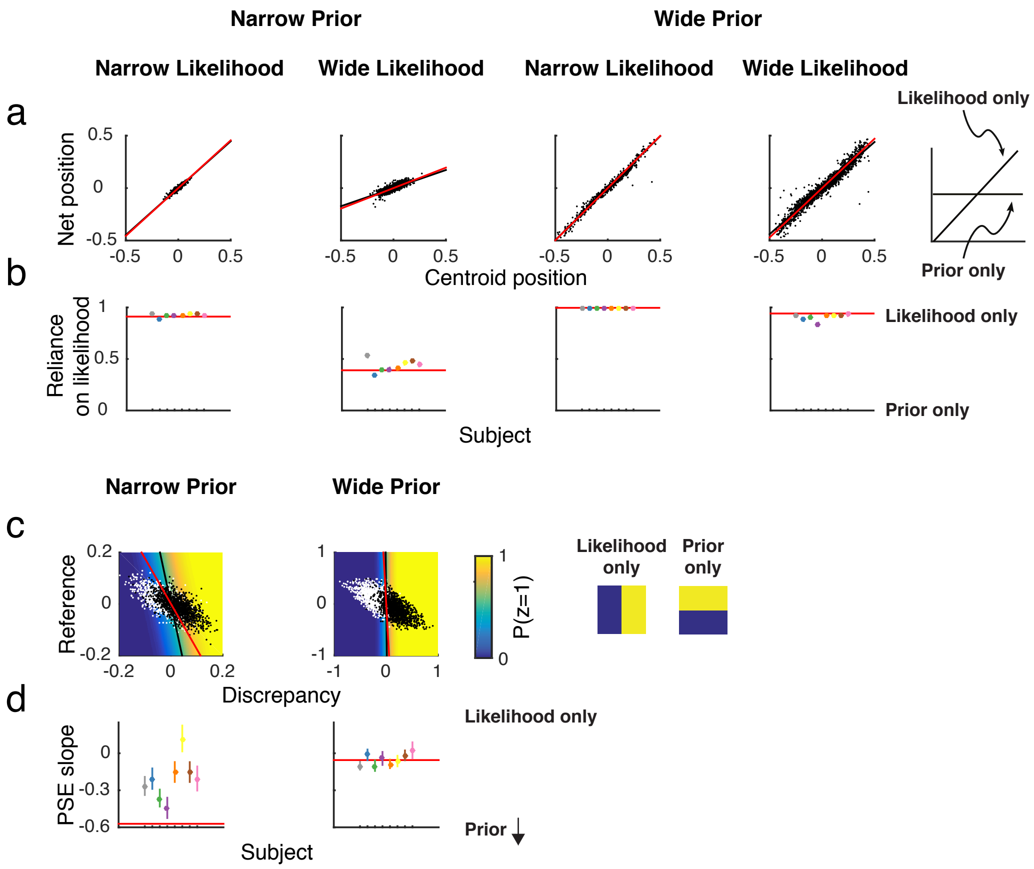

Figure 2. Estimation and 2-Afc data. (**a**) Estimation data overlaid with linear fit for a representative subject (subject #2 shown in blue in (**b**) and (**d**)). The net position as a function of the centroid of the likelihood is shown for each trial (black points). The estimation-slope fit (black line) and the optimal (red line) are displayed. Each panel displays estimation data for one condition, with overlaid fitted (black line) and optimal (red line) functions. (**b**) The median bootstrapped *Reliance on the likelihood* is shown for all subjects (error bars = 95% CI, not always visible since CI are narrow). The optimal *Reliance on the likelihood* values are shown (red line). A *Reliance on the likelihood* of 1 indicates complete reliance on the likelihood and a *Reliance on the likelihood* of 0 indicates complete reliance on the prior (**c**) 2-Afc data for the representative subject in (**a**) with one panel per condition. Raw binary decision data (white points, z=0, probe target to the left; black points, z=1, probe target to the right) and fitted psychometric functions describing the probability of the probe target being reported to the right (*P(z=1)*). The fitted PSE (black line) and optimal PSE (red line) are shown. (**d**) The median bootstrapped *PSE slope* is shown for all subjects (error bars = 95% CI), with optimal values (red line). As the prior width decreases so does the *PSE slope* (shown by arrow).

