

# How much to trust the senses: Likelihood learning

Yoshiyuki Sato

Graduate School of Information Systems,  
University of Electro-Communications, Japan



Departments of Physical Medicine and Rehabilitation,  
Physiology, and Applied Mathematics,  
Northwestern University,  
and Rehabilitation Institute of Chicago, IL, USA



Konrad P. Kording

**Our brain often needs to estimate unknown variables from imperfect information. Our knowledge about the statistical distributions of quantities in our environment (called priors) and currently available information from sensory inputs (called likelihood) are the basis of all Bayesian models of perception and action. While we know that priors are learned, most studies of prior-likelihood integration simply assume that subjects know about the likelihood. However, as the quality of sensory inputs change over time, we also need to learn about new likelihoods. Here, we show that human subjects readily learn the distribution of visual cues (likelihood function) in a way that can be predicted by models of statistically optimal learning. Using a likelihood that depended on color context, we found that a learned likelihood generalized to new priors. Thus, we conclude that subjects learn about likelihood.**

## Introduction

To make accurate estimates regarding relevant variables in our environment, we need to combine sensory evidence with our prior knowledge. For example, if we try to catch a fish, the fish location may be blurred because of waves in the water. If we know that the fish is likely to stay close to rocks, we can increase our chance of catching the fish by aiming somewhere near a rock in addition to relying on the blurred fish image. In the fish-catching task, the location of the fish is a quantity that is estimated using the blurred visual input. Bayesian inference enables us to optimally combine observations with prior knowledge. Bayesian models of the brain have been studied extensively in many fields of neuroscience, such as vision (Kersten, Mamassian, & Yuille, 2004), multisensory perception (Trommershauser, Kording, & Landy, 2011), and sensorimotor integration (Berniker & Kording,

2011). How the brain accomplishes this statistical computation is an important question in neuroscience.

A Bayesian model of the fish-catching task can be formulated as follows. Let us denote the fish location as  $x$  and the blurred visual input as  $y$ . Using Bayes' theorem, the posterior probability distribution of  $x$  after observing  $y$  can be written as  $P(x|y) \propto P(y|x)P(x)$ . Here,  $P(x)$  characterizes our prior knowledge about the statistics of the task-relevant quantity  $x$ , and represents information about  $x$  before observing  $y$ . Another factor  $P(y|x)$  (called the likelihood function) represents how likely each  $x$  causes  $y$ ; essentially, it represents the information about  $x$  that is obtained from observation. By combining the prior and the likelihood, we can optimally decide how much we should trust observations versus prior knowledge. Note that once the target variable and the observed variable are defined, we can clearly define the likelihood and the prior for the task. These functions are not arbitrarily set, as is often seen in applied statistics literature.

A likelihood function is derived from the conditional probability  $P(y|x)$ , which formulates how an observed quantity can deviate from the true value. In the example above, waves and streams blur the visual image and thus impose uncertainty on the visual information, resulting in a wider likelihood function. Other factors, like the refraction of light in the water, could cause a shift in the visual image. In this paper, for clarity and simplicity, we focus on the aspect of likelihood as representing uncertainty in sensory input. Highly uncertain sensory input indicates that we cannot place a high degree of trust in observation, thus likelihood also represents sensory reliability.

Our nervous systems adapt to changes in the world, and a central problem in behavioral neuroscience is the dynamics of this adaptation process. Many studies have shown that priors are learned from past experience (Berniker, Voss, & Kording, 2010; Jazayeri &

Citation: Sato, Y., & Kording, K. P. (2014). How much to trust the senses: Likelihood learning. *Journal of Vision*, 14(13):13, 1–13, <http://www.journalofvision.org/content/14/13/13>, doi:10.1167/14.13.13.

Shadlen, 2010; Körding & Wolpert, 2004; Miyazaki, Nozaki, & Nakajima, 2005). In those studies, it was assumed that the subjects knew the likelihood. It is known that sensory reliability (likelihood) plays an important role in the adaptation of sensory (Burge, Girshick, & Banks, 2010) and motor systems (Burge, Ernst, & Banks, 2008; Wei & Körding, 2010). The quality of sensory information changes over time because of injury, aging, and other internal and external factors. Our brain thus has to constantly evaluate the reliability of our sensory information. However, there is little experimental data that explains how likelihood itself is learned.

There are two factors that contribute to the interpretation of incoming sensory data. The first is the physical precision of our senses, which can be affected by many factors, such as retinal properties or noise in neuronal pathways. This precision might change, e.g., with light levels. Second, even precise sensory abilities might not lead to a complete set of information about a relevant variable of interest. In the fish catching task, visual information about the fish's location is blurred and uncertain even if one's retinal image is highly precise. The reliability of the visual image depends on factors in the environment, such as wave height, flow speed, and the contours of the water surface. In this and many other situations, we have to learn about the reliability of visual information. Although the first factor is usually referred to as "sensory uncertainty," both of these two factors, precision of sensory systems and reliability of sensory input define the likelihood function of the relevant quantity about the object. In this paper, we focus on the latter factor. While sensory precision itself varies over time, the reliability of sensory input can vary widely across tasks, and the brain has to efficiently address such variations.

Here, we used a sensorimotor task in which we manipulated the uncertainty represented in the likelihood to examine whether and how human subjects learned about the likelihood. First, we found that the subjects readily learned the likelihood, especially when they were told the likelihood would change. We compared the experimental data to three models and found that the optimal Bayesian learning model best described the observed behavior. Next, we confirmed that the subjects actually learned the likelihood, and not just the relative weight put on observation, by showing that the subjects learned multiple color-cued likelihoods and generalized the learned likelihoods to a new prior.

## Related research

Perhaps most similar in spirit to this study are recent studies of cue combination. Some of these reports

found that subjects can estimate the right cue weights directly without learning (Alais & Burr, 2004; Ernst & Banks, 2002). In those studies, the reliability of each cue is embedded (thus can be estimated) in the stimulus itself, and they are combined optimally. In other studies, experience changed the weights given to cues (Ernst, Banks, & Bühlhoff, 2000; Jacobs & Fine, 1999; Seydell, Knill, & Trommershäuser, 2010; Van Beers, van Mierlo, Smeets, & Brenner, 2011). In those studies, the reliability of each cue was controlled externally by variations of procedures, and the subjects adaptively updated the weights. As we will describe in the Methods section, in our experiment, the reliability of the observation is not embedded in the stimulus but is controlled over trials, thus our paradigm is similar to the latter kind of experiments. Sensory representations can also shift after repeated exposure to conflicted stimuli (Wozny & Shams, 2011). The observed time-scale of learning in one report (van Beers et al., 2011) was similar to that we report here. However, cue combination and prior-likelihood integration experiments are quite different. Cue combination refers to the instantaneous integration of observed multiple inputs while prior-likelihood integration is the integration of a current observation with prior knowledge stored in the memory. These two types of experimental paradigms provide different insights into the information processing mechanisms in the brain.

While these experiments have revealed important implicit insights into likelihood learning, they deal with a different kind of situation and leave a lot of important questions open which we address here.

