

RESEARCH ARTICLE | *Control of Movement*

Sensorimotor priors are effector dependent

Cong Yin,¹ Huijun Wang,^{2,3,4,5} Kunlin Wei,^{2,3,4,5} and Konrad P. Körding⁶

¹Capital University of Physical Education and Sports, Beijing, China; ²School of Psychological and Cognitive Sciences, Peking University, Beijing, China; ³Beijing Key Laboratory of Behavior and Mental Health, Beijing, China; ⁴Key Laboratory of Machine Perception, Ministry of Education, Beijing, China; ⁵Peking-Tsinghua Center for Life Sciences, Beijing, China; and ⁶Department of Neuroscience and Department of Bioengineering, University of Pennsylvania, Philadelphia, Pennsylvania

Submitted 3 April 2018; accepted in final form 12 May 2019

Yin C, Wang H, Wei K, Körding KP. Sensorimotor priors are effector dependent. *J Neurophysiol* 122: 389–397, 2019. First published May 15, 2019; doi:10.1152/jn.00228.2018.—During sensorimotor tasks, subjects use sensory feedback but also prior information. It is often assumed that the sensorimotor prior is just given by the experiment and that the details for acquiring this prior (e.g., the effector) are irrelevant. However, recent research has suggested that the construction of priors is nontrivial. To test if the sensorimotor details matter for the construction of a prior, we designed two tasks that differ only in the effectors that subjects use to indicate their estimate. For both a typical reaching setting and an atypical wrist rotation setting, prior and feedback uncertainty matter as quantitatively predicted by Bayesian statistics. However, in violation of simple Bayesian models, the importance of the prior differs across effectors. Subjects overly rely on their prior in the typical reaching case compared with the wrist case. The brain is not naively Bayesian with a single and veridical prior.

NEW & NOTEWORTHY Traditional Bayesian models often assume that we learn statistics of movements and use the information as a prior to guide subsequent movements. The effector is merely a reporting modality for information processing. We asked subjects to perform a visuomotor learning task with different effectors (finger or wrist). Surprisingly, we found that prior information is used differently between the effectors, suggesting that learning of the prior is related to the movement context such as the effector involved or that naive models of Bayesian behavior need to be extended.

Bayesian theory; learning rate; motor learning; prior learning; sensorimotor learning

INTRODUCTION

For any sensorimotor behavior, we rely on our current sensory inputs (likelihood) as well as knowledge we have accumulated from previous experience (prior). Bayesian statistics describe how prior and likelihood information can be optimally combined as a function of their respective distributions to yield a posterior-probability estimate. How the agent constructs his or her prior is one of the central questions for decision making in general and sensorimotor control in particular (Knill and Richards 1996; Körding 2007; Russell and Wefald 1991). Bayesian studies typically assume that experi-

ence in daily life shapes our prior (Chambers et al. 2017; Geisler and Kersten 2002; Mamassian et al. 1998; Stocker and Simoncelli 2006). However, for experimental studies, the prior is conveniently defined by imposing some movement or stimulus statistics on the agent, who learns it over a series of trials (Jazayeri and Shadlen 2010; Körding and Wolpert 2004; Vilares et al. 2012). For example, if subjects are exposed to a narrow distribution of target locations in a spatial localization task, we presume that they will develop a narrow prior, which can be and typically is confirmed by experimental data (Bejjanki et al. 2016). Thus movement-related statistics can be learned and serve as a prior for later movements.

However, two recent studies suggest that a prior is not purely informational but critically dependent on the learning context. A motor adaptation study has found that a prior learned in one movement direction is only partially generalized to other directions (Fernandes et al. 2014), suggesting that the use of the prior might depend on the movement context. Similarly, a prior learned by a sensorimotor estimation task is not completely generalized to a perceptual judgment task even when both tasks are performed with identical sensory feedback (Chambers et al. 2017). Notably, a time perception task showed that subjects form multiple priors if distributions are coupled with distinct motor outputs (Roach et al. 2017). Hence, construction of a prior appears to depend on its learning context. This is in stark contrast to how the prior is typically viewed in the Bayesian framework of human decision making. For perceptual decision making, people have found that two qualitatively different tasks appear to share the same prior (Bejjanki et al. 2016). For movement control, Bayesian information processing is assumed to relate to motor planning, and the effector is treated merely as a reporting modality of this planning process (Trommershäuser et al. 2011). Thus the prior is often assumed to be shared across effectors in the sensorimotor control domain. However, aforementioned recent findings of context-dependent prior would suggest that the prior acquired by motor experience might not be equally utilized for different effectors. However, this remains to be examined.