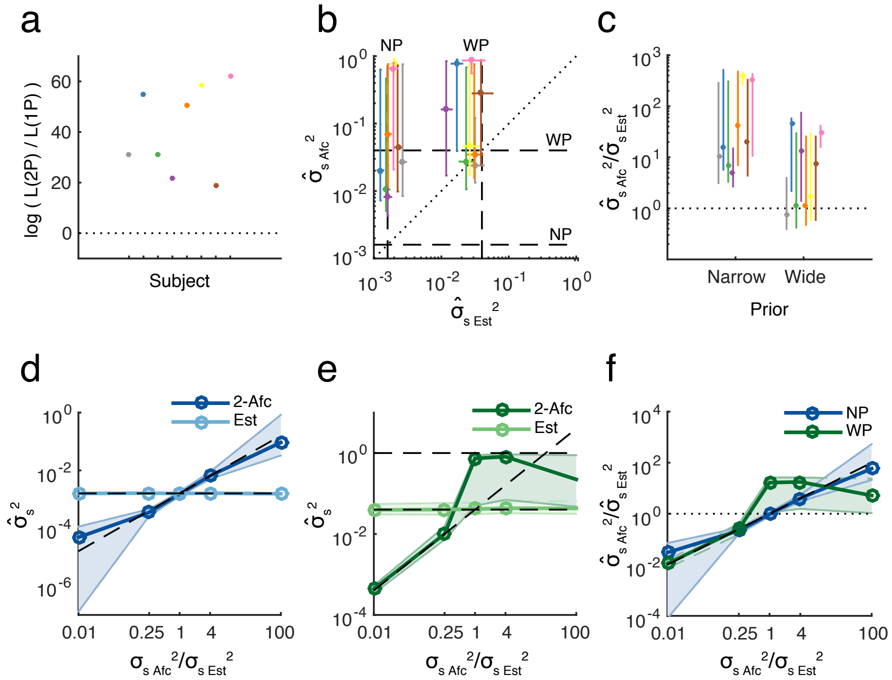

Figure 3. Model comparison and parameter estimation (**a**) The log-likelihood ratio between the 2-Prior model and the 1-Prior model computed from each subject's data is shown (colored circles). Positive values indicate that the 2-Prior model outperforms the 1-Prior model. (**b**) For each subject, prior variance inferred from the 2-Afc task data ($\hat{\sigma}^2_{s\,Afc}$) is shown as a function of the learned prior variance inferred from the estimation task data ($\hat{\sigma}^2_{s\,Est}$). The median bootstrap is shown (error bar=95% CI). Broken lines show the theoretical prior variances and the dotted line shows the diagonal, for which prior variances in the tasks are equal. (**c**) The median ratio between priors inferred from the 2-Afc and estimation data ($\hat{\sigma}^2_{s\,Afc}/\hat{\sigma}^2_{s\,Est}$) is shown for Narrow and Wide-Prior conditions. $\hat{\sigma}^2_{s\,Afc}/\hat{\sigma}^2_{s\,Est}$=1 is shown by the dotted line. (**d**-**e**) Prior-variance parameters estimated from the data of simulated Bayesian observers in the Narrow-Prior (**d**) and Wide-Prior (**e**) conditions. Simulated subjects use the theoretical prior variance in the estimation task and either the same prior variance in the 2-Afc task (1-Prior model, simulated $\sigma^2_{s\,Afc}/\sigma^2_{s\,Est}$=1) or a different prior variance in the 2-Afc task (2-Prior model, simulated $\sigma^2_{s\,Afc}/\sigma^2_{s\,Est}$=.01, .25, 4, 100. The median inferred prior is shown (shaded area=95% CI). Broken lines show the theoretical prior variances and the upper bound at 1. (**f**) The median $\hat{\sigma}^2_{s\,Afc}/\hat{\sigma}^2_{s\,Est}$ inferred from simulated data as a function of the simulated $\sigma^2_{s\,Afc}/\sigma^2_{s\,Est}$ for Narrow-Prior condition in blue and the Wide-Prior condition in green (shaded area=95% CI).

