Learning and inference using complex generative models in a spatial localization task

Vikranth R. Bejjanki
Department of Brain & Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY, USA
Department of Psychology and Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA

David C. Knill
Department of Brain & Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY, USA

Richard N. Aslin
Department of Brain & Cognitive Sciences and Center for Visual Science, University of Rochester, Rochester, NY, USA

A large body of research has established that, under relatively simple task conditions, human observers integrate uncertain sensory information with learned prior knowledge in an approximately Bayes-optimal manner. However, in many natural tasks, observers must perform this sensory-plus-prior integration when the underlying generative model of the environment consists of multiple causes. Here we ask if the Bayes-optimal integration seen with simple tasks also applies to such natural tasks when the generative model is more complex, or whether observers rely instead on a less efficient set of heuristics that approximate ideal performance. Participants localized a "hidden" target whose position on a touch screen was sampled from a location-contingent bimodal generative model with different variances around each mode. Over repeated exposure to this task, participants learned the a priori locations of the target (i.e., the bimodal generative model), and integrated this learned knowledge with uncertain sensory information on a trial-by-trial basis in a manner consistent with the predictions of Bayes-optimal behavior. In particular, participants rapidly learned the locations of the two modes of the generative model, but the relative variances of the modes were learned much more slowly. Taken together, our results suggest that human performance in a more complex localization task, which requires the integration of sensory information with learned knowledge of a bimodal generative model, is consistent with the predictions of Bayes-optimal behavior, but involves a much longer time-course than in simpler tasks.

Introduction

Humans (and other animals) operate in a world of sensory uncertainty, created by noise or processing inefficiencies within each sensory modality or by variability in the environment (Knill & Pouget, 2004). Given the presence of such internal and external uncertainty, task performance is limited by the quality of sensory information that is available on a given trial. Consider for example, the task of estimating the location of a target based on visual information, such as the problem faced by a baseball player trying to hit a rapidly approaching ball. Due to uncertainty in the visual signal available to the observer, an ideal approach to estimating the ball's location at the point of impact with the bat would be to combine the uncertain sensory information available on each pitch with prior knowledge about the likely locations of the ball, such as the areas in the strike zone that a given pitcher prefers to target, learned over prior encounters with that pitcher. Bayesian inference provides a principled approach for accomplishing this in a statistically efficient fashion; that is, in a manner that maximally utilizes the available information (Bernardo & Smith, 1994; Cox, 1946; Yuille & Bülthoff, 1996).

Citation: Bejjanki, V. R., Knill, D. C., & Aslin, R. N. (2016). Learning and inference using complex generative models in a spatial localization task. *Journal of Vision, 16*(5):9, 1–13, doi:10.1167/16.5.9.

doi: 10.1167/16.5.9 Received August 29, 2015; published March 10, 2016 ISSN 1534-7362

Several studies have shown that human observers combine prior information with uncertain sensory information in a manner predicted by Bayes-optimal behavior, in tasks including visual motion perception (Stocker & Simoncelli, 2006; Weiss, Simoncelli, & Adelson, 2002), visuo-motor integration (Kording & Wolpert, 2004; O'Reilly, Jbabdi, Rushworth, & Behrens, 2013; Sato & Kording, 2014; Tassinari, Hudson, & Landy, 2006; Vilares, Howard, Fernandes, Gottfried, & Kording, 2012), timing behavior (Jazayeri & Shadlen, 2010; Miyazaki, Nozaki, & Nakajima, 2005), cue combination (Adams, Graf, & Ernst, 2004; Jacobs, 1999; Körding et al., 2007), categorical judgments (Bejjanki, Clayards, Knill, & Aslin, 2011; Huttenlocher, Hedges, & Vevea, 2000), and movement planning (Hudson, Maloney, & Landy, 2007; Kwon & Knill, 2013). For example, considering the task of estimating the spatial location of a target based on uncertain visual information, previous studies have shown that human observers combine visually presented prior information (in the form of a Gaussian blob) with uncertain sensory information, in a Bayes-optimal fashion (Tassinari et al., 2006). Furthermore, several studies have shown that human observers are also capable of learning prior distributions over repeated exposure to a task (when receiving a single sample from the prior distribution on each trial), and integrating this learned knowledge with uncertain sensory information in a Bayes-optimal fashion (Berniker, Voss, & Kording, 2010; Kording & Wolpert, 2004; O'Reilly et al., 2013; Vilares et al., 2012). It is important to note that these studies used tasks in which observers were faced with uncertain sensory information about stimuli drawn from relatively simple prior distributions, or generative models in the language of Bayesian analysis, such as unimodal Gaussian distributions. An example of a task with such a simple generative model would consist of searching for a reward at a single hiding place, where the reward is replenished according to a Gaussian delay interval. Even in cases where the distributional properties of the generative model were changed over the course of the experiment (Berniker et al., 2010; O'Reilly et al., 2013; Vilares et al., 2012), this was done in a blocked manner, such that in a given block of trials, observers were faced with stimuli drawn from a unimodal Gaussian distribution.

In contrast to the tasks considered by these previous studies, in the natural environment, human observers are typically faced with tasks involving more complex generative models. Specifically, the stimulus of interest on each trial is usually drawn from one of a number of potential environmental causes, in a randomly interleaved (i.e., not blocked) fashion. If we assume that the statistical properties of the stimulus distribution associated with a given environmental cause are represented by a Gaussian distribution, then such a complex generative model is best described by a mixture model over multiple Gaussian distributions. An example of a task with such a complex generative model would consist of searching for a reward inside one of several hiding places, each of which has its own Gaussian delay interval. Furthermore, the sensory information available to observers in their natural environment typically includes several levels of uncertainty, again in a randomly interleaved fashion. Given these task properties, ideal behavior involves learning the complex generative model (i.e., the mixture model over all possible causes) for the task by learning the underlying distributions pertaining to each cause, associating the uncertain sensory information available on each trial with the relevant cause that might have generated it, and efficiently integrating the sensory information with the learned generative model on a trial-by-trial basis.

Given the significant increase in complexity associated with learning and efficiently using such complex generative models, it remains an open question as to whether human observers continue to behave in a Bayes-optimal manner. For instance, as the number of causes goes up, it might be prohibitively expensive from a computational perspective (i.e., the combinatorial explosion problem) to track and learn the complete generative model for the task. Observers might instead choose to employ suboptimal heuristics or strategies to simplify task performance, such as basing their performance on only the mean of the sensory information available across a series of trials, and not on the variance of that sensory information. It is worth noting that results from a prior study using a bimodal generative model (i.e., simulating a mixture model over two environmental causes) suggest that human observers can learn and utilize such complex generative models in an approximately Bayes-optimal manner (Kording & Wolpert, 2004). However, this experiment only involved one level of sensory uncertainty and the two mixture components of the generative model did not differ in their variance (i.e., in the reliability with which they signaled the underlying cause-specific stimulus information). As a result, since it did not vary the reliability of either the sensory information or the prior information, this experiment was not amenable to fully testing the predictions of the Bayesian model in tasks involving complex generative models; that is, the prediction that the weight assigned to the two sources of information should vary, on a trial-by-trial basis, in proportion to the reliability of each source of information. Furthermore, this experiment did not examine the dynamics of the process by which the complex generative model was learned.