In the present study, we used a visuomotor adaptation paradigm to test whether the learning of priors is effector dependent. Over a sequence of trials, subjects learned a novel visuomotor mapping between a moving effector and a cursor representation. These trials were performed by either the finger or the wrist in an interleaved fashion while the statistics of visuomo-

Address for reprint requests and other correspondence: K. P. Körding, Dept. of Neuroscience and Dept. of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104 (e-mail: kording@upenn.edu).

tor mapping remained consistent across these two effectors. We found that subjects could learn the task effectively by either of the two effectors, taking both prior and likelihood into consideration on a trial-by-trial basis. However, when fitting a Bayesian model, their use of prior information was different. Model comparison results suggest that subjects used two different precision degrees of priors for the finger and the wrist despite the fact that they were exposed to the same prior statistics or, alternatively, that their behavior was not Bayesian in a simple way.

METHODS

Subjects. Eight right-handed healthy subjects (3 men, 5 women; age 24.5 ± 2.8 yr) participated in the experiment. All subjects were naive to the purpose of the experiment and were paid according to their performance. The experimental protocol was approved by the Northwestern University Institutional Review Board and in accordance with the Northwestern University Institutional Review Board's policy statement on the use of human subjects in experiments. Written informed consent was obtained from all subjects before the experiment.

Experimental protocol. We adopted the experimental paradigm in one of our previous studies where subjects were required to control the movement of a cursor (Körding and Wolpert 2004). The critical manipulation was that subjects learned a novel visuomotor mapping (from feedback to eventual compensation) between the cursor movement and limb movement. In separate trials, subjects controlled the position of a cursor either by moving their right index finger or by rotating their right wrist (Fig. 1). The vision of their hand and forearm was excluded during the experiment, because they were covered by a semitransparent mirror. For the finger effector, a sensor was attached to the tip of the right index finger, whose movement on a smooth

tabletop was continuously captured. The position of the cursor would overlay with the actual finger position and move in the tabletop plane as the projector-and-mirror system was calibrated. For the wrist effector, three sensors were mounted to a metal sheet attached flat to the back of the right hand. Subjects were instructed to rest their right elbow on the table and rotate their wrist in the air. We calculated the angles between the hand plane defined by the three sensors and the horizontal Cartesian plane. The rotation angle along the Cartesian x -axis was used to drive the left-right cursor movement, and the rotation angle along the y -axis was used to drive the forward-backward cursor movement. The mapping between angular and positional displacement was linear in both directions, because 1° rotation would drive the cursor to move 5 mm. The cursor would appear at the center of the screen when the two wrist angles were zero. The motion capture was performed by using an Optotrak 3D Investigator motion capture system.

On each trial, subjects saw a yellow cross (12 mm along each direction) representing the starting location, a white circle (6-mm radius) representing the target, and a green circle (3-mm radius) representing the cursor position. The target was 100 mm from the starting position to the straight-ahead direction. Subjects were instructed to make a ballistic movement to hit the target accurately in each trial. To initiate a new trial, subjects had to reposition the cursor to the starting cross. The cursor was visible only within a distance of 10 mm from the center of the starting yellow cross. For the finger trials, the starting position was obtained by putting the right index finger close to the center of the desk, aligned with the body midline. For the wrist trials, subjects could bring the cursor to the starting position by keeping their hand horizontal, which effectively made the two rotation angles zero. Once the cursor stayed at the starting point for 500 ms, the target would appear, and a computer speaker played a beep sound to signal the subjects to move. Because some subjects had difficulty returning the cursor (only visible within 10-mm perimeter of the starting cross) to the starting position, the cursor was set to flash every second for 50 ms if it did not return in 2 s after the beginning of a trial. When the finger left the start position (>14 mm) or the wrist left the horizontal plane (either of the two angles $>2.9^\circ$), the cursor was extinguished. When the hidden cursor moved past half distance (5 cm), the cursor was briefly shown for 100 ms. Unknown to the subject, this midpoint cursor was displaced to the right by an amount that was drawn each trial from a Gaussian distribution with a mean of 1 cm and a standard deviation of 1 cm. This was the imposed prior distribution of the cursor shifts, which was learned by the subjects over trials. Of note, the visual perturbation was not imposed on the veridical position of the cursor. Instead, it was imposed relative to the midline connecting the starting point and the target. This was because the wrist movement exhibited more motor variability than the finger movement in our piloting data. To keep the variance of the mid-way visual feedback consistent between the two effector conditions, we applied visual perturbations without reference to the actual movements.