Model comparison results showed that priors were different. We wanted to know how they were different. To see how priors differed across tasks, we inferred the parameters of the best-fitting 2-Prior model, $\theta_{2\,prior} = \{\sigma^2_{s\,NP\,Est},\ \sigma^2_{s\,WP\,Est},\ \sigma^2_{s\,NP\,Afc},\ \sigma^2_{s\,WP\,Afc},\ \lambda\}$. Since we were interested in the relationship between the learned prior variances in the 2-Afc and estimation tasks, we then computed the ratio between the prior variances inferred from the model ($\hat{\sigma}^2_{s\,Afc}/\hat{\sigma}^2_{s\,Est}$). $\hat{\sigma}^2_{s\,Afc}/\hat{\sigma}^2_{s\,Est}$ is consistently above 1 for individual subjects in the Narrow-Prior condition (95% CI do not overlap with 1, Fig. 3c) This shows that the 2-Afc-task prior is wider than the estimation-task prior. This was not consistent for individual subjects in the Wide-

Prior condition (95% CI overlap with 1 in 5/8 subjects, Fig. 3c). Thus, we find that priors are broader in the Narrow-Prior condition, but not in the Wide-Prior condition. Regarding our hypothesis on the generalization of the learned prior from the estimation task to the 2-Afc task, this shows that the prior does not always generalize fully. Therefore, our model-based analysis supports policy representation rather than knowledge representation.

The difference between priors in the Narrow-Prior condition could have been caused by a bias in the parameter-estimation technique. We ensured that this was not the case by simulating a Bayesian observer whose prior variance was the theoretical value for both tasks and examining whether the simulation falsely demonstrated a difference in priors. In a subject simulated from the 1-Prior model, we found no difference between $\hat{\sigma}_{s\,Est}^2$ and $\hat{\sigma}_{s\,Afc}^2$ in Narrow-Prior and Wide-Prior conditions (Fig. 3d-f). Therefore, the effect found in the Narrow-Prior condition cannot be attributed to bias in the data analysis.

We failed to find a difference in inferred priors across tasks in the Wide-Prior condition. This could have been due to a genuine lack of difference or alternatively factors such as data-set size and the experimental parameters chosen. To check for this, we simulated subjects who used different prior variances in the two tasks, using the same amount of data as in the behavioral data set. In the simulated data, we replicated the lack of difference across tasks in the Wide-Prior condition when the simulated $\sigma_{s\,Afc}^2$ is wider than the theoretical value (Fig. 3e,f). Therefore, the lack of difference in the Wide-Prior condition could have been due to insufficient data or the experimental parameters that we used.

Discussion

We examined if a prior distribution learned during a sensorimotor estimation task generalized to a computationally-equivalent 2-Afc decision task. A model-based analysis showed that there was a difference in priors across tasks. The inferred prior was close to optimal in the sensorimotor estimation task where subjects learned the prior. The finding of a wider prior in the 2-Afc task shows that the prior did not generalize fully from the situation where subjects provided a continuous estimate of location to different task where subjects compared two object locations. This shows that sensorimotor priors are not knowledge, in the sense that they do not generalize fully across modalities.

Our results may seem trivial if people cannot properly integrate priors and likelihoods in the 2-Afc task. However, the 2-Afc task has been used to investigate probabilistic integration in cue combination, with findings of near optimal integration of information (Ernst & Banks, 2002; Hillis, Watt, Landy, & Banks, 2004). With regards to our experiments, this implies that it should be possible to incorporate the prior into 2-Afc judgments after learning occurs within this task. Therefore, it is not the 2-Afc task itself that leads to non-optimal use of priors. Rather, the learned prior does not generalize from one task to another.

One problem that emerged from our results was that our estimates of the prior variance from the 2-Afc decision data were noisy. This prevented us from concluding whether there was a difference in priors across tasks in the Wide-Prior condition. This could be due to insufficient data or the parameters used, as our simulations suggest.

Another caveat is that we assume that the brain uses maximum a-posteriori (MAP) to compute decisions. MAP is widely-used in the decision-making literature and is a plausible choice of mechanism since it maximizes reward for simple cases (Maloney, 2002; Mamassian, Landy, & Maloney, 2001). Other decision-making mechanisms include sampling from probability distributions and have been explored in previous work (Acuna et al., 2015; Vul, Goodman, Griffiths, & Tenenbaum, 2014). While the choice of MAP may be reasonable in the case of unimodal Gaussian posterior distributions as in the current study, MAP is less adapted to cases of multimodal or broadly-distributed posteriors. Further work is needed on which decision rules that the brain uses.