In the current study, we examined human behavior when faced with stimuli drawn from a complex generative model, using a spatial localization task that allowed us to vary the uncertainty implicit in both the sensory information and the prior information, on a trial-by-trial basis. Our design simulated a mixture model over two environmental causes, and, as a first logical step, considered the scenario in which the participant is perfectly cued as to the underlying cause on each trial. We implemented this using a location-contingent generative model—the task-relevant stimulus on each trial was drawn from a mixture of two underlying Gaussian distributions, with one distribution always centered at one location on the screen and the other always centered at a distinct second location on the screen, and with one distribution being more variable (i.e., representing a less reliable a priori signal about the likely location of a target drawn from that distribution) than the other. On each trial, the location of the target was drawn randomly from this bimodal spatial distribution, and participants were presented with a cluster of dots that was randomly drawn from a separate Gaussian distribution centered on this true "hidden" location. The cluster of dots, depending on its variance (which could be low, medium, or high), provided a more or less reliable estimate of the target's hidden location on each trial (reliability being inversely proportional to the variance of the dot cluster). Across trials, observers could integrate information from the observed cluster of dots (i.e., the likelihood), with trial-by-trial feedback they received on previous trials about the underlying bimodal distribution governing the likely locations of the target (i.e., the "prior"). This experimental design allowed us to better study the computational mechanisms used by human observers in tasks involving complex generative models than was possible with previous studies. Specifically, while aspects of our task are similar to those used in previous studies (especially Berniker et al., 2010; Kording & Wolpert, 2004; Tassinari et al., 2006; and Vilares et al., 2012) by randomly interleaving multiple underlying distributions (which varied in their reliability) and multiple levels of sensory uncertainty, on a trial-by-trial basis, our design allows us to fully test the predictions of the Bayesian model in tasks involving complex generative models, and to rule out alternative suboptimal strategies that could potentially be used to carry out the task.

There are several computational models that participants could potentially use to estimate the location of the hidden target on each trial of this task. As mentioned above, for example, due to the complexity inherent in the statistics that govern target location across trials, participants might choose to ignore this "latent" source of information and instead base their estimates solely on the sensory information available on each trial. For instance, they might simply choose the centroid of the cluster of dots as their estimate for the target location on each trial. An alternate strategy might be to learn some summary statistic, such as the average location, that describes each mixture component of the generative model and to bias their estimates toward that source of information. Both these kinds of strategies would, however, be statistically suboptimal because they would not maximally exploit all the sources of information available to the observer. Instead, a statistically optimal approach would entail learning the distributional statistics (both the mean and the variance) characterizing each mixture component of the generative model, and integrating this information with the sensory information available on each trial, in a manner consistent with the predictions of Bayes-optimal behavior. This approach predicts that the weight assigned by participants to the sensory information should drop as the reliability of the sensory information goes down (i.e., the variance of the cluster of dots goes up), and, crucially, this drop should be greater when the stimulus is drawn from the mixture component that has lower variance (i.e., greater reliability). By following this strategy, participants would maximally utilize all the sources of information available on each trial by appropriately weighting each source of information by the uncertainty implicit in it.

We provide evidence from two experiments that, over multiple exposures to the task of estimating the location of a hidden target, human observers not only learn the complex distributional statistics about the a priori locations of the target (i.e., they learn the location-contingent bimodal generative model), but also integrate this knowledge with uncertain sensory information on a trial-by-trial basis, in a manner consistent with the predictions of Bayes-optimal behavior. Importantly, our experimental design allowed us to validate the predictions of the Bayesian model in a very specific manner, and thus rule out alternative suboptimal strategies that could be used to carry out the task. Consistent with the predictions of the Bayesian model, we found that participants assigned a smaller weight to the sensory information as the reliability of sensory information went down, and that this drop was greater when the stimulus was drawn from the more reliable mixture component of the complex generative model used in the task. Our design also allowed us to examine the dynamics underlying the process by which participants learned the complex generative model. Extending prior work using simple generative models (Berniker et al., 2010), we found that participants quickly learned the mean locations for the bimodal generative model (i.e., the centers of the two distributions that made up the mixture components of the generative model for the task), but it took them much longer to learn the relative variances of the mixture components of these two modes. Furthermore, our results show that in tasks with such complex generative models, the learning rate is dramatically