Implicit signs of cue reliability learning have also been found in sensorimotor integration experiments. In one study, subjects were required to estimate the random-walking discrepancy between their hand and a cursor from noisy observation of the cursor (Baddeley, Ingram, & Miall, 2003). Their relatively constant performance among difference noise levels could be attributed to learning of the noise parameter. In another study, time-dependent variability was imposed to the visual feedback of a target in a fast-reaching task, and the subjects optimally combined the cue reliability and their motor variability to determine the best timing to start moving (Battaglia & Schrater, 2007). In their analysis, additional internal noise parameters were fitted to behavioral data while assuming the true noise parameter was known to subjects, thus it is not directly clear whether and how well the subjects estimated the noise parameter. Those studies did not show how the noise parameters, and thus likelihood, were learned. While the learning of likelihood has been implied by these studies, it is important to precisely understand its learning dynamics.

Several theoretical studies have also suggested that some types of sensory adaptation can be interpreted as the learning of likelihood (Sato & Aihara, 2011; Sato, Toyozumi, & Aihara, 2007; Stocker & Simoncelli, 2006). However, there is little direct evidence that human subjects actually learn the quality of sensory input (likelihood) over trials and combine it with the prior, which is an essential factor in the Bayesian paradigm that enables us to flexibly make use of learned information about the world.

Here we investigate likelihood learning in the context of prior/likelihood integration using a radically simplified setting. Using switches in likelihood width, we characterize the learning process in great detail. We analyze the role of instructions for the subjects. We also use generalization tests to provide solid evidence for the learning of likelihoods.

## Materials and methods

All experimental protocols were approved by the Northwestern University Institutional Review Board, were in accordance with Northwestern University's policy statement on the use of humans in experiments, and conformed to the Declaration of Helsinki. All participants were naive to the goals of the experiment, provided signed consent forms, and received monetary compensation for their participation.

### Experiment 1

In this experiment, we investigated whether and how subjects learned about likelihood by manipulating the width of a likelihood function several times during an experiment. We used a coin location estimation task (Berniker et al., 2010; Tassinari, Hudson, & Landy, 2006). Slow prior learning has been reported in a similar task setting (Berniker et al., 2010), which can obstruct our main focus of likelihood learning. Thus, in our experiments, we visually displayed a Gaussian prior representing the distribution of target coins.

Sixteen subjects (12 males and four females, aged 20–37 years) participated in Experiment 1. The subjects were seated 50 cm from a vertical computer monitor (13.3-in. diagonally). They were instructed to view the screen as a surface of water and to locate an unseen coin that a person had dropped into the water. We used an arbitrary unit to define the horizontal location of the left and right edges of the screen as  $-0.5$  and  $0.5$  respectively, and all length values given in this paper relate to these coordinates. We used a program in Matlab to control the stimulus presentation and record the subject responses. At the beginning of each trial, a

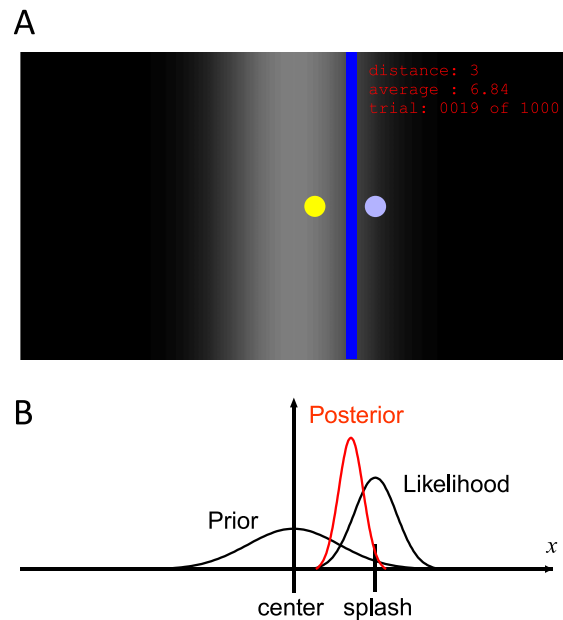


Figure 1. Overview of the task and its simple Bayesian model. (A) An example of the screen. Subjects were asked to locate an unseen target coin (yellow circle) after observing a splash (blue circle) caused by the coin. The coin location was displayed as feedback after the subjects made a response. Subjects moved the blue line horizontally and estimated the location of the coin. The graded background shows the prior over the coin position. Circles and texts are enlarged for better visibility. (B) Bayesian model of the task. The estimated coin location is the mean of posterior obtained by combining the prior at the center of the screen and the likelihood centered at the splash.

blue circle (diameter = 0.006), which represented a splash caused by a coin dropped into the water, appeared on a graded monochrome background, which represented the prior probability distribution of the coin (Figure 1A). The subjects used a computer mouse to move a vertical blue bar (width = 0.01) horizontally on the screen, and clicked the button on the mouse to indicate their estimation of the coin location. After the subjects made a response, the location of the coin was presented as a yellow circle (diameter = 0.006) on the screen for 1 s. The absolute distance ( $\times 100$ ) between the subject's response and the coin location in the last trial was displayed on the right upper corner of the screen as a score, in addition to the average score for all completed trials and the current trial number. The subjects were instructed to keep the average distance as low as possible. All coins and splashes were displayed on the vertical center of the screen, thus only their horizontal locations were important.

The location of the coin was drawn randomly from a Gaussian distribution with a mean of 0 and a standard deviation of 0.1. This probability distribution defined the prior of the coin and was explicitly shown to the subjects. The location of the splash was drawn from a

Gaussian distribution centered at the coin, and this distribution defined the likelihood of the coin. Its standard deviation was either 0.05 (narrow condition) or 0.2 (wide condition) and the condition switched every 70–100 trials. The switching intervals were sampled from a uniform distribution within that range and the initial condition was chosen randomly. Each subject performed 1,000 trials, resulting in 40–60 min of participation, depending on the subject's response speed. The subjects were permitted to take a break every three switches (approximately every 10 min).

The subjects were told that the height at which the coins had been dropped determined the spread of the splashes, i.e., the width of the likelihood. The subjects were divided into two groups: those in the with-instruction group ( $N = 9$ ) received information about how the height was determined, while those in the without-instruction group ( $N = 7$ ) did not.

The task instructions were as follows: “Assume that the screen is a surface of water. In each trial, it will appear as if someone has dropped a coin from above, and you will not see the coin nor the person but only one splash caused by the coin. Coins and splashes will always be displayed on the vertical middle line, so the horizontal location is important. That person determines the horizontal location of the coins according to the graded background randomly from trial to trial. The person tends to drop the coins from brighter locations. The width of the splash distribution is determined by the height at which the person drops the coin. The person is not seeing your responses, and your responses do not affect their behavior. Your task is to move the blue line to estimate the location of the coin. Keep the average distance between your response and the coin location as low as possible.” For those in the with-instructions group, we added additional instructions: “The person dropping the coins chooses the height from two options, high and low, and sometimes switches between these options. The person drops the coins from the same height for maybe tens of trials, and then switches to the other height.” For those in the without-instructions group, we either said, “I will tell you nothing about how the person determines the height,” or said nothing about how the person dropping the coins determines height.

## Experiment 2

Experiment 2 was similar to Experiment 1 except that multiple likelihoods of coin location were cued by the splash color and learned simultaneously. The prior variance of coin location was changed during the experiment to test whether subjects could generalize a learned likelihood to a new prior.