Why are the priors different? The tasks may engage distinct neural systems, with the estimation task having a stronger sensorimotor component ('Where is the object in relation to me?'), whereas the 2-Afc task is a perceptual task and concerns relationships between objects in the outside world ('Where is one object in relation to another?'). Therefore, partly independent neural representations may lead to incomplete generalization across tasks (Aglioti, DeSouza, & Goodale, 1995; Knill, 2005). In this view, partial generalization comes from partly distinct neural systems.

One implication of our finding is that priors cannot be assumed to generalize even when the difference between learning and testing conditions/tasks is subtle. For example, previous work investigating decision-making mechanisms quantifies the prior in an estimation task and measures the influence of the subjective prior variance in a 2-Afc task (Acuna et al., 2015). The findings of this work therefore rest on the assumption that the prior is the same across tasks and the conclusions of this paper and others with the same assumption should be revisited.

Importantly, our finding is inconsistent with the view that the brain acquires fully generalizable knowledge, in the form of priors that can automatically be incorporated into behavior regardless of the task. While high-level conceptual representations may fit the definition of knowledge (Battaglia et al., 2013; Perfors et al., 2011; Tenenbaum et al., 2011), our findings show that learning in a sensorimotor task has a strong policy component, with a prior being partly confined to the task where it was learned. In naturalistic situations, the use of policies may be functionally beneficial, allowing for learning to be flexible and optimized for the task at hand.

These results are compatible with a learning framework, rather than a high-level Bayesian view of the brain's computations, where one set of priors (knowledge) is used for different output behaviors. Multi-layer neural networks provide a flexible way of modeling diverse kinds of behavior based on function optimization (Lecun, Bengio, & Hinton, 2015; Marblestone, Wayne, & Kording, 2016). Within a broader network, sub-networks that implement specialized learning could produce patterns of generalization or non-generalization across conditions and tasks. Importantly, a system that learns by gradient descent will approximate Bayesian behavior without explicitly implementing Bayesian computations (Rao, 2004), simply because it is the optimal strategy for estimations under uncertainty. Our finding thus casts doubt on the view that Bayesian computation is at the core of the neural code (Fiser, Berkes, Orbán, & Lengyel, 2010; Ma, Beck, Latham, & Pouget, 2006; Pitkow & Angelaki, 2017; Zemel, Dayan, & Pouget, 1998).

Methods

Experimental details

The results presented here use data from previous work (Acuna et al., 2015), augmented with newly collected data on the same paradigm. A complete description of the methods is given in previous work and will be described here.

We required computationally-equivalent tasks that allowed us to infer the variance of priors used by subjects. We used a "coin-catching" task (Acuna et al., 2015; Berniker et al., 2010; Vilares et al., 2012), where on each trial, subjects guessed the location of a hidden target ("coin") on the screen based on an uncertain visual cue ("splash") and a prior learned through feedback on target location. Varying the prior and likelihood variance allowed us to assess whether subjects weighed prior and likelihood information according to their relative uncertainties during sensorimotor estimation and decision making.

Before starting the experiment, subjects were presented with the instructions that someone was throwing coins into a pond, represented by the screen; and that their aim was to guess where the coin target landed (Fig. 1). On each trial, they were presented with "splash" stimuli and were told that it was caused by a hidden coin target. On estimation trials, subjects provided an estimate of the second target's location on the horizontal axis by placing a vertical bar where they thought that the target landed. On 2-Afc trials, subjects compared the locations of the inferred target locations and indicated which target was further to the right. Subjects were paid based on their performance on the estimation task, as quantified using the distance between their estimates and the true target location.

Eight subjects performed the experiment, including the seven subjects from an existing data set (Acuna et al., 2015) and one additional subject to increase the power of group statistics (statistical results were the same with and without this subject). This subject corresponds to subject 4 in purple on each plot where individual results are shown. The experiment lasted 10,000 trials over 5 days. On each day, they were seated in front of a computer monitor (52 cm wide, 32.5 cm high) in a quiet room. Stimuli were generated by sampling visual targets from a Gaussian prior distribution defined over spatial location, with a mean at the center of the screen and standard deviation of .04 or .2 in units of screen width. The target was hidden from view. Instead, they were presented with a visual cue with experimentally-controlled uncertainty (splash stimulus). The splash consisted of four dots sampled from a Gaussian likelihood distribution centered on the target location. The likelihood distribution could have a standard deviation of .025 or .1 in units of the screen. In all trials, two consecutive splashes were displayed for .025 s, each followed by a visual mask for .5 s. The standard deviation of the likelihood was either the same across presentations within a trial (both at .025 or both at .1) or varied within trial (.025 and .1) and presented in randomized order. We refer to the broader likelihood as the reference and the narrower likelihood as the probe.