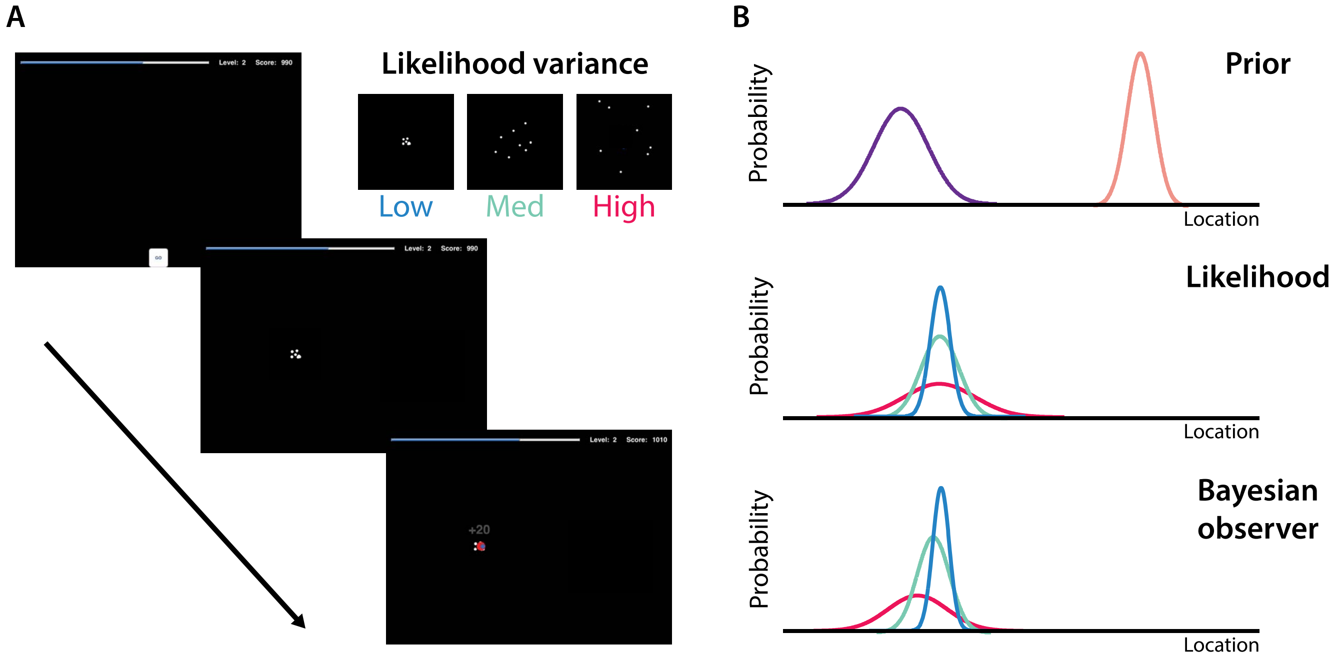

Figure 1. Learning and inference in a spatial localization task. (A) An illustration of a typical trial. After participants touched a "GO" button, they were presented with uncertain sensory information in the form of a cloud of dots (the likelihood) with one of three levels of variance (low variance shown here; see inset for an illustration of the three levels). Participants estimated the location of a hidden target, drawn from a complex generative model (a mixture over two Gaussian distributions: the prior), by touching a location on the display. Feedback was provided posttouch. (B) An illustration of Bayes-optimal behavior. Considering the example of a target drawn from the broad prior distribution, the ideal observer would learn the mean and variance of the prior and integrate this learned knowledge with the likelihood on each trial, to estimate the location of the hidden target.

slower than that observed previously with simple generative models, presumably due to the significantly increased complexity involved in tracking and learning the means and variances of the generative model, especially when sensory uncertainty was also variable. Taken together, our results therefore provide compelling evidence in support of the hypothesis that, when faced with complex environmental statistics, human observers are capable of learning such distributions, and utilizing this learned knowledge in a statistically efficient fashion.

General methods

We developed a spatial localization task where the participants' goal on each trial was to estimate the location of a hidden target by touching the appropriate location on a touch-sensitive display (Figure 1A). Several aspects of our task were similar to the task used in Kording and Wolpert (2004), Tassinari et al. (2006), and Vilares et al. (2012). The horizontal and vertical coordinates of the target location on each trial were independently drawn from a complex generative model: a location-contingent mixture distribution over two underlying isotropic two-dimensional Gaussian distributions. These Gaussian distributions differed in their mean locations (one of them was always centered in the left half of the display while the other was always centered in the right half of the display), and their relative variances (one of them had a standard deviation of 40 pixels while the other had a standard deviation of 20 pixels). Across participants, the distributions were always centered at the same locations on the left and right of the display, but the variance assigned to each was counterbalanced with the higher variance distribution centered in the left half of the display for 50% of the participants, and in the right half of the display for the other 50%.

On each trial, participants were not shown the true target location; instead, they were presented with uncertain sensory information about the target location in the form of a two-dimensional cloud of eight small dots (Figure 1A). This cloud was independently generated on each trial by drawing samples from a separate two-dimensional isotropic Gaussian distribution centered at the true target location for that trial (i.e., the location of the target drawn for that trial from the complex generative model). The variance of this cloud of dots was manipulated to generate sensory information with three levels of reliability, with reliability inversely proportional to the variance of the dot distribution. Specifically, the dots were drawn from a distribution that either had low variance (*SD* of 10 pixels), medium variance (*SD* of 60 pixels), or high variance (*SD* of 100 pixels; see inset of Figure 1A for an illustration of the three levels of variance). After participants provided their response by touching a location on the display, feedback on their accuracy on that trial was provided by displaying a new dot at the touched location and a second new dot at the true target location. In addition, participants received feedback in the form of numerical points, with the magnitude of the points varying based on their accuracy.

Experiment 1: Can human observers learn and efficiently use complex prior knowledge?

In Experiment 1, participants were exposed to 1,200 trials of the spatial localization task, with half the trials involving targets drawn from each of the two underlying Gaussian distributions (the mixture components of the location-contingent generative model). These two sets of 600 trials were randomly intermixed and within each set, 200 trials were randomly presented with one of the three levels of reliability; in other words, six total conditions were randomly interleaved throughout the experiment. In each trial, participants therefore had access to two sources of information—the uncertain sensory information available on that trial (i.e., the cloud of dots that corresponded to the likelihood) and any knowledge they might have learned up to that point in the experiment, based on feedback received in previous trials, about the likely locations of the target (i.e., the mixture distribution over the two underlying Gaussian distributions that corresponded to the prior). Participants received two cues on each trial as to which of the two underlying Gaussian distributions the target on that trial was drawn from: the side of the display in which the sensory information was presented (the two underlying distributions were always centered in opposite halves of the display) and the color of the sensory information (the color of the dot cloud was set either to green or white, depending on the target distribution). Since the components of the complex generative model were perfectly separable given the location contingent cues on each trial, we can characterize the prior as involving two separable distributions, conditioned on the trial-specific location and color of the sensory information. Accordingly, in the rest of this article, we refer to the two Gaussian distributions that make up the mixture components of the complex generative model, as the "broad" prior distribution (the distribution with the larger standard deviation) and the "narrow" prior distribution (the distribution with the smaller standard deviation).

There are several computational models that participants could potentially use to carry out this task. One possibility is that participants just ignore the underlying prior distributions and base their estimates solely on the sensory information available on each trial. That is, given the complexity implicit in the generative model governing target location, participants can ignore this source of information (rather than expend the computational resources to track and learn the location-contingent prior distributions) and instead just choose the centroid of the cluster of dots (which are visible on each trial) as their estimate for the hidden target location. This strategy would predict that participants would assign a weight of one to the likelihood (i.e., base their responses solely on the sensory information), irrespective of the reliability of the sensory information or the reliability (or inverse variance) of the prior distribution for that trial. Furthermore, we would expect to see no change in participants' weights as a function of exposure to the task, since this strategy does not involve any learning (cf., the likelihood distribution is drawn independently on each trial). Importantly, however, such a computational strategy (henceforth referred to as Model 1) would be suboptimal for two reasons: first, as the reliability of the sensory information goes down, the probability that the centroid of the sensory information will correspond exactly to the true target location also goes down; and second, by ignoring the information provided by the prior distributions, participants cannot take advantage of the fact that not all spatial locations are a priori equally likely.

A second possibility is that participants are sensitive to the underlying generative model but given the complexity involved in the task, they choose to learn only the mean locations for the two prior distributions while ignoring their relative variances (i.e., their relative reliability)—learning only the means is statistically easier and requires exposure to far fewer trials. By learning the average position for where the target tends to be, for each prior distribution participants can use this knowledge to bias their estimates, particularly as the reliability of the sensory information goes down. This model (Model 2) predicts a smaller weight to the likelihood, as the reliability of the sensory information goes down, and it also predicts that as participants gain greater exposure to the underlying distributions (thus allowing them to better learn the two prior means), their behavior should show a greater sensitivity to the underlying prior distributions. However, a key prediction of this model is that for each level of sensory reliability, we should see identical weights being assigned to the likelihood, for targets drawn from the two different prior distributions, because the variances of the prior distributions are ignored by this model.