Fifteen subjects (10 males and five females, aged 20–37 years) participated in Experiment 2. We excluded one subject whose score was greater than two times the standard deviation of the mean score across all subjects and who, upon postexperiment questioning, had not understood the instructions. We used Matlab with the Psychtoolbox-3 extension (Brainard, 1997; Kleiner, Brainard, & Pelli, 2007; Pelli, 1997) to generate the stimuli and record the subject responses. The experiment lasted 40–60 min depending on the subject's response speed. The subjects were encouraged to take a break every 250 trials.

The prior over the coin location was a Gaussian distribution with a mean of 0 (center of the screen) and a standard deviation of either 0.06 or 0.18, as described below. The prior was always represented by the density of white one-pixel dots on black in the background of the screen, rather than as intensity. The probability of the background dots appearing in any horizontal location was linear to that of the prior at that location, and normalized to a maximum of 0.95 to prevent an impression of probability saturation. Each splash was represented by a green or red circle (diameter = 0.01) with a horizontal location drawn from a Gaussian distribution centered at the coin location with a standard deviation of either 0.06 or 0.18, depending on the color. The correspondence between the colors and the standard deviations was counter-balanced across subjects.

Each subject performed 750 trials divided into three phases: initial likelihood learning (trials 1–400), new prior learning (trials 401–550), and generalization test (trials 551–750). During the initial likelihood learning phase, the standard deviation of the coin prior was either 0.06 or 0.18 (counterbalanced across subjects). The color of each splash was randomly chosen every trial from green and red with equal probabilities. After the subject made a response, the location of the coin was presented as a yellow circle (diameter = 0.01) for 1 s. In subsequent phases, the prior standard deviation was switched to the other value. In the new prior learning phase, subjects saw only one splash color and thus only one likelihood condition (counter-balanced across subjects). Feedback about the actual coin location was given as in Experiment 1. In the generalization test phase, both splash colors were used with equal probabilities but subjects did not receive feedback about coin locations and scores. Thus, in the new prior learning phase, subjects learned the new prior in only one context (color) of likelihood. The generalization of this likelihood to the other context was tested in the generalization test phase.

In addition to the instructions received by both groups in Experiment 1, subjects were instructed as follows: “In each trial, one of two people will drop a coin. The color of the splash specifies which person

dropped the coin. The two people will drop the coins from different heights and they will each maintain the same height throughout the experiment. The density of the dots in the background shows the overall distribution of coin location. The people tend to drop the coins to locations with more dots, and they randomly choose coin locations from trial to trial. Both people produce the same overall coin distribution. This distribution can change, but the height at which the coins are dropped does not change. At some point in the session, the location of the coins will no longer be displayed, but your score will be calculated in the background. So try to minimize the average distance.”

## Models of human behavior

There are several possible ways the brain could deal with changes in likelihood width. (a) No adaptive learning. The brain does not adaptively estimate the reliability of the sensory information following a likelihood change. Instead, it utilizes a static strategy to estimate the coin location. (b) Adaptive but suboptimal learning. The brain simply learns the likelihood of the coin location from recent experience. This simple computation is easy to implement and less computationally demanding, but is statistically suboptimal. (c) Bayesian inference with likelihood learning. The brain adaptively learns the likelihood and maintains optimal inferences accordingly. We tested these hypotheses to elucidate the mechanisms underlying likelihood learning and assess the applicability of Bayesian models in explaining estimation behaviors.

We compared three models of human behavior that correspond to the three possibilities above: static nonlinear mapping from splash to coin, computing the slope from recent trials, and a fully Bayesian optimal model.

### Static mapping model

In this model, in a given trial, the mapping from the splash location to the coin location is learned by fitting a continuous piece-wise linear function to the observed data up to the trial in a least squared error manner. We divided the screen into 10 sections horizontally, each separated by 11 points (two edges and nine intermediate points) and fitted the mapping in each piece with a linear function

$$f_i(y) = b_i + c_i(y - a_i) \quad (i = 1, \dots, 10),$$

where  $y$  is the splash location,  $f_i$  is the mapped (estimated) coin location when the splash location is in the  $i$ -th piece,  $a_i$  is the left edge of the  $i$ -th piece, and  $b_i$  and  $c_i$  are parameters.

We can remove  $c_i$  from the constraint of continuity  $f_i(a_{i+1}) = b_{i+1}$ . We modeled learning as the act of finding the remaining  $b_1, \dots, b_{11}$  values by minimizing the cost function

$$\sum_{i=1}^{10} \sum_{j=1}^{m_i} (f_{i,j} - x_{i,j})^2 + \frac{\lambda}{2} \sum_{i=1}^{10} (b_i - b_{i+1})^2,$$

where  $m_i$  is the number of data points in the  $i$ -th piece and  $f_{i,j}$  and  $x_{i,j}$  are the estimated coin location and the actual coin location corresponding to  $j$ -th splash in the  $i$ -th piece. We added the second normalization term for a smooth mapping function to avoid initial instability. After each trial, the model learns the piecewise linear function up to the trial. The only free parameter of this mapping model is the coefficient  $\lambda$  of the normalization factor.

### Recent slope learning model

In this model, in every trial, a linear regressed line is calculated from recent  $n$  trials by linearly regressing feedback about coin locations from the splash locations. In each trial, the observer simply predicts the coin location based on the observed splash location using the regressed line. There is only free parameter,  $n$ .

### Bayesian learning model

In this model, previous observations affect the current estimation via estimation of the likelihood width. A generative model of the task is depicted in Figure 5A, where  $h^t$  is one of the two possible heights (1 or 2),  $\sigma_1$  and  $\sigma_2$  are the likelihood widths for each possible height, and  $\sigma_t^i$  is the likelihood width in trial  $t$ . The height at the beginning of an experiment is defined as  $h^1 = 1$ . In trial  $t$ , the optimal observer’s task after observing  $y^t$  is to compute the mean value of the posterior probability distribution  $P(x^t | x^{1:t-1}, y^{1:t})$ , where  $y^{1:t}$  denotes all  $y$ s from time 1 up to time  $t$ . Assuming the probabilistic structure of the generative model of the task, after some calculation, this posterior can be written as

$$\begin{aligned} P(x^t | x^{1:t-1}, y^{1:t}) \\ \propto \sum_{h^t=1,2} \int P(y^t | x^t, \sigma_t^i = \sigma_{h^t}) \\ P(x^t) P(h^t, \sigma_{h^t} | x^{1:t-1}, y^{1:t-1}) d\sigma_{h^t}. \end{aligned}$$

We choose the estimated value of  $x^t$ ,  $\hat{x}^t$ , by calculating the mean value of this posterior. To compute this posterior of  $x^t$ , the observer needs to predict  $h^t$ ,  $\sigma_1$ , and  $\sigma_2$  (the last term). After observing  $x^t$  and  $y^t$ , the prediction can be updated as

$$\begin{aligned}
 &P(h^{t+1}, \sigma_1, \sigma_2 | x^{1:t}, y^{1:t}) \\
 &\propto \sum_{h^t=1,2} P(h^{t+1} | h^t) P(y^t | x^t, \sigma_i^t = \sigma_{h^t}) \\
 &P(h^t, \sigma_1, \sigma_2 | x^{1:t-1}, y^{1:t-1}).
 \end{aligned}$$

this updated prediction can be used to estimate  $x^{t+1}$  in the next trial. The initial priors of  $\sigma_1$  and  $\sigma_2$  are assumed to be inversely proportional to their values (Jeffreys priors). The likelihood functions are assumed to be Gaussian.