Critically, we changed the quality of the cursor feedback on a trial-by-trial basis to modify its sensory uncertainty. One of four types of feedback quality (σ_S , σ_M , σ_L , σ_∞) was randomly selected for each trial. The relative frequencies of the four types across trials were set at 3, 1, 1, and 1, respectively. The σ_S feedback (small uncertainty) was a typical green circle with a 3-mm diameter. Giving a cloud-type impression, the σ_M feedback (medium uncertainty) was three green circles, which were distributed as a two-dimensional Gaussian with an standard deviation of 1 cm in two directions. The σ_L feedback (large uncertainty) was analogous but with a standard deviation of 4 cm. No cursor feedback was provided in the σ_∞ case. Subjects were instructed to take into account what they saw at the half distance and to get as close to the target as possible. The trial was finished when the cursor moved 5 cm further past the half distance. No feedback of the final cursor location was provided except for the σ_S condition; in this condition, the position of the final cursor feedback was calculated as




Fig. 1. Experimental setup. Subjects are required to use either their finger or wrist to control a cursor to accurately hit a target (white circle) on the projection screen. In the finger condition, the spatial coordinates of the finger are mapped on the spatial coordinates of the cursor. In the wrist condition, the wrist rotation around the x -axis and the y -axis are mapped onto the x and y coordinates of the cursor. The starting position is marked by a yellow cross. The actual cursor trajectory and the displayed trajectory are indicated as the dashed line and the solid line, respectively. These trajectories are not shown to the subject; instead, only a single cursor (the σ_S condition) or 3 cursors (the σ_M and σ_L conditions; shown as green dots) are shown when the actual cursor passes the half-distance line. In the σ_∞ condition, no such midpoint feedback is given.

the position of veridical cursor feedback at the final point minus the position of the veridical cursor feedback at the midpoint, plus the lateral shift with respect to the midline. Here, veridical cursor feedback refers to the location of cursor based on the movement. Because end-point feedback is critical for learning of the mapping, the σ_S condition occupied half of the trials. A beep sound was played twice, once when subjects reached midpoint and once when they completed the movement. If the green cursor successfully touched the target, the target turned red and subjects were rewarded by having a point added to their score. If a successful reach was achieved in those trials without end-point feedback, the score point was accumulated without showing.

Before the formal experiment, subjects familiarized themselves with 100 σ_S -condition trials without perturbations (50 trials for the finger and the wrist conditions, respectively). The formal experiment consisted of 2,040 trials for each subject. Trials were organized in alternating blocks where 10 successive finger trials were followed by 10 successive wrist trials, and so on. On each trial, a text prompt was displayed beneath the starting position to indicate whether it was a finger or wrist trial. Thus the same prior distribution was learned by two effectors with interleaving trial blocks. The experiment lasted ~2.5 h.

Data analysis. The lateral deviation of the final finger position or wrist rotation in each trial indicated subjects' perceived miscalibration, informed by feedback at the midpoint. It was repeatedly shown in our previous studies that people can combine their previous knowledge about the distribution of lateral shift (the prior) with the feedback that they receive midway through the movement (the likelihood) to estimate the current position of the hidden cursor in a nearly optimal fashion (Fernandes et al. 2012; Körding and Wolpert 2004). This means that subjects combine the means of the learned prior and likelihood, weighted by their relative precision (the inverse of the variance). Thus, according to Bayesian theory, the optimal estimate of the lateral shift can be written as

$$\hat{x} = \frac{\sigma_f^2}{\sigma_p^2 + \sigma_f^2} x_p + \frac{\sigma_p^2}{\sigma_p^2 + \sigma_f^2} x_f, \quad (1)$$

where $\hat{x}$ is the estimated lateral shift, x_p and x_f denote the mean of the prior and the sensed midpoint feedback position, and σ_p^2 and σ_f^2 denote the variance of the prior and the midpoint feedback. As a proxy of sensed cursor position, x_f was defined as the midpoint cursor position in the σ_S condition and as the centroid of the cursor cloud in the σ_M and σ_L conditions. In each feedback condition for each effector, we performed a linear regression of the estimated position ($\hat{x}$) as a function of the sensed cursor position (x_f). The estimated lateral shift was measured as the lateral deviation of the veridical cursor position based on movement between midpoint and end point. The regression slope gives us a measure of the relative reliance on midpoint feedback, which serves as the basis of our analysis. Thus we call it weight of feedback (w_f) in the following text. Importantly, if subjects behave in a "Bayes-optimal" way, then the weight of feedback should be equal to the coefficient in Eq. 1 as

$$w_f = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_f^2} \quad (2)$$

If a subject's prior uncertainty is much higher than the likelihood uncertainty, then this value will approach 1; i.e., the subject will rely almost exclusively on the midpoint feedback information. If, on the other hand, the prior is much less uncertain than the likelihood, then this value will approach 0, and subjects should almost disregard the midpoint feedback information. Only the second half of the formal experiment (the last 1,020 trials) were used for analysis, because subjects needed to learn the prior distribution over trials (see a similar treatment in Körding and Wolpert 2004).