Thus, this model is still suboptimal since the narrow prior distribution is a more reliable indicator of the likely target location across trials, compared to the broad prior distribution, and an optimal observer should take this into account.

Finally, a third possibility is that participants learn both the mean positions and the relative reliabilities of the two prior distributions, and integrate this learned knowledge with the sensory information available on each trial in a manner that is consistent with the predictions of Bayes-optimal behavior (Model 3). If participants use such a strategy, we should see a smaller weight to the likelihood as the reliability of the sensory information goes down, and this drop should be greater for trials in which the target is drawn from the narrow prior, than when the target is drawn from the broad prior. Furthermore, this model predicts that, as participants gain greater exposure to and thus learn more about the underlying distributions, participants' behavior should show a greater sensitivity to the statistical properties (both mean location and relative variance) of the underlying distributions (see Figure 1B for an illustration). Formally, if the mean and the variance of the sensory information on a given trial is given by μ_l and σ_l^2, and the mean and the variance of the underlying target distribution for that trial is given by μ_p and σ_p^2, then the target location $\hat{t}$ predicted by this model would be (Bernardo & Smith, 1994; Cox, 1946; Jacobs, 1999; Yuille & Bülthoff, 1996):

$$\hat{t} = w_l \mu_l + (1 - w_l)\mu_p \quad (1)$$

where w_l, the weight assigned by the observer to the sensory information, should be:

$$w_l = \frac{1/\sigma_l^2}{1/\sigma_l^2 + 1/\sigma_p^2} \quad (2)$$

By examining the weights assigned by participants to the likelihood on each trial, given the relative reliability of the sensory information and the prior information on that trial, we sought to distinguish between the three potential strategies described above (i.e., Models 1, 2, and 3). In addition, we sought to characterize learning-induced changes in participants' weights, as a function of exposure to the task, by splitting the total number of trials into four temporal bins, each of which included 300 trials of exposure to the task (with the six conditions randomly interleaved in each bin).

Methods

Participants

Eight undergraduate students at the University of Rochester participated in this experiment, in exchange for monetary compensation. Each participant had normal or corrected-to-normal vision, was naïve to the purpose of the study, and provided informed written consent. The University of Rochester's institutional review board approved all experimental protocols.

Procedure

Before the start of the experiment, participants were provided with the following task instructions:

> In this study, you will be playing a game on the iPad. The story behind this game is that you are at a fair and there is an invisible bucket that you are trying to locate. Sometimes the bucket is going to be located on the left side of the display and at other times the bucket is going to be located on the right side of the display. Now, given that the bucket is invisible, you can't see where it is. However, on each trial you will see some locations that other people have previously guessed the bucket is located at. These "guesses" will show up as white or green dots on the screen. Now, it is important to note that you don't know which (if any) of the dots actually correspond to the location of the bucket. Indeed, all of the dots could be just random guesses. Your job on each trial is to try to figure out where the bucket is actually located. Once you decide on a location, you can just touch it, at which point you will see two more dots: a red dot, which shows you the true location of the bucket on that trial, and a blue dot, which shows you the location that you touched. If the blue dot is right on the red dot, then you correctly guessed the location of the bucket and you will get 20 points. If you don't exactly guess the location of the bucket but you still get close, then you will get 5 points. When you see this, it means that you need to try just a little harder to get the right location—you are close. Finally, if your guess is very far away, you will get no points.

Participants were seated at a comfortable viewing distance from a touch-sensitive display (iPad 2; Apple, Inc., Cupertino, CA), with a resolution of 768 pixels (vertical) × 1024 pixels (horizontal). Participants provided responses using their finger, having been instructed to pick a favorite finger and to consistently use that finger throughout the experiment. At the start of each trial, they were presented with a "GO" button centered at the bottom of the display. Once they touched this button, the trial started with the presentation of a cluster of white or green dots (the "guesses"). Participants were told to estimate the location of the hidden target as rapidly and accurately as possible by touching a location on the screen. After participants provided a response, the cluster of dots

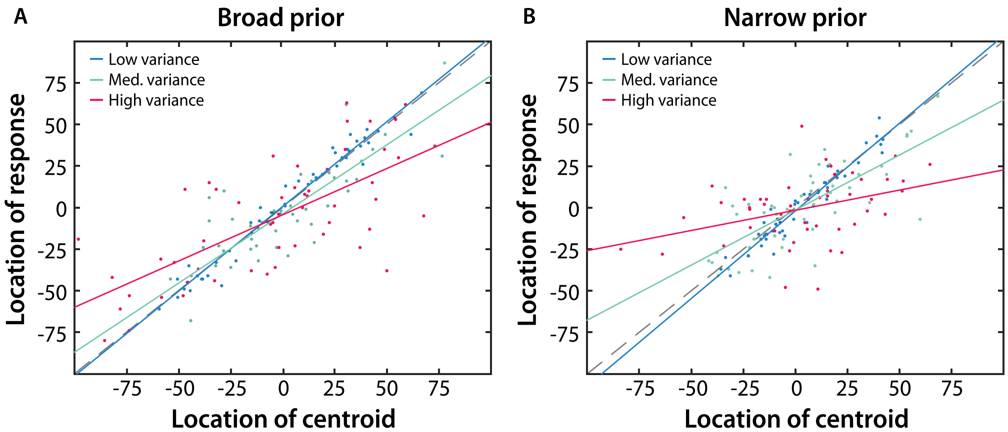

Figure 2. Data from the final temporal bin, for a representative participant in Experiment 1. Consistent with the predictions of Bayes-optimal behavior, as likelihood variance increased (blue to green to red), participants shifted from selecting the centroid of the sensory information as their estimate for the target location (the gray dashed line), toward selecting the mean of the underlying prior distribution (zero on the y-axis). This was true in both the broad (A) and narrow (B) prior conditions, but the shift was more pronounced for the narrow prior condition than for the broad prior condition. For illustrative purposes, the mean of the underlying prior distribution was removed from participants' response locations, and from the locations of the centroid of the sensory information. Each dot represents a trial and solid lines represent the best-fit regression lines in each condition.

disappeared and feedback was provided. The game included multiple levels: once a participant accumulated a total of 600 points on each level, they were shown a congratulatory screen and allowed to "level up." Progress toward the next level was always shown at the top of the screen, using a progress bar. This ability to level up was used solely to motivate participants and did not represent any change in the experimental procedure. Participants carried out a total of 1,200 trials, split across four experimental blocks. Short breaks were allowed between blocks, and the total experimental duration, including breaks, was approximately 50 min.