The model cannot know when switches will take place. As such, it needs to consider all possibilities for when these switches could occur. In this model, for simplicity, the transition probability  $P(h^{t+1} | h^t)$  is assumed to be constant over time:  $P(h^{t+1} | h^t) = 1 - \alpha$  if  $h^{t+1} = h^t$ , and  $\alpha$  otherwise. Here,  $\alpha$  is the switching probability and is the only free parameter in this learning model. In the implementation of the model, the likelihood function is estimated at every time step and updated, taking into account both the possibility of there being no switch and the possibility of a switch, each weighed by their respective probability.

Numerically, this model is implemented by discretizing the probability distributions. Distributions over  $x$  were discretized in the range  $[-1.5, 1.5]$  with steps of 0.01 (301 points), and distributions over  $\sigma$ s were discretized in the range  $[e^{-5}, e]$  with steps of 0.01 in power (601 points). We approximated the integrals by summing up the discretized probability distributions with appropriate weights and normalizing them afterwards.

## Results

We characterized likelihood learning in a coin location estimation task in which subjects estimated the position of a hidden coin from a cue “splash” (Figure 1A). The splash was drawn from a Gaussian distribution centered on the target, and by changing the variance of this distribution, we were able to vary how informative the splash was (and thus change the width of the likelihood function). In the first experiment, the likelihood function changed over time while the prior stayed constant. In the second experiment, the color of the cue splash indicated the width of the likelihood function and the prior was switched during the experiment to probe generalization. For each trial, we used the coin location estimate to infer how much the subjects trusted the prior versus the likelihood. This relative reliance then enabled us to indirectly characterize the subject’s estimate of the likelihood in each trial.

## Measuring the brain’s estimation of the likelihood

To understand how the likelihood is learned, we need to have a way of measuring the brain’s estimation of the likelihood. Although there is no direct way to measure the estimated likelihood, it can be estimated indirectly. When the likelihood function is wide, observation cannot be trusted and subjects are more likely to achieve a higher score if they rely strongly on the prior. When the likelihood function is narrow, the opposite is true. A simple Bayesian model of prior-likelihood integration (Figure 1B) allowed us to calculate how strongly subjects should rely on the prior. This derivation relies on the assumption that the prior-likelihood combination is actually a good approximation of a Bayesian model, an issue that has been tested in a good number of previous studies (Körding & Wolpert, 2004; Miyazaki et al., 2005; Tassinari et al., 2006). The weight put on the visual cue position when estimating the target position can be expressed as

$$\hat{x} = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_l^2} y,$$

where  $\hat{x}$  is the best estimate,  $y$  is the cue position,  $\sigma_p^2$  and  $\sigma_l^2$  are the variance of the prior and the likelihood, respectively, and the prior mean is 0. The coefficient provides a measure of how strongly the observer relies on the observation and, assuming that the prior variance is already known, it reflects the observer’s estimation of the likelihood variance. If the coefficient changes when the likelihood width is changed while the prior variance is held constant, we can say that the subject changed their reliance on observation versus prior knowledge in response to the change in likelihood. We thus have a measurable quantity that we can use to track a subject’s estimate of likelihood width over time.

To experimentally determine how strongly the subject relies on the splash feedback when estimating the coin position, we plotted the cue positions versus the subject’s estimates over the trials and calculated the regressed slope (Figure 2A). A slope close to 0 would indicate that the subject’s estimated likelihood was wide, and a slope close to 1 would indicate a narrow estimated likelihood.

## Experiment 1: Dynamics of likelihood learning and its optimality

In Experiment 1 we investigated whether and how likelihood was learned by alternating the true likelihood width several times between two values. The prior was held constant throughout the experiment.

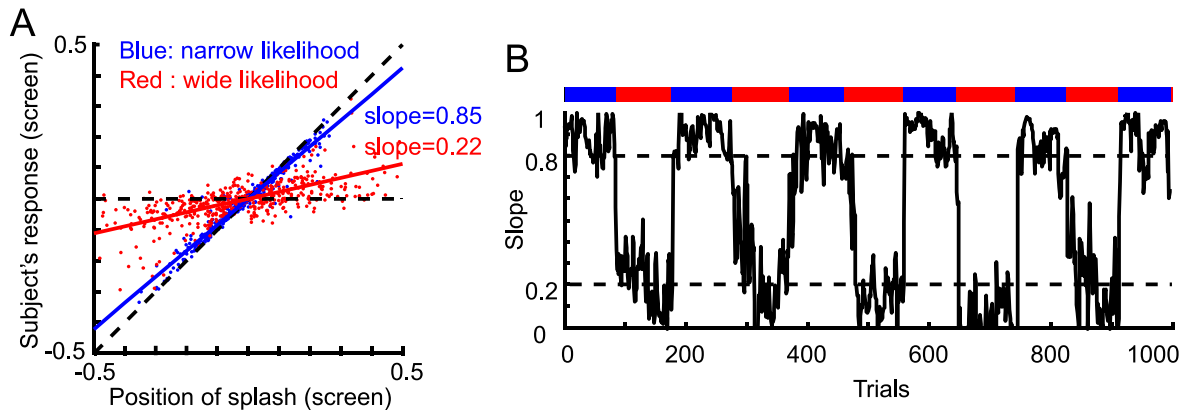


Figure 2. Data from a representative subject in the with-instruction group for Experiment 1. (A) The subject's responses in all trials as a function of the splash position. Steeper slopes indicate that the subject relied more strongly on the likelihood. We used the slope as a measure of the subject's estimate of likelihood width. For reference, the diagonal dashed line shows the prediction when the subject relied only on the cue position and ignored the prior, and the horizontal dashed line shows the prediction when the subject ignored the cue. (B) The time-course of the slope. The slope in each trial was calculated by regressing the following five trials. Dotted lines are the expected slopes (0.2 and 0.8 for wide and narrow conditions, respectively) from the simple Bayesian model if the subject learned the correct likelihood. The bar above represents the likelihood condition in each trial.

First, we checked whether the subjects had obtained the correct prior mean by calculating the  $y$ -intercept of the regressed line from the pooled data (See Figure 2A) for each subject. If the subject had a prior mean other than 0, we expected this to be reflected in the  $y$ -intercept value. We found a small bias in the narrow likelihood condition for the with-instructions group ( $t$  test across subjects,  $p = 0.01$ ). However, the value of this bias was much smaller ( $0.0025 \pm 0.0023$   $SD$ ) than the width of the vertical blue bar (0.01) used to indicate response making, thus the bias was only marginal. We found no significant  $y$ -intercept values for the other conditions (wide likelihood condition for the with-instructions group and both conditions for the without-instructions group). Therefore, we can reasonably assume that the subjects obtained a nearly correct prior mean.

The subjects changed how strongly they relied on the position of the cues based on the likelihood condition (Figure 2A). This indicates that the subjects learned the likelihood variance and adjusted their behavior appropriately. To characterize the dynamics of this likelihood learning, we quantified the changes in slope in response to changes in likelihood width. As the prior was held constant throughout the experiment, we can reasonably assume that all the learning we observed was due to the changes in the likelihood. A representative subject quickly learned the likelihood after each switch (Figure 2B). This quick learning was also apparent when averaged across subjects (Figure 3). The learning occurred roughly within the first 20 trials after each switch. We also found that the learning speed was faster when switching from the narrow to wide likelihood conditions. The experiment was designed such that a simple Bayesian estimator with the correct

likelihood and prior would use slopes of 0.8 and 0.2 for the narrow and wide conditions, respectively. However, the learned slope was significantly different from those values ( $t$  test for data from 30–70 trials after each switch,  $p < 10^{-5}$  for all conditions). For the narrow likelihood condition, both groups placed insufficient weight on vision. While we found some weak biases, subject data rapidly converged to near-optimal slopes. We will discuss these deviations later in the Discussion.