On each trial, subjects performed one of two tasks, as defined by the question at the end of the trial. On estimation trials, subjects were asked "Where was the coin located?" and they indicated where they thought the second coin target was using a vertical bar ("net"), which was 2% screen width and extended from the top to the bottom of the screen. On 2-Afc trials, subjects were asked "Which coin was further to the right?" and using a key-press they indicated if they thought the first or second coin target was further to the right. At the end of estimation trials only, feedback was provided on the exact location of the target, but not on 2-Afc trials, allowing us to ask if the prior learned in the estimation task generalizes to the 2-Afc task.

There were four conditions in the estimation task: Narrow Prior – Narrow Likelihood, Narrow Prior – Wide Likelihood, Wide Prior – Narrow Likelihood, Wide Prior – Wide Likelihood. In the 2-Afc trials, conditions were defined by the width of the prior, Narrow Prior and Wide Prior, and whether likelihoods were equal within trial, Equal Likelihoods (both narrow or both wide) or Unequal Likelihoods (one narrow and one wide). We only used Unequal Likelihood trials in the present study. Therefore in our analysis, there were two conditions for the 2-Afc trials: Narrow Prior and Wide Prior.

On each day of the experiment subjects performed two 1,000-trial blocks. The prior over target location switched from block to block (e.g., from wide to narrow on one day, from narrow to wide on the subsequent day, and so on). Each block contained 500 estimation trials and 500 2-Afc trials in a random order all generated from the same prior. In order to aid with learning the prior, estimation trials made up the first half of each block (375 estimation trials and 125 2-Afc trials), and 2-Afc trials made up the second half of each block (125 estimation trials and 375 2-Afc trials).

Model-based Data Analysis and simulations

We asked whether the use of prior information differed between psychophysical tasks. To answer this question, we compared the performance of two models in accounting for the behavioral data: a model in which the variance of the prior was fixed across tasks (1-Prior model), and a second model which contained task-specific prior variance parameters (2-Prior model). To examine how priors related to the theoretical values, we then estimated the parameters for the best model using bootstrapped parameter estimation. To ensure that the data analysis produced unbiased results, we performed the same analysis on data simulated from a Bayesian observer.

Simulations

We wanted to ensure that our analysis led to unbiased model comparison results and unbiased estimates of the prior variance, which we denote σ_s^2. We therefore simulated data from the 1-Prior model with subjects who used the theoretical prior in both tasks (Narrow Prior: $\sigma_s^2{}_{Est}$=.0016, $\sigma_s^2{}_{Afc}$=.0016; Wide Prior: $\sigma_s^2{}_{Est}$=.04, $\sigma_s^2{}_{Afc}$=.04), expecting to find a better fit for the 1-Prior model. We also simulated data from the 2-Prior model where subjects used the theoretical prior in the estimation task (Narrow Prior: $\sigma_s^2{}_{Est}$=.0016, Wide Prior: $\sigma_s^2{}_{Est}$=.04) and a different prior variance in the 2-Afc task (Narrow Prior: $\sigma_s^2{}_{Afc}$=1.6×10^{-5}, .0004, .0064, or .16; Wide Prior: $\sigma_s^2{}_{Afc}$= .0004, .01, .16, or 4), expecting to find a better fit for the 2-Prior model. We simulated 1000 subjects for each condition. To make sure that we could obtain reasonable parameter estimates, we performed bootstrapped parameter estimation on one simulated subject for each condition.