Data analysis

To quantify the computational mechanisms used by participants in this experiment, we estimated the extent to which their behavior depended on the likelihood versus the prior on each trial. Specifically, we used linear regression to compute the weight (w_l) assigned to the centroid of the cluster of dots (the likelihood), with the weight assigned to the mean of the underlying target distribution for that trial (the prior) being defined as $(1 - w_l)$. Thus, on each trial, given the centroid of the sensory information (μ_l), the mean of the underlying target distribution (μ_p) and participants' estimate for the target location ($\hat{t}$), the weight assigned by participants to the centroid of the sensory information (w_l) was estimated using:

$$\hat{t} = w_l\mu_l + (1 - w_l)\mu_p + \text{noise} \quad (3)$$

We focused on participants' performance in the vertical dimension to eliminate the potential influence of variability that may have been introduced by an interaction between participants' handedness and the horizontal separation of the prior locations. Moreover, their performance in the horizontal dimension was similar to their performance in the vertical dimension—we found no interaction between dimension (horizontal vs. vertical) and exposure bin, across prior and likelihood conditions (all ps > 0.05).

Results and discussion

Consistent with the predictions of the Bayes-optimal model (Model 3), by the end of exposure to the task (i.e., in the final temporal bin) observers assigned a reliably smaller weight to the likelihood (and thus, a greater weight to the prior) as the sensory information decreased in reliability (Figure 2). In a 2 (prior variance) $\times$ 3 (likelihood variance) repeated measures ANOVA carried out over participants' weights in the final temporal bin, there was a main effect of likelihood variance, $F(2, 14) = 18.89$, $p < 0.001$.

Furthermore, for each prior condition, there was a reliable interaction between exposure and likelihood variance: as participants gained more exposure to the task, they assigned a smaller weight to the likelihood, and this drop in weight was greater as the likelihood variance increased (Figure 3). Specifically, in a 4 (temporal bin) $\times$ 3 (likelihood variance) repeated measures ANOVA carried out over participants' weights in the broad prior condition, there was a main effect of temporal bin, $F(3, 21) = 4.00$, $p = 0.02$, a main

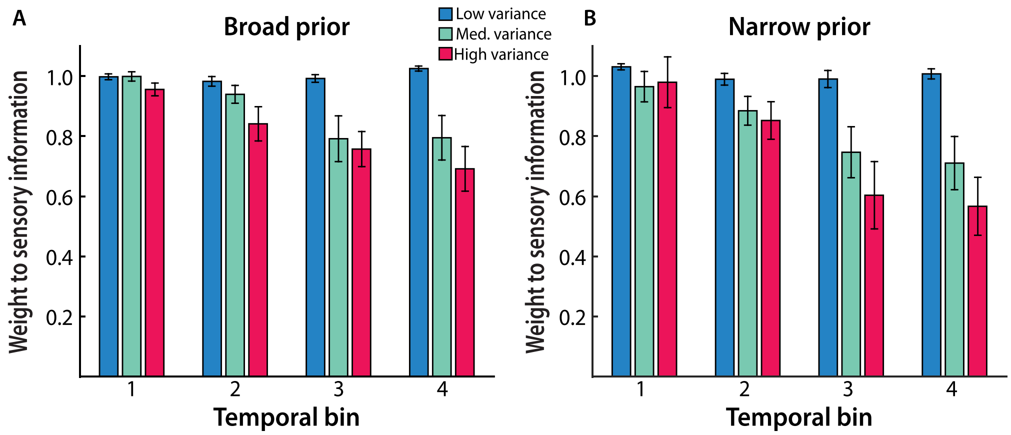

Figure 3. Weights assigned to the centroid of the sensory information in Experiment 1. Consistent with the predictions of Bayes-optimal behavior, participants in Experiment 1 relied less on the sensory information, and more on the priors, as the variance of the sensory information increased. This was true for both broad (A) and narrow (B) prior conditions, with the drop in the weight to the sensory information being greater for the narrow prior condition than for the broad prior condition. Furthermore, for both prior conditions, as variance increased, the drop in the weight to the sensory information was greater as participants gained more exposure to the task. Each temporal bin included 300 trials split between the two prior conditions, and the four bins of trials are depicted in temporal order. Columns represent means and error bars represent *SEM* across participants.

effect of likelihood variance, $F(2, 14) = 21.17$, $p < 0.0001$, and an interaction between the two factors, $F(6, 42) = 5.22$, $p < 0.001$. Similarly, in a 4 (temporal bin) × 3 (likelihood variance) repeated measures ANOVA carried out over participants' weights in the narrow prior condition, there was a main effect of temporal bin, $F(3, 21) = 9.06$, $p < 0.001$, a main effect of likelihood variance, $F(2, 14) = 12.24$, $p < 0.001$, and an interaction between the two factors, $F(6, 42) = 5.47$, $p < 0.001$. These findings are inconsistent with Model 1, which predicts a weight of 1 to the sensory information across all the likelihood and prior conditions and no change in participants' weights as a function of exposure to the task. Furthermore, in contrast to the predictions of Model 2, we found that the drop in the weight assigned by participants to the likelihood as a function of exposure to the task, particularly in the high likelihood variance condition, was greater for trials in which the target was drawn from the narrow prior, than when the target was drawn from the broad prior. In a 4 (temporal bin) × 2 (prior variance) repeated measures ANOVA carried out over participants' weights in the high likelihood variance condition, we saw a main effect of temporal bin, $F(3, 21) = 8.29$, $p < 0.001$, and an interaction between temporal bin and prior variance, $F(3, 21) = 3.07$, $p = 0.05$. This finding suggests that participants are sensitive to, and learn, both the means and the relative variances of the underlying prior distributions. Taken together, this pattern of results provides clear evidence in support of Model 3—the hypothesis that human observers learn the complex generative model and use this learned knowledge in a manner that is consistent with the predictions of Bayes-optimal behavior.

Experiment 2: Exploring the dynamics of learning

While the results from Experiment 1 support the hypothesis that participants learn the complex generative model and use this knowledge in a statistically efficient manner, the dynamics of this learning remain unclear. For instance, prior work (Berniker et al., 2010) has shown that when presented with stimuli drawn from simple generative models, human observers rapidly learn the prior mean, but it takes them many more trials to learn the prior variance. In Experiment 2, with a new group of participants, we examined the dynamics of learning with complex generative models by including an extra condition, randomly interleaved with all the conditions used in the first experiment, in which no sensory information was presented. In this condition, observers thus had to estimate the position of the target based solely on their learned knowledge (up to that point) about where the target was likely to occur (i.e., based on their prior knowledge), thereby allowing us to analyze their evolving learned knowledge of the location-contingent generative model as a function of exposure to the task. Participants carried out 400 trials of this condition, split evenly between the two prior distributions. On each trial, they were cued to estimate the location of the target on either the left or the right of the display (with a box that was colored green or white), thus requiring them to draw on their learned knowledge of the location-contingent prior distribution associated with that trial. As in Experiment 1, the 1,600 trials in Experiment 2 (all conditions from Experiment 1 plus 400 trials in this extra condition) were again split up into the relevant prior and likelihood conditions (the no sensory information condition is statistically equivalent to a likelihood condition with infinite uncertainty), and with each condition being further split into four temporal bins to study the influence of task exposure.