For the σ_S condition, there is minimal feedback uncertainty for locating a single cursor. If the theoretical variance of likelihood

approaches 0, the weight of feedback should be 1 according to Eq. 2. For the σ_M and σ_L conditions, a cloud of three cursors was shown with a standard deviation of 1 and 4 cm, respectively. We assume this manipulation added extra feedback uncertainty to the likelihood. Theoretically, the amount of the added variance equals the Gaussian variance of the cursor cloud divided by the number of cursors (Tassinari et al. 2006). For the σ_∞ condition without any midpoint feedback, the variance of likelihood should be infinite, and the weight of feedback should be 0.

In addition to computing the weight of feedback from the linear regression, we can compute the variance of the regression residual for each condition. As typically assumed for linear regression, the residual was Gaussian distributed; we can thus obtain the following:

$$\sigma_R^2 = \sigma_M^2 + w_f^2 \sigma_f^2, \quad (3)$$

where σ_R^2 and σ_M^2 denote the variance of residual and motor noise, respectively. $w_f^2 \sigma_f^2$ can be viewed as planning noise for visuomotor transformation where the visual feedback leads to the motor response. Leveraging on the fact that the theoretical weight of feedback for the σ_∞ condition is 0, we can estimate motor noise by calculating the variance of the residual in the σ_∞ condition. This motor noise is computed for the finger and the wrist condition separately.

To infer the mean of the prior, we can rearrange Eq. 1 to obtain the following:

$$x_f = x_p + \frac{\sigma_p^2 + \sigma_f^2}{\sigma_f^2} (x_f - \hat{x}). \quad (4)$$

In other words, we can estimate the mean of subjects' prior (x_p) as the intercept of the linear regression of x_f as a function of $x_f - \hat{x}$. We built four models to test whether different degrees of precision for the priors are essential to account for the observed weights of feedback in different conditions. Among them, two models assume a common prior for the two effectors, and the other two models assume two different precision degrees for the two priors. For the specification of feedback likelihood, we assume that the likelihood uncertainty of visual feedback is equal to the sum of baseline likelihood variance ($\sigma_{baseline}^2$) and added variance (i.e., uncertainty caused by using the cloud of cursors in the σ_M and σ_L conditions). We can estimate the baseline likelihood variance $\sigma_{baseline}^2$ as the σ_f^2 (Eq. 2) in the σ_S condition where no added noise is applied. By using Eq. 2 and assuming the prior uncertainty equals the imposed prior variance (i.e., $\sigma_p^2 = 1$), we can obtain σ_f^2 with the observed weight of feedback in the σ_S condition. Thus $\sigma_{baseline}^2$ is estimated by using trials from both effector conditions. The baseline variance is not the same as motor noise because it is a single parameter for both effectors, whereas motor noise estimated from the regression analysis is for two effectors separately. Subsequently, the likelihood uncertainty for the σ_M and σ_L conditions can be specified as $\sigma_{baseline}^2 + 1^2/3$ and $\sigma_{baseline}^2 + 4^2/3$, respectively. Note that 1 and 4 are the standard deviations of the underlying Gaussian distributions where three cursors are drawn. Alternatively, we can make $\sigma_{baseline}^2$ a free parameter. Thus we have four models: *M1* uses an estimated $\sigma_{baseline}^2$ and a common prior for two effectors, *M2* makes $\sigma_{baseline}^2$ a free parameter and uses a common prior for two effectors, *M3* uses an estimated $\sigma_{baseline}^2$ but with two effector-specific priors, and *M4* makes $\sigma_{baseline}^2$ a free parameter and uses two effector-specific priors. Hence, these four models have one, two, two, and three free parameters, respectively (Table 1).