To test how likelihoods are learned, we compared our experimental results to a model-based analysis. With this approach, we were able to analyze the dynamics of human likelihood learning and ultimately compare experimental data with model predictions. One of the models to be compared was a Bayesian learning model, which was derived from the generative model of the task. To test the extent to which the assumed generative model was affected by instructions, one subject group was told that the likelihood would switch throughout the trials (“the person throws coins from two heights”) while the other group received no information about likelihood change. We found the instructions to have a clear influence on learning behavior (Figure 3). On average, the acquired slope for the two conditions was closer in the without-instructions group (two sample  $t$  test for data from 30–70 trials after switch,  $p < 10^{-5}$ ), indicating that less learning occurred in the condition without instructions about likelihood change. In the with-instructions group, all of the subjects showed a clear difference between the two conditions. However, in the without-instructions group, some subjects showed a large difference as in the with-instructions group, while other subjects showed a minimal difference (data not shown). Averaging the data from the two types of subjects in the

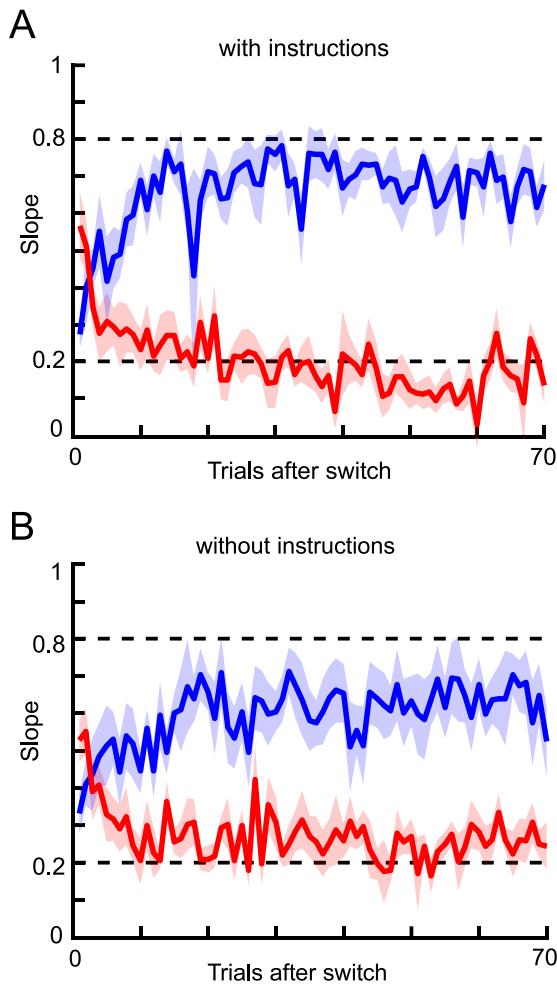


Figure 3. Mean slope after each switch. (A) Results for the subjects who received instructions regarding changes in likelihood width. Blue and red lines show the results for the narrow and wide conditions, respectively. The shaded area indicates the standard error across subjects. Each slope was calculated from pooled data for corresponding trials within subjects, and then averaged across subjects. (B) Results for the subjects who received no instructions regarding changes in likelihood width.

without-instructions group resulted in a weaker overall learning level for that group. Instructions thus appear to be important for this type of learning task.

How do subjects combine information from past trials into an estimate of the likelihood? We predicted that trials would influence subsequent trials in a way that decayed as a function of the number of intermediate trials. We found that increasing deviations between the cue splash and target coin were associated with decreased reliance on the visual splash in subsequent trials (Figure 4A). Moreover, this effect decreased as the number of intermediate trials increased (Figure 4B). We also observed that this likelihood learning was slower for the wide likelihood function (Figure 4B, blue vs. red). Thus, there are clear

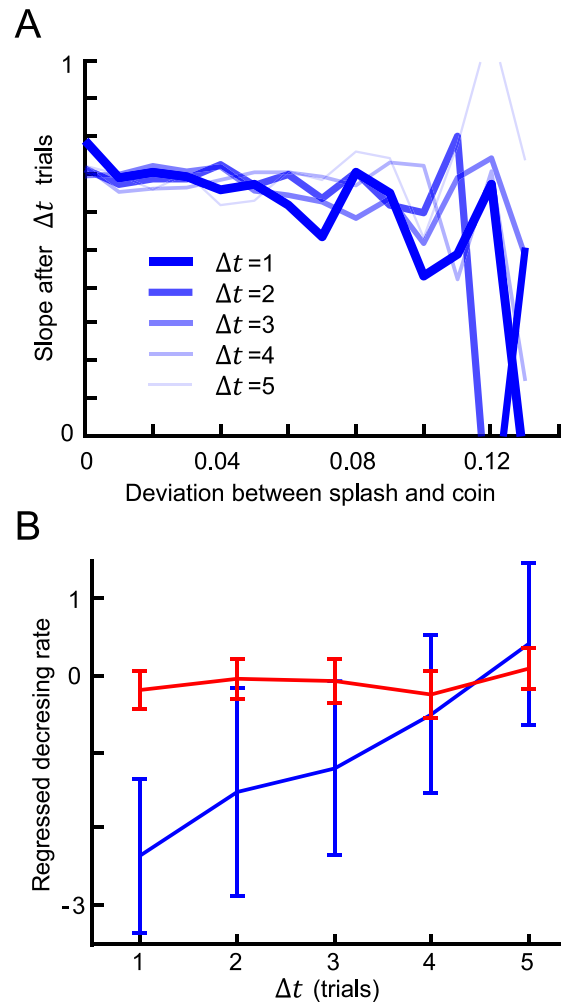


Figure 4. Effect of the deviation between splash and target coin on subsequent trials. (A) Results for the narrow condition averaged across subjects in the with-instructions group. The horizontal axis represents the deviation between a subject's response and the coin location in a trial. The vertical axis represents the slope calculated from the pool of trials  $\Delta t$  trials after the trials with the corresponding deviation. (B) Decreasing rate of the slope values in (A) for each  $\Delta t$ . The rate was obtained by regressing the slope values as a function of the deviation in (A). The error bars show the 95% confidence interval of the regressed decreasing rate. The blue and red lines show the results for the narrow and wide conditions, respectively.

signs of trial-by-trial likelihood learning over the time course of several subsequent trials.