Methods

Participants

Eight undergraduate students at the University of Rochester participated in this experiment, in exchange for monetary compensation. Each participant had normal or corrected-to-normal vision, was naïve to the purpose of the study, provided informed written consent and did not participate in Experiment 1. The University of Rochester's institutional review board approved all experimental protocols.

Procedure

Before the start of the experiment, participants were provided with task instructions that were nearly identical to those provided to participants in Experiment 1. Specifically, in addition to the instructions from Experiment 1, they were also provided with the following instructions:

> There will also be some trials in which you will be the first person to guess the location of the bucket. In these trials, rather than seeing dots on the screen, you will see a briefly flashed white or green rectangle, which will indicate the side of the screen that the invisible bucket is located at—the bucket could be located anywhere inside that rectangle. Again, your job is to try to figure out where the bucket is actually located.

On each trial of the experiment, the procedure was identical to that used in Experiment 1, except for the extra condition in which no sensory information was presented. In the trials corresponding to this condition, after participants touched the GO button, they saw a briefly flashed green or white rectangle, spanning either the left half or the right half of the display. After participants provided a response, feedback was provided as in all the other conditions. Participants carried out a total of 1,600 trials, split across four experimental blocks. Short breaks were allowed between blocks, and the total experimental duration, including breaks, was approximately an hour.

Data analysis

As in Experiment 1, for the trials in which sensory information was available, we used linear regression to compute the weight (w_l) assigned to the centroid of the cluster of dots (the likelihood), with the weights assigned to the mean of the underlying location-contingent target distribution for that trial (the prior) being defined as $(1 - w_l)$. For the trials in which no sensory information was available, we computed participants' mean responses, across all trials in each temporal bin. As in Experiment 1, we again focused on performance in the vertical dimension.

Results and discussion

We first examined the weights assigned by participants to the likelihood and the prior for the trials in which sensory information was available, and found that their behavior was similar to that in Experiment 1 (Figure 3). Specifically, in line with the Bayes-optimal model (Model 3), by the end of exposure to the task in Experiment 2 (i.e., in the final temporal bin), observers assigned a reliably smaller weight to the likelihood (and a greater weight to the prior) as the sensory information decreased in reliability (Figure 4). In a 2 (prior variance) × 3 (likelihood variance) repeated measures ANOVA carried out over participants' weights in the final temporal bin, there was a main effect of likelihood variance, $F(2, 14) = 45.84$, $p < 0.0001$, and an interaction between the two factors, $F(2, 14) = 3.83$, $p = 0.047$. Furthermore, for each prior condition, we again found an interaction between exposure and likelihood variance. Specifically, in a 4 (temporal bin) × 3 (likelihood variance) repeated measures ANOVA carried out over participants' weights in the broad prior condition, there was a main effect of temporal bin, $F(3, 21) = 7.49$, $p = 0.001$, a main effect of likelihood variance, $F(2, 14) = 34.21$, $p < 0.0001$, and an interaction between the two factors, $F(6, 42) = 2.84$, $p = 0.02$. Similarly, in a 4 (temporal bin) × 3 (likelihood variance) repeated measures ANOVA carried out over participants' weights in the narrow prior condition, there was a marginal main effect of temporal bin, $F(3, 21) = 2.87$, $p = 0.06$, a main effect of likelihood variance, $F(2, 14) = 17.72$, $p < 0.001$, and a marginal interaction between the two factors, $F(6, 42) = 2.30$, $p = 0.052$. Finally, participants assigned a smaller weight to the likelihood, particularly in the high likelihood variance condition, when the target was drawn from the narrow prior, than when the target was drawn from the broad prior. In a 4 (temporal bin) × 2 (prior

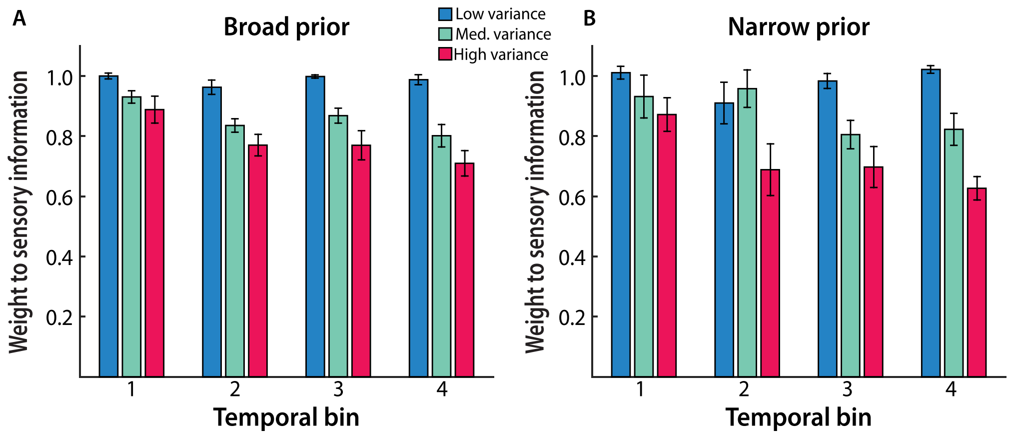

Figure 4. Weights assigned to the centroid of the sensory information in Experiment 2. As in Experiment 1, participants in Experiment 2 relied less on the sensory information and more on the priors as the variance of the sensory information increased, in line with the predictions of Bayes-optimal behavior. This was true for both broad (A) and narrow (B) prior conditions, with the drop in the weight to the sensory information being greater for the narrow prior condition than for the broad prior condition. Furthermore, for both prior conditions, as variance increased, the drop in the weight to the sensory information was greater as participants gained more exposure to the task. Each temporal bin again included 300 trials split between the two prior conditions, and the four bins of trials are depicted in temporal order. Columns represent means and error bars represent *SEM* across participants.

variance) repeated measures ANOVA carried out over participants' weights in the high likelihood variance condition, there was a main effect of temporal bin, $F(3, 21) = 5.55$, $p < 0.001$, and a main effect of prior variance, $F(1, 7) = 7.89$, $p = 0.026$. These findings therefore represent a replication of our results from Experiment 1. To further confirm that these results replicate the results from Experiment 1, we carried out a 2 (experiment) $\times$ 4 (temporal bin) $\times$ 3 (likelihood variance) mixed ANOVA, with experiment as a between-participants factor and temporal bin and likelihood variance as within-participant factors, over participants' weights in both the broad and narrow prior conditions. In each case, we found no interaction between experiment and any other factor (all $ps > 0.17$).