To simulate behavioral responses, we follow procedures used in previous work(Acuna et al., 2015; Kording & Wolpert, 2004), described here. We simulated the same amount of data as in the behavioral data set. Stimuli were generated, where targets, denoted *s*, were sampled from a prior distribution $N(\mu, \sigma_s^2)$ and subjects were given uncertain information on the target's location in the form of *n* samples from a likelihood distribution centered on the target, $N(s, \sigma_l^2)$. Simulated subjects estimated the centroid of the four samples, *c*, distributed according to the variance of the sample mean, σ_l^2/n. In both the estimation and 2-Afc tasks, simulated subjects optimally combined information from their learned prior, $N(\mu, \sigma_s^2)$, and the likelihood, $N(c, \sigma_l^2/n)$. They weigh the two sources of information according to their relative precision, resulting in a posterior distribution. Simulated responses in the estimation task, $\hat{s}$, were samples from the distribution of the maximum a posteriori (MAP):

$$\hat{s} = \mathrm{N}\left(\frac{\frac{\mu}{\sigma_s^2} + \frac{c}{\sigma_l^2/n}}{\frac{1}{\sigma_s^2} + \frac{1}{\sigma_l^2/n}}, \qquad \frac{\frac{1}{\sigma_l^2/n}}{\left(\frac{1}{\sigma_s^2} + \frac{1}{\sigma_l^2/n}\right)^2}\right) \quad (1)$$

The mean of the MAP distribution, μ_{MAP}, is therefore a precision-weighted mean of the prior mean, μ, and the splash centroid, *c*. The slope of the fitted μ_{MAP}

with c provides an indicator of how much subjects relied on the likelihood to form their estimate, which we refer to as the *Reliance on the likelihood*.

For each trial of the 2-Afc task, we generated the probability of the subject deciding that the probe target (splash with centroid c_{l1} and narrower likelihood, σ_{l1}^2) is to the right of the reference target (splash with centroid c_{l2} and broader likelihood, σ_{l2}^2). This probability, $P(z=1)$, was computed from the following psychometric function:

$$P(z=1) = \frac{\lambda}{2} + (1-\lambda)\left(\frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{c_{l1}(\sigma_{l2}^2/n+\sigma_s^2) - c_{l2}(\sigma_{l1}^2/n+\sigma_s^2)}{\sqrt{2}\sqrt{\sigma_{l1}^2/n(\sigma_{l2}^2/n+\sigma_s^2)^2 + \sigma_{l2}^2/n(\sigma_{l1}^2/n+\sigma_s^2)^2}}\right)\right]\right) \quad (2)$$

Derivations of Equation 2 are provided in a previous paper (Acuna et al., 2015). This function assumes a MAP decision rule. From the probabilities obtained from Equation 2, binary decisions were generated using Bernoulli sampling.

Parameter fitting, model comparison and parameter estimation

We used maximum-likelihood estimation to find the best-fitting parameters under the 1-Prior and 2-Prior models. We computed the likelihood of the model for a given set of parameters for the 1-Prior model, $\theta_{1\,prior} = \{\sigma_{s\,NP}^2, \sigma_{s\,WP}^2, \lambda\}$, and the 2-Prior model, $\theta_{2\,prior} = \{\sigma_{s\,NP\,Est}^2, \sigma_{s\,WP\,Est}^2, \sigma_{s\,NP\,Afc}^2, \sigma_{s\,WP\,Afc}^2, \lambda\}$, by generating predictions from Equations 1 and 2 and comparing these with the decision data.

In the estimation task, the log-likelihood of each estimate, $log(P(\hat{s}|\theta))$, was computed from the parameters of the MAP distribution in Equation 1:

$$\log(P(\hat{s}|\theta)) = \log\left(\frac{1}{\sqrt{2\pi\sigma_{MAP}}} e^{\frac{(\hat{s}-\mu_{MAP})^2}{2\sigma_{MAP}^2}}\right) \quad (3)$$

Since the 2-Afc task resulted in binary decisions, the log-likelihood was computed as the log-loss, where z is the decision outcome and p is $P(z=1)$ computed from Equation 2.

$$\log(P(z|\theta)) = z(\log(p)) + (1-z)\log(1-p)) \quad (4)$$

The log-likelihood was summed across individual trials to give an aggregate measure of the likelihood across tasks for the data of each individual subject. We found the parameters of the model using a maximum-likelihood estimation algorithm. To ensure an equal weighting of conditions and tasks, the number of trials was equalized by selecting the condition with the lowest number of trials and fixing the number of trials for remaining conditions. Bounds of 0 and 1 were used to constrain the prior variance parameters. Bounds of 0 and .1 were used for the lapse parameter. Initial parameters were randomly selected uniformly from the range of possible values.