Our results indicate signs of likelihood learning. Thus, we set out to determine which models could account for the observed learning effects. We compared three models that corresponded to our three hypotheses about likelihood learning (see Methods for mathematical details). (a) Static mapping: Because the overall distribution of the deviation is wider in the wide likelihood condition, a learned slope could differ between conditions if the subjects learned deviation-



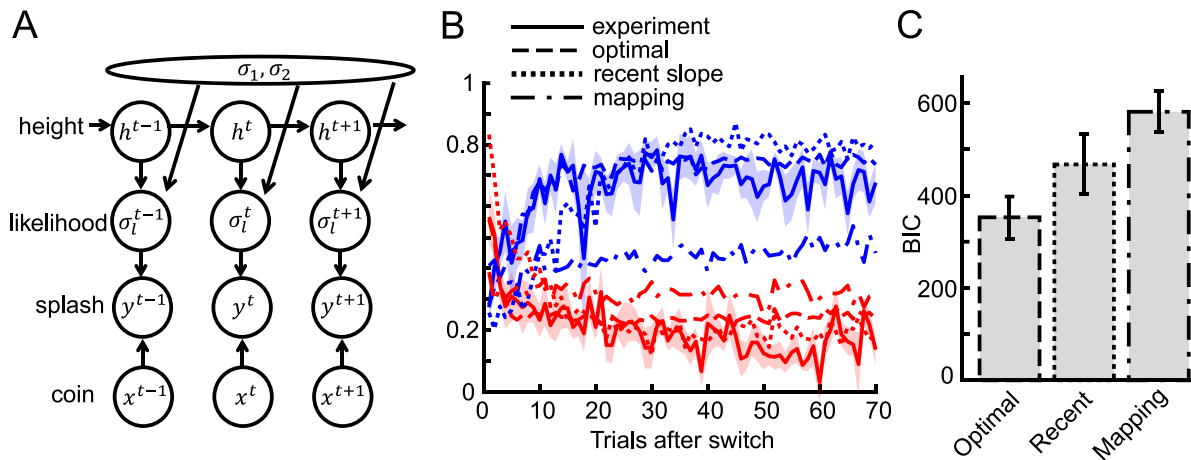


Figure 5. Description of the optimal model and fitted results. (A) Schematic diagram of the generative model of the task. This diagram depicts the probabilistic structure of the variables in the task. At trial  $t$ , the fully Bayesian model makes use of this generative model and infers the coin location  $x^t$  from the observed splash  $y^t$  and all past splashes and coin location feedback. (B) Fitted results of three models for the with-instructions group. Blue and red lines show the results for the narrow and wide conditions, respectively. Dashed, dotted, and dash-dot lines correspond to optimal, recent slope, and static mapping models, respectively. (C) For model comparison we plotted the mean BIC of each model across subjects in the with-instructions group. The error bars show the standard error.

dependent slopes, even if their strategy was static with no switching behavior. In this model, the observer nonlinearly maps the splash position onto the coin location estimations. This model learns the parameters initially, but does not incorporate any additional adaptive processes. (b) Recent slope learning: For this simple and flexible learning model, we implemented a system that learns the slope (and essentially likelihood) from recent  $n$  trials. This model can adapt to changes in likelihood. This strategy is less computationally demanding and easier to implement than the fully Bayesian model below. (c) Bayesian learning: We used a fully Bayesian optimal model that takes likelihood switching into account. This model contains a generative model of the task (Figure 5A), which it uses to make estimations about the coin location. At trial  $t$ , the splash location  $y^t$  is observed and the optimal observer estimates the unknown coin location  $x^t$ , and to do so, it has to estimate the likelihood width  $\sigma_l^t$ .

Among the three models, the fully optimal model was the closest match to human behavior (Figure 5B). The optimal model replicated the slower learning phenomenon seen in the wide-to-narrow switch condition compared with the narrow-to-wide switch condition, and the underestimation of slope in the narrow condition. We compared the fitness of the three models quantitatively in terms of Bayesian Information Criterion (BIC), which is a criterion for comparing different models (Figure 5C). We calculated BIC by fitting each model to all trials for each subject, then averaged BIC across subjects. We found that all models were significantly different from one another, and the optimal model was the best (one-way repeated mea-

asures ANOVA,  $p < 10^{-8}$ , and posthoc Tukey-Kramer test with 95% confidence interval). The estimated free parameter, the switching probability, of the optimal model ( $0.051 \pm 0.020$  SE in the with-instruction group), was a little higher than the actual average value of the experimental setting,  $1/85 \approx 0.012$ , but was not significantly different ( $t$  test,  $p = 0.09$ ). The best-fitted parameter  $n$  of the recent slope model was  $22.1 \pm 2.6$  SE. This best fit of the optimal model suggests that the subjects learned the likelihood in a very efficient way. The result for the without-instructions group was similar, although the fitted parameter was much more broadly distributed across subjects ( $0.26 \pm 0.12$  SE). We also checked the performance of the optimal model taking the mode of the posterior distribution as the estimation of the coin location (MAP estimate) instead of the mean. Their performance was similar but the model that takes the mean was significantly better (Supplementary Figure S1).

## Experiment 2: Context dependent likelihood learning, not slope learning

Although we reported that the subjects displayed adaptive behavior in response to changes in likelihood, one could argue that the subjects just learned the slope, rather than the likelihood, and switched between different learned slopes. This issue is deeply related to how uncertainty is represented and how learning of uncertainty occurs in the brain. To show that the subjects actually learned the likelihood, we conducted Experiment 2, in which we tested whether the learned

likelihood was generalized to a new prior using a context dependent learning paradigm.

Trials in this experiment were divided into three phases: the likelihood learning phase, the new prior learning phase, and the generalization test phase. In the likelihood learning phase, the likelihood width was indicated by the color (green or red) of the splash, which was randomly chosen on a trial-by-trial basis. The prior width was held constant. The subjects received feedback about the correct coin location in the first two phases. The representative subject successfully learned two contextual likelihoods (Figure 6A). At the beginning of the new prior learning phase, the prior width changed to a different value (in Figure 6A, narrow to wide prior), after which it remained fixed in subsequent trials. Either one of the two likelihoods (one color only) was used in this phase. The representative subject came to rely more on observation, as predicted theoretically. The change was almost instant because we explicitly showed the prior on the screen. In the generalization test phase, both likelihoods were used but no feedback was given.

After the prior switch, subjects were not able to directly learn the new slope for the color that was not used during the second phase, because the subjects never get feedback for that color in the latter two phases. If slopes were learned instead of the likelihoods then we would expect the slope to be unchanged. According to Bayesian predictions, if likelihoods were learned then this manipulation should affect the slope for the untrained color.

We measured the generalization effect using the difference between the slope before and after the prior switch. The representative subject generalized the learned likelihood (that had not been used with the new prior) to the new prior (Figure 6A). The generalization effect was significant when averaged across subjects (Figure 6B, paired two sample  $t$  test,  $p = 0.0002$ , comparison between the slope calculated from the last 100 trials in the first phase and the slope from the first 100 trials in the last phase). Note that even though the two lines during the last phase (Figure 6B) look close, each subject is involved in either red or blue line, not both. The slopes were significantly different when analyzed for each subject (paired two sample  $t$  test,  $p = 0.0004$ ).