Importantly, in the conditions in which no sensory information was available (i.e., no cluster of dots was presented), we found that within the first temporal bin (the first 100 of the 400 prior-only trials) participants' estimates of the target location were on average indistinguishable from the true prior mean locations (plus or minus motor noise estimated from the low likelihood variance conditions; broad prior: $t_7 = 1.73$, $p = 0.13$; narrow prior: $t_7 = 1.24$, $p = 0.25$). Participants were therefore able to rapidly learn the prior means even when presented with a complex generative model (Figure 5), extending results from previous work examining such behavior in the presence of simpler generative models. It is important to note, however, that in the conditions where sensory information was available, which were randomly interleaved with the prior-only conditions, participants' weights continue to change throughout the experiment (Figure 4). This

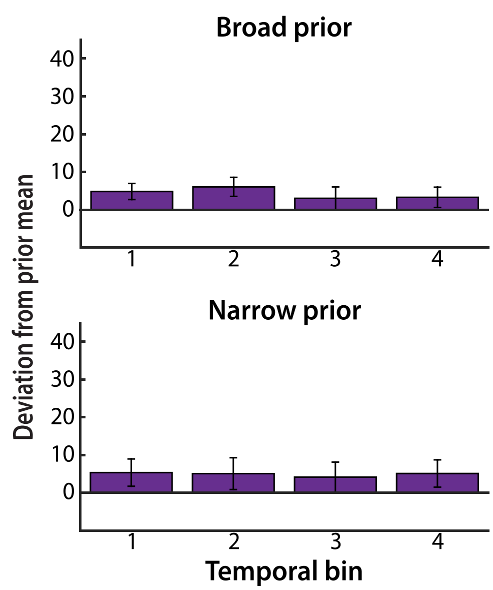

Figure 5. Mean response location in the absence of sensory information. Participants in Experiment 2 rapidly learned the true prior mean for both the broad (top) and narrow prior conditions (bottom). Participants' mean response location in the trials in which no sensory information was available was indistinguishable from the true prior means, plus or minus motor noise, within the first temporal bin. For illustrative purposes, the y-axis represents the mean deviation from the true mean in each prior condition. Each temporal bin included 100 trials split between the two prior conditions, and the four bins of trials are depicted in temporal order. Columns represent means and error bars represent *SEM* across participants.

pattern of results suggests that, while participants learn the prior means very rapidly, it takes much more exposure to learn the relative variances of the two prior distributions. Moreover, the finding that participants learn the true prior means within the first temporal bin, but their weights continue to change throughout the experiment, further contradicts the predictions of Model 2. If participants were only learning and using the prior means, and not the relative variances of the prior distributions (as predicted by Model 2), then we would expect to see no change in participants' weights beyond the first temporal bin, since they show no change in their knowledge of the prior means beyond this bin (as determined by performance in the no-sensory information condition).

Taken together, the results from Experiments 1 and 2 provide compelling evidence in support of the hypothesis that observers in our study learn and utilize complex generative models, in combination with sensory information, in a manner that is consistent with the predictions of Bayes-optimal behavior (i.e., Model 3).

General discussion

Across two experiments, we used a spatial localization task to examine the computational mechanisms employed by human observers when faced with tasks involving complex generative models. We obtained compelling evidence in support of the hypothesis that human observers learn such complex generative models, and integrate this learned knowledge with uncertain sensory information in a manner consistent with the predictions of Bayes-optimal behavior. At the outset of our experiments, participants' behavior was primarily driven by sensory information (i.e., the centroid of the cluster of dots) and we saw no difference in the weights assigned to sensory information across likelihood (i.e., differences in cluster variance) and prior conditions. However, by the end of both experiments (i.e., the final temporal bin), when participants had gained extensive exposure to the complex generative model, they assigned a significantly lower weight to the sensory information (and thus a greater weight to the mean of the trial-specific prior distribution) as the sensory information decreased in reliability, in a manner consistent with the predictions of the Bayesian model. This pattern of performance rules out alternative suboptimal strategies that could be used in this task. Furthermore, extending prior findings with simple generative models (Berniker et al., 2010), we found that even when faced with a complex generative model, participants were able to rapidly learn the true prior means based on just a few hundred trials of exposure, but it took them much longer to learn and use the relative variances of the prior distributions. To our knowledge, this is the first study that has explicitly examined the extent to which human observers behave in a manner consistent with the predictions of the Bayesian model in tasks that have the degree of complexity approaching that encountered in their natural environment; that is, tasks involving multiple, randomly interleaved levels of sensory uncertainty and where the stimulus of interest is drawn from a mixture distribution over multiple, randomly interleaved environmental causes.

In addition to demonstrating that our participants' behavior is consistent with the predictions of the Bayesian model, we can also evaluate the extent to which participants' empirical weights approach the ideal weights that would be assigned by a Bayes-optimal observer, given the uncertainty implicit in the two sources of information available to the participant: the noisy sensory information and the learned knowledge of the location-contingent generative model. How close are our participants' weights to this quantitative ideal? In order to accurately compute these ideal weights, however, we would need to characterize the uncertainty implicit in participants' *internal* estimates of the prior and likelihood distributions, thereby ensuring that the ideal observer has access to the same quality of information as our participants. Since participants were presented with multiple samples from the likelihood distribution (i.e., the cloud of dots) on every trial, the uncertainty implicit in it is computable by the participant on a trial-by-trial basis (see Sato & Kording, 2014, for a scenario in which this was not true). It is therefore reasonable to approximate participants' internal estimate of likelihood uncertainty by the true uncertainty implicit in the distributions used to generate the cloud of dots on each trial (i.e., the standard deviation of the dot-cloud distribution divided by the square root of the number of dots).

Approximating the uncertainty implicit in participant's internal estimates of the prior distributions is a much more challenging endeavor. Participants' internal estimates of the prior distributions could only be obtained by integrating across trials, since feedback on each trial only provided a single sample from the trial-specific prior distribution. Thus, we would expect participants' internal estimates of the true prior distributions to evolve over the course of the experiment, and only when participants had received enough samples from the underlying distributions to attain perfect knowledge of the prior distributions could we approximate their internal estimates of these distributions with the true prior distributions. The challenge stems from the fact that we had no independent way of assaying participants' internal estimates of the prior distributions, and thereby determining at what point in the experiment a given participant had received enough samples to attain such perfect knowledge.

Despite the foregoing limitation, we can compute the ideal weights that a Bayes-optimal observer, who is assumed to have perfect knowledge of the prior distributions, might assign to the sensory information in Experiment 2, and evaluate the extent to which participants' weights approach this quantitative ideal. The ideal observer in this analysis is therefore assumed to have knowledge about the underlying target distributions that is an upper bound on the knowledge that participants could learn during the course of the experiment. Comparing the ideal weights computed in this manner to participants' empirical weights across prior and likelihood conditions, we found that their empirical weights moved closer to the ideal weights as a function of exposure to the task. Furthermore, by the final temporal bin, which represented the most exposure to the underlying generative model, participants' weights showed the same qualitative trends as the ideal observer, but they differed quantitatively from the ideal weights in a manner that interacted with both prior and likelihood variance. Specifically, while participants' weights matched the ideal weights in the low likelihood variance condition, they assigned a greater weight to the sensory information, in comparison to the ideal observer, in the medium and high likelihood variance conditions, with the difference increasing with an increase in likelihood variance (i.e., the change in participants' weights as a function of likelihood variance was shallower than that predicted by the ideal observer; see Tassinari et al., 2006, for a similar finding). Indeed, this difference was also greater for the narrow prior condition, than for the broad prior condition (Supplementary Fig. S1). This pattern of results suggests that, even by the end of Experiment 2, participants may not have received enough exposure to samples from the underlying target distributions to attain the perfect knowledge of the prior distributions assumed by the ideal observer in this analysis. Indeed, if participants began the experiment with flat internal estimates for the priors (which is reasonable given that they did not have a priori reason to prefer one location on the display to another), and sharpened these internal estimates as a result of exposure, we would expect to see exactly this pattern of results. Consistent with this hypothesis, in a follow-up experiment (see supplementary information), we found that when participants were provided with double the number of samples from each prior distribution, their weights moved significantly closer to the ideal weights (Supplementary Fig. S2). Taken together, these results suggest that, when faced with tasks involving such complex generative models, the rate at which human observers learn the generative model is dramatically slower than that observed previously with simple generative models. This slowdown is most likely due to the significantly increased complexity involved in tracking and learning the complex generative model.