To compare the performance of the 1-Prior model and the 2-Prior model, we used cross-validated model comparison. Individual data were divided into 10 folds, with each fold containing a random subset of trials. Data were stratified so that each fold contained the same number of trials from each condition and so that binned reference locations were equally represented in each fold. Parameters were fit to 9 out of the 10 folds using the parameter fitting. As a measure of the model's performance, the log-likelihood summed across the remaining held-out data set was used. This was performed once for each of the ten held-out sets.

To estimate the parameters of the best-fitting model, we used bootstrapped parameter estimation where we performed the procedure presented above 1000 times on data sets resampled with replacement from each individual's data. In the simulation results presented, we simulated one subject per Prior-Ratio condition.

Estimating the Reliance on the likelihood and PSE slope from behavioral data

In order to examine the use of probabilistic information in subjects' estimations, the MAP distribution described in Equation 1 was fit to subjects' estimation data using maximum-likelihood estimation and the parameters were used to compute the *Reliance on the likelihood* in the estimation task. In the estimation task, subjects gave a continuous estimate of target location. We wanted to quantify how much subjects relied on the learned target location (prior) or visual information (likelihood). To do so, we computed the relationship between the likelihood's center and estimations, which we termed the *Reliance on the likelihood*. If someone were to rely only on likelihood information to judge target location, on average, estimations should correspond to the center of the likelihood (*Reliance on the likelihood* =1). If someone were to ignore the likelihood entirely and rely only on their prior to judge target location, there should be no relationship between the likelihood's center and estimations (*Reliance on the likelihood* = 0). The Bayesian optimal strategy is to weigh the prior and likelihood according to their relative precision, *Reliance on the likelihood* = $\sigma_s^2/(\sigma_s^2 + \sigma_l^2/n)$.

In the 2-Afc task, subjects were given probabilistic information on target location exactly as in the estimation task. On each trial, they compared the locations of two targets with different uncertainties, a probe stimulus with Narrow Likelihood and a reference stimulus with Wide Likelihood. Uncertainty should influence the judgment of target location in the same way as in the estimation task. A Bayesian observer judges the more uncertain target to be shifted further to the prior mean than the more certain target. This, in turn, influences decisions about relative target location (z=0, probe target left; z=1, probe target right). Therefore, use of the prior can be inferred from subjects' 2-Afc data.

Consider the psychometric function that describes the comparison of targets with unequal variances (Fig. 2c). The psychometric function is the probability that the probe target is reported to the right, $P(z=1)$, as a function of difference between the likelihood centroids (*Discrepancy*), and the *Reference location*. The subject's prior influences the *Discrepancy* at which the targets are perceived as equal (point of subjective equality, *PSE*). For a Bayesian observer, the *PSE* arises when the reference is more distant from the prior's center than the probe. The *PSE* further deviates from zero as the distance between the prior and the reference increases. Importantly, the slope of this linear relationship, the *PSE slope*, is related to the width of the subject's prior – a *PSE slope* of 0 shows that subjects relied only on visual information from the likelihood; and the more negative the *PSE slope*, the narrower the subject's prior. The optimal *PSE slope* is given by *PSE slope* = $(\sigma_{l1}^2/n - \sigma_{l2}^2/n)/(\sigma_{l2}^2/n + \sigma_s^2)$(Acuna et al., 2015).

We fit psychometric functions (the cumulative Gaussian function) to each subject's decision data. The *PSE slope* (m_{PSE}) estimated from this function provides an indicator of the variance of the subject's prior. We model the probability of a decision as:

$$\mathrm{P(z=1)} = \frac{\lambda}{2} + (1-\lambda)\left(\frac{1}{2}\left[1 + \mathrm{erf}\left(\frac{\delta - c_{l2} \times m_{PSE}}{\sqrt{2}\,\sigma}\right)\right]\right) \quad (5)$$

where δ is the discrepancy between stimuli, c_{l2} is the location of the reference stimulus with broader likelihood, λ is a lapse parameter used to account for attentional lapses, σ describes the deviation of the function and assumes a MAP decision rule (Acuna et al., 2015). We find the values of m_{PSE}, λ and σ using a maximum-likelihood estimation algorithm.