One argument poses that if the learned quantity was slope, the learned slope for the color not used in the new prior learning phase could have been gradually forgotten during the phase, and then the slope was attracted toward the newly learned slope for the other color. This could explain the red line in Figure 6B, but not the blue line. The generalization effect was significant even when we analyzed only the data for the blue line condition (paired two sample  $t$  test,  $p = 0.03$ ). Thus, even this extended model of slope learning

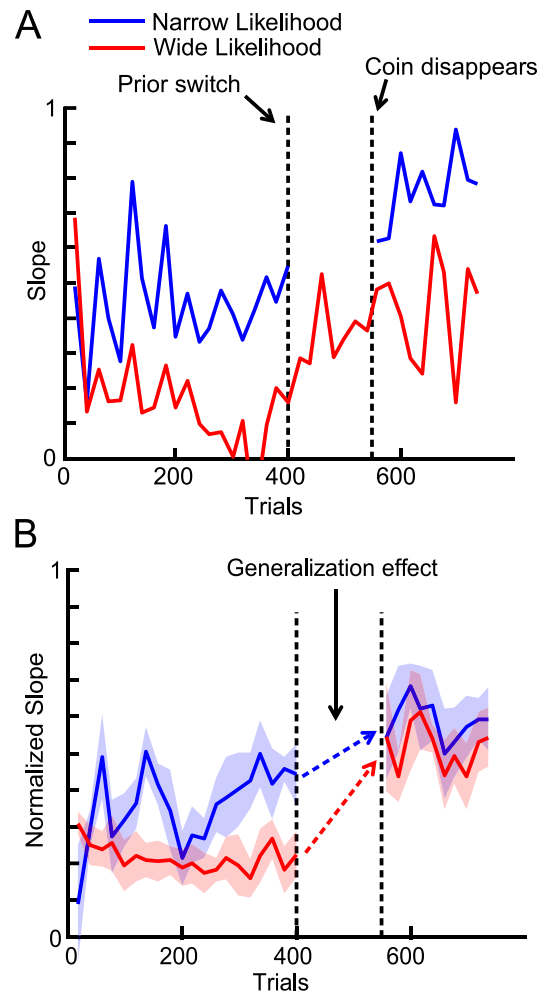


Figure 6. Generalization of likelihood learning (Experiment 2). (A) Data for a representative subject. Blue and red lines correspond to the narrow and wide likelihood conditions, respectively. The slope was calculated by dividing the trials into 10-trial bins and then categorizing the trials in each bin based on condition. We then calculated the slope for each condition. For this subject, the wide likelihood was used to learn the new prior. The predicted slope values from the simple Bayesian model were 0.1 (red) and 0.5 (blue) before the switch and 0.5 (red) and 0.9 (blue) after the switch. (B) Normalized slope for the likelihood that was not used during new prior learning (e.g., the blue line for the subject in [A]) averaged across all subjects. The slope was normalized in the sense that, for the subjects whose initial prior was wider, the obtained slope value was subtracted from 1 to flip the results vertically and compare them directly across different initial prior conditions. Clearly, training with one likelihood condition in the new prior learning phase carried over to the other likelihood condition (dashed arrows).

cannot explain the data. These results clearly show that the subjects learned the likelihood, not the slope, in a context-dependent way, and combined the learned likelihood with the new prior.

## Discussion

In this study, we changed the variance of sensory observations and thus examined likelihood learning. We found that subjects learned to rely more on reliable visual cues. Indeed, this should follow naturally from the Bayesian prior/likelihood combination. We continued with a quantification of the dynamics of likelihood learning. Subjects integrated subsequent trials into an estimate of likelihood, and this integration process was affected by task instructions. Our human behavioral data were well fit by a Bayesian model that incorporated the process of likelihood switching during the experiment. Using a generalization strategy, we showed that the subjects learned the likelihood, and not the direct slope, in a context dependent way, as the data cannot be explained by simple error correcting algorithms.

One could argue that, although the likelihood width was the quantity to be learned in this experiment, learning might have been achieved by acquiring a prior over some hidden parameter that implicitly determined the likelihood, making this a prior learning task. We used the words likelihood and prior to refer to the likelihood function for the task in each trial, i.e., the likelihood of the coin location obtained from the splash location, and the prior for the task, i.e., the overall distribution of coins, in each trial. Thus, the likelihood is essentially information about a task-relevant target (coin location) obtained from current observations, and the prior is information before the current observations. In this paper, we used a task in which the variable to be estimated and the current observation were clearly defined in each trial to investigate whether and how human subjects learned how strongly they could rely on observation versus the prior to perform the task efficiently. We found evidence of learning about the likelihood of the task, regardless of what the underlying process is. How subjects identify the quantity to be estimated when it appears to be arbitrary and how subjects learn about more detailed stochastic structures embedded in a likelihood and a prior are important topics, although they are not the main focus of our study.

In terms of a Bayesian model of learning, prior learning and likelihood learning take essentially the same form. Here, we showed that they share the same features. For example, it has been shown theoretically (DeWeese & Zador, 1998) and experimentally (Berniker et al., 2010; Miyazaki et al., 2005) that the learning speed of the variance parameter of a prior is faster when the distribution is changed from narrow to wide, as opposed to the other way around, which we observed. The computational similarity of prior and likelihood learning raises the question of how similar the two learning phenomena are.

In a recent study on prior learning, Berniker et al. (2010) used a similar experimental paradigm to that in this study. Despite a small number of differences (instructions, screen background, parameters) it seems meaningful to compare the results. There, prior learning was also found to follow the predictions of a Bayesian learning model. However, the initial learning in that experiment required about 200 trials, which is about ten times slower than the likelihood learning observed in the current study. This might be partially because both the prior and likelihood were unknown in the study by Berniker et al. (2010). However, it is implausible that this is the sole reason for the difference. Another possibility is that, in their experiment, learning of the prior mean slowed down the learning process. Although subjects also had to learn the center of the likelihood even in our experiment, one possible argument is that it is probably more natural for the subjects to guess that the splash location matches the likelihood center in our experiment, than to guess the screen center is the prior center in Berniker et al.'s experiment. However, they showed that the learning of the prior mean was very fast (roughly within the first ten trials), so the need to learn the mean is unlikely the reason of the much slower learning. Computationally, the difference between the two experiments suggests that the natural initial assumptions made about the likelihood function are more vague than those made about the prior. This might also suggest that learning about the prior and likelihood involves different neural mechanisms, which is consistent with a recent finding that prior and likelihood uncertainty is represented in different regions in the brain (Vilares, Howard, Fernandes, Gottfried, & Kording, 2012).

Although the mean slope of the behavior and the optimal model matched well, the standard deviation of the slope was larger for the behavioral data (Supplementary Figure S2). There could be several possible reasons. First, even if it is a very simple task to point to the middle point between two clearly visible dots, the subjects' responses certainly have some variance. This variance might come from calculation variability and motor variability. It is possible to incorporate such additional variance parameters to the model. However, we constructed the model with minimum number of free parameters to show our main point as clearly as possible. Second, not all subjects seemed to do their best to minimize the distance between their response and the coin on every trial. Despite the instructions, some of them appeared to make rapid responses to shorten the experiment. Such subjects might still have made optimal responses on average but with a greater variance.

We found that the subjects who were not given instructions about the likelihood switch showed weaker

learning on average than those who were given instructions. It appears that the subjects in the without-instructions group did not learn the likelihood effectively, but this result was expected because the subjects could adapt an arbitrary generative model as opposed to the optimal model with the knowledge of likelihood switching. Indeed, subjects who reported that they noticed the switching of the likelihood function tended to show stronger learning than subjects who did not notice switching. It is possible that noticing the switching of the likelihood function involved some kind of meta-learning about the probabilistic structure of the task. This issue may be an important topic for future research.