For each model, we first calculated the squared difference between the model-predicted weights of feedback (Eq. 2) and the observed weights of feedback across the six conditions (2 effectors $\times$ 3 feedback uncertainty conditions, excluding the σ_∞ condition, which had no likelihood information provided). By minimizing the sum of the squared difference for each model, we could infer the model parameters for each subject. Taking the model complexity into consider-

Table 1. Four different models with their respective free parameters

Model	Free Parameters	Prior	Baseline Likelihood
M1	σ_{prior}^2	=	Estimated
M2	$\sigma_{prior}^2, \sigma_{baseline}^2$	=	Free parameter
M3	$\sigma_{prior, finger}^2, \sigma_{prior, wrist}^2$	$\neq$	Estimated
M4	$\sigma_{prior, finger}^2, \sigma_{prior, wrist}^2, \sigma_{baseline}^2$	$\neq$	Free parameter

The models assume either one or two priors for the finger and the wrist effectors. The baseline likelihood uncertainty is either estimated from regression analysis or left as a free model parameter.

ation, we compute the Akaike information criterion (AIC) for individual subjects for each model. To compare the performance of these four models, we compute the overall residual error of model fitting across the six conditions and the variance explained. Model comparison was then performed by one-way ANOVA on subjects' AIC, fitting residuals, and variance explained for independent samples. The significance level was set at $\alpha = 0.05$.

RESULTS

First, we sought to test when the subjects learned the new mapping using their actual shift at midpoint as an indicator (Fig. 2). We divided all the 2,040 learning trials into 34 bins (in each bin, there were 30 trials for each effector). For the finger trials, subjects moved -4.91 ± 1.55 and -6.23 ± 1.57 mm to the left of the midline at midpoint for the first and second bins, respectively, which was smaller than the mean perturbation ($t_7 = 3.29$ and 2.40 , $P < 0.05$). However, from the third bin, subjects' shift was no longer different from the mean perturbation (-7.51 ± 1.55 mm; $t_7 = 2.09$, $P = 0.08$). That is, for the finger trials, subjects started to compensate for the mean shift of the cursor within 90 trials.

In contrast, for the wrist trials, even for the first bin, subjects moved -10.07 ± 4.02 mm to the left of the midline at midpoint, which was not different from the mean perturbation ($t_7 = -0.02$, $P = 0.99$). We then explored the shift at midpoint for the first three wrist trials: they were -3.01 ± 3.87 , -8.37 ± 3.91 , and -8.85 ± 3.26 mm, which had no difference from the mean perturbation. Because subjects first learned the 10 finger trials and then 10 wrist trials, there may be some transfer of learning from the finger to the wrist. In sum, for the wrist trials, subjects compensated for the mean shift within three trials, much faster than for the finger trials.


Before examining a possible difference in prior between effectors, we first established that our behavioral results were

in line with previous findings (For the analysis below, we used only data from the second half of the formal experiment). We found that subjects made trial-by-trial adjustments to imposed visual lateral shifts (Fig. 3). Following a positive perturbation (leftward shift with respect to the mean of the prior, which serves as the likelihood), the veridical cursor location tended to be located to the right in the second half of the movement and vice versa, resulting in a negative correlation between the lateral shifts and the deviation of cursor position from midpoint to end point.

To estimate the weight of feedback and the mean of the prior for the four feedback uncertainty conditions, we pooled the data from the two effector conditions. We found that the weight of feedback decreased as the feedback uncertainty increased, which was qualitatively consistent with Bayes-optimal behavior. The weights of feedback were 0.80 ± 0.04 , 0.54 ± 0.04 , 0.26 ± 0.03 , and 0.01 ± 0.02 (averages $\pm$ SE) for the σ_S , σ_M , σ_L , and σ_∞ conditions, respectively. The weight of feedback for the σ_∞ condition was not significantly different from the theoretical value of 0 ($t_7 = 0.27$, $P = 0.79$). One-way ANOVA on weights of feedback revealed a significant difference between different uncertainty conditions ($F_{3, 21} = 151.56$, $P < 0.001$); post hoc pairwise comparisons confirmed that the weight of feedback significantly decreased with increasing level of feedback uncertainty (σ_S – σ_M : $t_7 = 5.98$, $P < 0.001$; σ_M – σ_L : $t_7 = 6.34$, $P < 0.001$; σ_L – σ_∞ : $t_7 = 5.74$, $P < 0.001$). The means of priors were 9.90 ± 0.18 , 9.97 ± 0.31 , 9.46 ± 0.37 , and 9.64 ± 0.40 mm for these four conditions, respectively (Fig. 4). They were not significantly different across conditions ($F_{3, 21} = 1.00$, $P = 0.45$) and not different from the experimentally imposed average shift size (σ_S : $t_7 = -0.55$, $P = 0.60$; σ_M : $t_7 = -0.11$, $P = 0.92$; σ_L : $t_7 = -1.44$, $P = 0.19$; σ_∞ : $t_7 = -0.91$, $P = 0.39$). This indicates that subjects learned the mean of imposed prior distribution for all the likelihood conditions. Furthermore, the feedback uncertainty affects behavior as it should: subjects place less weight on visual feedback when the feedback information is more uncertain.