References

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. http://doi.org/10.1167/15.3.7

Aglioti, S., DeSouza, J. F. X., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, *5*(6), 679–685. http://doi.org/10.1016/S0960-9822(95)00133-3

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 of the United States of America*, *110*(45), 18327–32. http://doi.org/10.1073/pnas.1306572110

Berniker, M., Voss, M., & Kording, K. (2010). Learning priors for bayesian computations in the nervous system. *PLoS ONE*, *5*(9), 1–9. http://doi.org/10.1371/journal.pone.0012686

Criscimagna-hemminger, S. E., Donchin, O., Gazzaniga, M. S., Shadmehr, R., Sarah, E., Donchin, O., & Michael, S. (2003). Learned Dynamics of Reaching Movements Generalize From Dominant to Nondominant Arm. *Journal of Neurophysiology*, 168–176.

Daw, N. D., & Doya, K. (2006). The computational neurobiology of learning and reward. *Current Opinion in Neurobiology*, *16*(2), 199–204. http://doi.org/10.1016/j.conb.2006.03.006

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–33. http://doi.org/10.1038/415429a

Fernandes, H. L., Stevenson, I. H., Vilares, I., & Kording, K. P. (2014). The generalization of prior uncertainty during reaching. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(34), 11470–84. http://doi.org/10.1523/JNEUROSCI.3882-13.2014

Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, *14*(3), 119–30. http://doi.org/10.1016/j.tics.2010.01.003

Haith, A. M., & Krakauer, J. W. (2013). Theoretical models of motor control and motor learning. In *Routledge handbook of motor control and motor learning.* (Routledge, pp. 7–28). London.

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), 967–992. http://doi.org/10.1167/4.12.1

Knill, D. (2005). Reaching for visual cues to depth: the brain combines depth cues differently for motor control and perception. *Journal of Vision*, *5*(2005), 103–115. http://doi.org/10.1167/5.2.2

Kording, K., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247. http://doi.org/10.1038/nature02169

Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*, 436–444. http://doi.org/10.1038/nature14539

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–8. http://doi.org/10.1038/nn1790

Maloney, L. T. (2002). Statistical decision theory and biological vision. In H. D & M. R (Eds.), *Perception and the physical world* (pp. 145–189). New York: Wiley.

Mamassian, P., Landy, M., & Maloney, L. T. (2001). Bayesian Modelling of Visual Perception. In R. P. N. Rao, B. A. Olhausen, & M. S. Lewicki (Eds.), *Probabilistic models of the brain: Perception and neural function* (pp. 13–34). Cambridge, MA: MIT Press.

Marblestone, A. H., Wayne, G., & Kording, K. P. (2016). Toward an Integration of Deep Learning and Neuroscience. *Frontiers in Computational Neuroscience*, *10*(September), 1–41. http://doi.org/10.3389/fncom.2016.00094

Perfors, A., Tenenbaum, J. B., & Regier, T. (2011). The learnability of abstract syntactic principles. *Cognition*, *118*(3), 306–338. http://doi.org/10.1016/j.cognition.2010.11.001

Pitkow, X., & Angelaki, D. E. (2017). Perspective How the Brain Might Work : Statistics Flowing in Redundant Population Codes. *bioRxiv*, *1702.03492*, 1–9.

Rao, R. P. N. (2004). Bayesian Computation in Recurrent Neural Circuits. *Neural Computation*, *16*(1), 1–38.

Shadmehr, R. (2004). Generalization as a behavioral window to the neural mechanisms of learning internal models. *Human Movement Science*, *23*(5), 543–568. http://doi.org/10.1016/j.humov.2004.04.003

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT press.

Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(40), 10154–63. http://doi.org/10.1523/JNEUROSCI.2779-06.2006

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318. http://doi.org/10.1016/j.tics.2006.05.009

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science (New York, N.Y.)*, *331*(6022), 1279–1285. http://doi.org/10.1126/science.1192788

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. http://doi.org/10.1016/j.cub.2012.07.010

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? Optimal decisions from very few samples. *Cognitive Science*, *38*(4), 599–637. http://doi.org/10.1111/cogs.12101

Xu, F., & Tenenbaum, J. B. (2007). Word Learning as Bayesian Inference. *Psychological Review*, *114*(2), 245–250.

Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, *10*(2), 403–430.