Context dependent learning enables us to flexibly adapt to new environments and switch between them. Many studies have shown that the brain can, at least partially, achieve context dependent learning in the visual (Seydell et al., 2010), tactile (Nagai, Suzuki, Miyazaki, & Kitazawa, 2012), multisensory (Yamamoto, Miyazaki, Iwano, & Kitazawa, 2012), and motor (Osu, Hirai, Yoshioka, & Kawato, 2004) modalities, among others. We showed that different likelihoods can be learned simultaneously in different contexts defined by color cues. Some studies have found color-cued contextual adaptation (Addou, Krouchev, & Kalaska, 2011; Osu et al., 2004; Wada et al., 2003), while others have shown that color-cue alone cannot induce contextual learning for multiple distributions (Nagai et al., 2012; Seydell et al., 2010). Because the tasks are quite different between these experiments and that in the current study, we cannot reasonably discuss the reason for this discrepancy. However, investigations of context enabled learning for multiple distributions will produce insights about the mechanisms of learning.

In our task, the correct feedback was always given to the subject. In reality, a feedback is not always given. In the fish-catching task example, there are different possible ways that the likelihood is updated when the catching fails. The catcher might use different sources of information across space, time, and sensory modalities providing the true fish location, or the catcher might be able to learn the likelihood only from successful trials. How the learning differs between in our simplified paradigm and in more realistic situation is an interesting future topic.

Here, we showed that human subjects could learn the quality of sensory input represented in likelihood over time and adjust their behavior in response to changes in likelihood. Many internal and external factors affect the precision and reliability of sensory information, and we adapt to such changes to efficiently perceive and interact with the world. How and where in the brain likelihood is learned and combined with the prior is an important topic for understanding human behavior.

*Keywords: Bayesian models, likelihood learning, sensorimotor integration, context-dependent learning*

## Acknowledgments

We thank Max Berniker and Iris Vilares for their comments on the manuscript. This work was supported by Grant-in-Aid for Young Scientists (B) (23700309) from the Japan Society for the Promotion of Science (YS) and the NIH: 1R01NS063399 (KPK). The authors declare no competing financial interests.

Commercial relationships: none.

Corresponding author: Yoshiyuki Sato.

Email: yoshi.yk.sato@gmail.com.

Address: Graduate School of Information Systems, University of Electro-Communications, Chofu, Tokyo, Japan.

## References

- Addou, T., Krouchev, N., & Kalaska, J. F. (2011). Colored context cues can facilitate the ability to learn and to switch between multiple dynamical force fields. *Journal of Neurophysiology*, *106*(1), 163–183.
- Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*, *14*(3), 257–262.
- Baddeley, R., Ingram, H., & Miall, R. (2003). System identification applied to a visuomotor task: Near-optimal human performance in a noisy changing task. *Journal of Neuroscience*, *23*(7), 3066–3075.
- Battaglia, P. W., & Schrater, P. R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. *Journal of Neuroscience*, *27*(26), 6984–6994.
- Berniker, M., & Kording, K. (2011). Bayesian approaches to sensory integration for motor control. *Wiley Interdisciplinary Reviews: Cognitive Science*, *2*(4), 419–428.
- Berniker, M., Voss, M., & Kording, K. (2010). Learning priors for Bayesian computations in the nervous system. *PLoS ONE*, *5*(9), e12686.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.
- Burge, J., Ernst, M., & Banks, M. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of Vision*, *8*(4):20, 1–19,

- <http://www.journalofvision.org/content/8/4/20>, doi:10.1167/8.4.20. [PubMed] [Article]
- Burge, J., Girshick, A. R., & Banks, M. S. (2010). Visual-haptic adaptation is determined by relative reliability. *Journal of Neuroscience*, *30*(22), 7714–7721.
- DeWeese, M., & Zador, A. (1998). Asymmetric dynamics in optimal variance adaptation. *Neural Computation*, *10*(5), 1179–1202.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.
- Ernst, M. O., Banks, M. S., & Bühlhoff, H. H. (2000). Touch can change visual slant perception. *Nature Neuroscience*, *3*(1), 69–73.
- Jacobs, R. A., & Fine, I. (1999). Experience-dependent integration of texture and motion cues to depth. *Vision Research*, *39*(24), 4062–4075.
- Jazayeri, M., & Shadlen, M. N. (2010). Temporal context calibrates interval timing. *Nature Neuroscience*, *13*(8), 1020–1026.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annual Review of Psychology*, *55*, 271–304.
- Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? In *Perception 36 ECVF Abstract Supplement*.
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.
- Miyazaki, M., Nozaki, D., & Nakajima, Y. (2005). Testing Bayesian models of human coincidence timing. *Journal of Neurophysiology*, *94*(1), 395–399.
- Nagai, Y., Suzuki, M., Miyazaki, M., & Kitazawa, S. (2012). Acquisition of multiple prior distributions in tactile temporal order judgment. *Frontiers in Psychology*, *3*, 276.
- Osu, R., Hira, S., Yoshioka, T., & Kawato, M. (2004). Random presentation enables subjects to adapt to two opposing forces on the hand. *Nature Neuroscience*, *7*(2), 111–112.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Sato, Y., & Aihara, K. (2011). A bayesian model of sensory adaptation. *PLoS ONE*, *6*(4), e19377.
- Sato, Y., Toyozumi, T., & Aihara, K. (2007). Bayesian inference explains perception of unity and ventriloquism aftereffect: Identification of common sources of audiovisual stimuli. *Neural Computation*, *19*(12), 3335–3355.
- Seydell, A., Knill, D. C., & Trommershäuser, J. (2010). Adapting internal statistical models for interpreting visual cues to depth. *Journal of Vision*, *10*(4):1, 1–27, <http://www.journalofvision.org/content/10/4/1>, doi:10.1167/10.4.1. [PubMed] [Article]
- Stocker, A., & Simoncelli, E. (2006). Sensory adaptation within a Bayesian framework for perception. *Advances in Neural Information Processing Systems*, *18*, 1289–1296.
- Tassinari, H., Hudson, T. E., & Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *Journal of Neuroscience*, *26*(40), 10154–10163.
- Trommershäuser, J., Kording, K., & Landy, M. S. (Eds.). (2011). *Sensory cue integration*. New York: Oxford University Press.
- van Beers, R. J., van Mierlo, C. M., Smeets, J. B. J., & Brenner, E. (2011). Reweighting visual cues by touch. *Journal of Vision*, *11*(10):20, 1–16, <http://www.journalofvision.org/content/11/10/20>, doi:10.1167/11.10.20. [PubMed] [Article]
- 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.
- Wada, Y., Kawabata, Y., Kotosaka, S., Yamamoto, K., Kitazawa, S., & Kawato, M. (2003). Acquisition and contextual switching of multiple internal models for different viscous force fields. *Neuroscience Research*, *46*(3), 319–331.
- Wei, K., & Körding, K. (2010). Uncertainty of feedback and state estimation determines the speed of motor adaptation. *Frontiers in Computational Neuroscience*, *4*, 11.
- Wozny, D. R., & Shams, L. (2011). Computational characterization of visually induced auditory spatial adaptation. *Frontiers in Integrative Neuroscience*, *5*, 75.
- Yamamoto, S., Miyazaki, M., Iwano, T., & Kitazawa, S. (2012). Bayesian calibration of simultaneity in audiovisual temporal order judgments. *PLoS ONE*, *7*(7), e40379.