We also quantified the motor noise level associated with each effector. Estimated by regression analyses for the σ_∞ condition, the motor noise for the wrist was 89.60 ± 41.02 mm², which was significantly larger than that for the finger (15.48 ± 2.71 mm²; $t_7 = -5.43$, $P < 0.001$). As expected, motor noise was effector dependent. In theory, this difference in motor noise should have no influence on the reliance on likelihood within a trial.

Fig. 2. Actual shift at midpoint over bins. Error bars denote SE. Red dots represent the finger condition, and blue dots represent the wrist condition.






Fig. 3. Actual correction from midpoint to end point plotted as a function of applied visual perturbations. A: data are from a typical subject with midpoint feedback (trials from σ_∞ condition are excluded). Each dot represents a trial. B: data from the 4 feedback uncertainty conditions (σ_S , σ_M , σ_L , σ_∞) for the same subject are shown separately. Red dots represent the finger trials, and blue dots represent the wrist trials. Red and blue lines are the corresponding linear regression fits.

Was there a difference in reliance on the likelihood between the two effectors? The weights of feedback for the finger were 0.69 ± 0.05 , 0.40 ± 0.04 , 0.18 ± 0.04 , and -0.01 ± 0.02 for the σ_S , σ_M , σ_L , and σ_∞ conditions, respectively (Fig. 5A). The weights of feedback for the wrist were 0.92 ± 0.04 , 0.70 ± 0.05 , 0.33 ± 0.04 , and 0.03 ± 0.03 for the σ_S , σ_M , σ_L , and σ_∞ conditions, respectively. Two-way repeated-measures ANOVA revealed that main effect of feedback uncertainty and main effect of effector were both significant ($F_{3, 21} = 183.5$, $P < 0.001$ and $F_{1, 7} = 23.47$, $P < 0.005$). The interaction was also significant ($F_{3, 21} = 6.98$, $P < 0.005$). Post hoc tests showed that for the three conditions with midpoint feedback (i.e., σ_S , σ_M , and σ_L conditions), the weight of feedback for

wrist was significantly larger than that for finger ($P < 0.005$, 0.005 , and 0.05 for σ_S , σ_M , and σ_L conditions, respectively); we did not find significant differences in the σ_∞ condition ($P = 0.36$). Simple main effect analyses confirmed that for both effectors, the weight of feedback decreased with increasing feedback uncertainty. The difference in weight of feedback between the two effectors can be highlighted when the weights of feedback for wrist and the weights of feedback for finger of each subject are plotted against each other (Fig. 5B). It is evident that the latter is larger than the former for almost all individual subjects and in all the conditions with midpoint feedback.

We then built four models to test whether different degrees of precision for the priors could explain the difference in reliance on the likelihood between the two effectors. Because the models have different numbers of parameters, we first compared them by using AIC (Fig. 6A). The AIC values were -15.64 ± 0.99 , -14.34 ± 1.09 , -20.33 ± 1.43 , and -19.15 ± 1.24 for $M1$, $M2$, $M3$, and $M4$, respectively. We found significant difference between these four models ($F_{3, 31} = 5.58$, $P < 0.005$). Importantly, post hoc comparisons revealed a significant difference between $M1$ and $M3$ ($P < 0.05$) and between $M2$ and $M4$ ($P < 0.05$). No significant difference existed between $M1$ and $M2$ ($P = 0.87$) or between $M3$ and $M4$ ($P = 0.90$). Thus the models assuming two different precision degrees of the priors were the best models to account for our data; also, whether the baseline variance was estimated from the data or model fitting did not make a difference.

We then compared the performance of the four models by examining their fitting residuals and variance explained. The subject-average residuals were 0.17 ± 0.03 , 0.15 ± 0.03 , 0.05 ± 0.02 , and 0.04 ± 0.01 for $M1$, $M2$, $M3$, and $M4$, respectively (Fig. 6B). One-way ANOVA revealed a significant difference between models ($F_{3, 28} = 7.21$, $P < 0.001$). Post hoc comparisons indicated that there was a significant differ-





Fig. 4. Scatterplot of the sensed cursor position (x_f) as a function of $x_f - \hat{x}$, where $\hat{x}$ is the estimated position, for an example subject. Data from the 4 feedback uncertainty conditions (σ_S , σ_M , σ_L , σ_∞) are shown separately. Each dot represents a trial. Solid lines are the corresponding linear regression fits. The intersection point of gray dashed line and y-axis shows the estimated mean of the prior.