Going forward, a logical next step to further understanding of the computational mechanisms used by human observers to learn from and use complex generative models, would be to introduce uncertainty in the environmental causes that make up the mixture components of the generative model. In the current study, as a first step, we made the environmental distribution perfectly location contingent (i.e., predictable), with one Gaussian distribution being centered on the left of the display and the other being centered on the right of the display. As a next step, including uncertainty with regard to the underlying mixture component that the target on each trial is drawn from (i.e., removing the location contingency) would introduce further complexity in the generative model and make the task even more representative of real-world tasks. Ideal behavior in such tasks can still be modeled in the Bayesian framework by considering a hierarchical process, where the ideal observer first solves the "causal inference" problem (what underlying cause produced the sample; Kayser & Shams, 2015; Knill, 2007; Rohe & Noppeney, 2015; Shams & Beierholm, 2010) and then uses the estimated sample to (a) build a model of the underlying distribution for the inferred cause, and (b) integrate that inferred causal distribution with sensory information on subsequent trials to finally infer the target location.

Keywords: learning complex models, Bayesian modeling, spatial localization

Acknowledgments

This work was supported by research grants from ONR (MURI-N00014-07-1-0937; Daphne Bavelier, PI) and NIH (HD-037082; Richard Aslin, PI). We are grateful to Holly Palmeri for help with recruiting and running participants.

Our dear friend and collaborator David Knill died tragically while this research project was in progress. We hope our conceptual approach, analytic methods, and write-up of the results lives up to his high standards. He will be sorely missed.

Commercial relationships: none.
Corresponding author: Vikranth R. Bejjanki.
Email: bejjanki@princeton.edu.
Address: Department of Psychology & Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA.

References

Adams, W. J., Graf, E. W., & Ernst, M. O. (2004). Experience can change the "light-from-above" prior. *Nature Neuroscience*, *7*, 1057–1058, doi:10.1038/nn1312.

Bejjanki, V. R., Clayards, M., Knill, D. C., & Aslin, R. N. (2011). Cue integration in categorical tasks: Insights from audio-visual speech perception. *PLoS ONE*, *6*(5), e19812, doi:10.1371/journal.pone.0019812.

Bernardo, J. M., & Smith, A. F. M. (1994). *Bayesian theory*. New York: Wiley.

Berniker, M., Voss, M., & Kording, K. (2010). Learning priors for Bayesian computations in the nervous system. *PLoS ONE*, *5*(9), e12686, doi:10.1371/journal.pone.0012686.

Cox, R. T. (1946). Probability, frequency and reasonable expectation. *American Journal of Physics*, *14*(1), 1–13, doi:10.1119/1.1990764.

Hudson, T. E., Maloney, L. T., & Landy, M. S. (2007). Movement planning with probabilistic target information. *Journal of Neurophysiology*, *98*, 3034–3046, doi:10.1152/jn.00858.2007.

Huttenlocher, J., Hedges, L. V., & Vevea, J. L. (2000). Why do categories affect stimulus judgment? *Journal of Experimental Psychology: General*, *129*, 220–241, doi:10.1037/0096-3445.129.2.220.

Jacobs, R. A. (1999). Optimal integration of texture and motion cues to depth. *Vision Research*, *39*, 3621–3629.

Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, *13*, 1020–1026, doi:10.1038/nn.2590.

Kayser, C., & Shams, L. (2015). Multisensory causal inference in the brain. *PLoS Biology*, *13*(2), e1002075, doi:10.1371/journal.pbio.1002075.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*, 712–719.

Knill, D. C. (2007). Robust cue integration: A Bayesian model and evidence from cue-conflict studies with stereoscopic and figure cues to slant. *Journal of Vision*, *7*(7):5, 1–24, doi:10.1167/7.7.5. [PubMed] [Article]

Körding, K. P., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., & Shams, L. (2007). Causal inference in multisensory perception. *PLoS ONE,* 2(9), e943, doi:10.1371/journal.pone.0000943.

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

Kwon, O.-S., & Knill, D. C. (2013). The brain uses adaptive internal models of scene statistics for sensorimotor estimation and planning. *Proceedings of the National Academy of Sciences, USA*, *110*(11), E1064–E1073, doi:10.1073/pnas.1214869110.

Miyazaki, M., Nozaki, D., & Nakajima, Y. (2005). Testing Bayesian models of human coincidence timing. *Journal of Neurophysiology*, *94*, 395–399, doi:10.1152/jn.01168.2004.

O'Reilly, J. X., Jbabdi, S., Rushworth, M. F. S., & Behrens, T. E. J. (2013). Brain systems for probabilistic and dynamic prediction: Computational specificity and integration. *PLoS Biology*, *11*(9), e1001662, doi:10.1371/journal.pbio.1001662.

Rohe, T., & Noppeney, U. (2015). Cortical hierarchies perform Bayesian causal inference in multisensory perception. *PLoS Biology*, *13*(2), e1002073, doi:10.1371/journal.pbio.1002073.

Sato, Y., & Kording, K. P. (2014). How much to trust the senses: Likelihood learning. *Journal of Vision*, *14*(13):13, 1–13, doi:10.1167/14.13.13. [PubMed] [Article]

Shams, L., & Beierholm, U. R. (2010). Causal inference in perception. *Trends in Cognitive Sciences*, *14*, 425–432.

Stocker, A. A. & Simoncelli, E. P. (2006). Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, *9*, 578–585.

Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *The Journal of Neuroscience*, *26*, 10154–10163, doi:10.1523/JNEUROSCI.2779-06.2006.

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*, 1641–1648, doi:10.1016/j.cub.2012.07.010.

Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, *5*, 598–604.

Yuille, A. L., & Bülthoff, H. H. (1996). Bayesian decision theory and psychophysics. In D. C. Knill W. Richards (Eds.), *Perception as Bayesian inference* (pp. 123–161). Cambridge, UK: Cambridge University Press.