Fig. 5. Weights of feedback for the finger conditions and the wrist conditions. **A**: average weight of feedback for the 4 feedback uncertainty conditions (σ_S , σ_M , σ_L , σ_∞). Red and blue bars represent the finger and wrist conditions, respectively. * $P < 0.05$; ** $P < 0.005$. n.s., not significant. **B**: weights of feedback for wrist and weights of feedback for finger are plotted against each other. Each dot represents data from a subject. Error bars denote average and SE across subjects. Dotted line denotes equal weights of feedback.



ence between *M1* and *M3* ($P < 0.05$) and between *M2* and *M4* ($P < 0.05$); there was no significant difference between *M1* and *M2* ($P = 0.96$) or between *M3* and *M4* ($P = 0.99$). Thus the models assuming effector-specific priors can better capture the data, independently of how the baseline variance was obtained. This conclusion was further supported by the variance explained by these four models, because *M3* and *M4* had substantially higher r^2 values than *M1* and *M2* (67.8%, 68.2%, 90.1%, and 91.6% for *M1*, *M2*, *M3*, and *M4*, respectively).

For *M1*, the baseline variance estimated from the data was $0.27 \pm 0.07 \text{ cm}^2$; correspondingly, the baseline variance obtained from model fitting (*M2*) was $0.64 \pm 0.23 \text{ cm}^2$. We did not find a significant difference between them, although this might result from a lack of power given the sample size ($t_{14} = 1.53$, $P = 0.14$). With *M3* and its estimated baseline variance, we found that the prior variance of the finger ($0.48 \pm 0.04 \text{ cm}^2$) was significantly smaller than that of the wrist ($2.31 \pm 0.33 \text{ cm}^2$; paired t -test, $t_7 = -5.24$, $P < 0.005$). Similar results were obtained for *M4*: the prior variance of the finger ($0.80 \pm 0.30 \text{ cm}^2$) was significantly smaller than that of the wrist ($2.84 \pm 0.53 \text{ cm}^2$; paired t -test, $t_7 = -5.92$, $P < 0.001$). In sum, whether the likelihood uncertainty was estimated from the data or estimated by model fitting did not make a difference in model performance; it was the number of priors that largely determined the performance of the model. The models assuming effector-specific priors (*M3* and *M4*) could better account for data. The prior uncertainty was significantly larger for the wrist, a result consistent with the finding that the

weights of feedback were larger in the wrist conditions than in the finger conditions.

We then analyzed all of the 2,040 learning trials to examine whether the finger and wrist priors were stationary over the experiment. When we separated all the trials into four bins (255 trials per bin for each effector), we found that for the finger trials, the feedback weights in the σ_S condition were 0.48 ± 0.08 , 0.64 ± 0.07 , 0.66 ± 0.06 , and 0.71 ± 0.05 for the four bins, respectively. One-way repeated-measures ANOVA showed a significant main effect of bins ($F_{3, 21} = 6.85$, $P < 0.01$). Pairwise comparisons revealed that only the weight in the first bin was smaller than those in all the other bins ($t_7 = -4.24$, -2.79 , and -3.73 , $P < 0.01$, 0.05 , and 0.01 , respectively, post hoc least significant difference test); the weights in the other bins were not different from each other. No difference between bins was found for the other three feedback conditions. In contrast, for the wrist trials, the feedback weights did not differ across bins for each of the four feedback conditions. Therefore, the feedback weight changed for the finger trials in the σ_S condition only, and it remained stationary for the second half of the learning trials in our experiment. These results thus indirectly support that both the finger and wrist priors were stationary during the second half of the experiment, as assumed by our data analyses and model comparisons.

Apart from the above-mentioned within-trial analysis, we also conducted a trial-by-trial analysis to explore the cross-trial correlations. In the σ_S condition (as there were no end point feedback for the other conditions) for each effector, we performed a linear regression of the midpoint shift in the next trial as a function of the prediction error in the current trial. The regression slope gives us a measure of the reliance on end point visual feedback across trials. Similar to our Bayesian results, we found that the absolute value of regression slope for the wrist (-0.16 ± 0.04) was significantly larger than that for the finger (0.03 ± 0.04 , $t_7 = 4.00$, $P < 0.01$). Thus, both for within-trial and cross-trial analyses, subjects tend to rely more on the visual feedback in the wrist than finger condition.

DISCUSSION

Our study shows that both effectors can learn a novel visuomotor mapping of a random amplitude, although the wrist performs the task with more motor noise than the finger. Both effectors can learn the mean of the prior accurately, and they




Fig. 6. Performance of 4 different models *M1*–*M4*. **A**: Akaike information criterion (AIC) values are plotted separately for the 4 models. The lower the value, the better the model performance. **B**: model residuals are plotted separately for the 4 models. Error bars denote SE. * $P < 0.05$. n.s., not significant.

show less reliance on the likelihood information as feedback uncertainty increases, in line with the predictions of Bayes-optimal behavior. However, although the task is computationally equivalent for the two effectors, the wrist exhibits more reliance on the new feedback than the finger in all the conditions with imposed feedback perturbations, suggesting that it is associated with a higher prior uncertainty. Through model comparisons, we find that the effector-specific prior is the most probable explanation, and the results support that prior uncertainty is significantly higher for the wrist than for the finger. These results could also be supported by the cross-trial analysis. Together, our findings suggest that the brain does not use a single and veridical prior for the same type of task that is executed by different effectors or, alternatively, that behavior is not according to a simple Bayesian model.

As for the model selection, although the learning task seems like a problem of dynamic state estimation, the task is simplified as static cue combination on each trial. Similar to the idea formulated in our 2004 paper (Körding and Wolpert 2004), this is a reasonable assumption given that our subjects make discrete motor decisions (where to direct the hand in the second half of the reach) based on briefly presented visual feedback at the half distance of a reach. We agree that the prior is probably learned over the trials (i.e., a dynamic process), but for each trial this prior is to be combined with the likelihood, and this simple integration is best modeled as a static cue combination case. Moreover, we agree that this finding can be explained in different frames, including failure of Bayesian transfer (Maloney and Mamassian 2009) or effector-specific learned compensation. It is just a matter of wording: we use Bayesian framework to model the data and examine within-trial learning, and thus effector-specific prior is the chosen term in this regard. The problem is simply that if we leave the normative space, there is a huge space of alternative models.

Empirical studies that demonstrated Bayes-like motor behaviors typically used simple or familiar motor tasks, including finger pointing (Körding and Wolpert 2004; Tassinari et al. 2006), a one-dimensional interceptive task without timing constraint (O'Reilly et al. 2013), timing behavior with key press (Jazayeri and Shadlen 2010; Miyazaki et al. 2005; Wing and Kristofferson 1973), spatial localization on a computer screen by rotating a dial (Berniker et al. 2010a; Sato and Körding 2014; Vilares et al. 2012), and grasping (Schlicht and Schrater 2007). For these well-practiced tasks, the Bayesian model typically assumes that motor execution is simple and straightforward, fulfilling the role of expressing our state estimates without biases. In the current study we show that for an unfamiliar wrist rotation task, people can still combine the sensory feedback with the prior information as reflected by decreasing feedback weights with increasing feedback uncertainty (Fig. 5B). In contrast to hand reaches, which are quite natural and common in our daily life, we seldom reach for a goal (or more specifically, to control a cursor) using wrist rotation. Thus these findings appear to support that even for unfamiliar movements, people can still integrate prior and likelihood information in a meaningful way (Stevenson et al. 2009). However, when we compare subjects' performance between the wrist condition and the finger condition, we find that the findings contradict the assumptions of simple Bayesian-optimal models.

Although subjects exhibited Bayes-like behavior in both effector conditions, they used the new information differently. Weight of feedback in the wrist task was larger than that in the finger task for all the feedback conditions. Because the visual feedback was the same between the two conditions, subjects should use the likelihood information identically if movement performance is solely driven by new feedback. Thus it is likely that subjects have different degrees of prior precision for different effectors when learning to make within-trial compensation of the shift. By evaluating different Bayesian models, we find that the models assuming effector-specific priors can best explain our data. The estimated prior uncertainty is higher for the wrist than for the finger, consistent with the difference in weight of feedback. Note the variances of the priors we estimated from model fitting are different from the imposed variance of 1 cm²; the variance for the finger case is lower, and that for the wrist case is higher, than the imposed variance. The actual prior uncertainty should be affected by this imposed variance as well as the natural variability when the task is performed without perturbations. We speculate that the prior difference between effectors may result from the higher uncertainty of performing this visuomotor task with the wrist. The indirect evidence to support this speculation is the higher motor variability in the wrist condition. The variance of the lateral deviation in the no-feedback condition (σ_{∞} condition) indicates the movement variability without the influence of visual perturbations. We indeed find that the wrist effector is associated with larger variance in this case. Apart from the motor noise, it is possible that the learning of the mapping (the initial learning of the 1-cm rightward shift over the 1,000 trials) between wrist rotation angle and planar cursor movement is more difficult and less straightforward than the finger's position-to-cursor mapping: there is more planning noise for the wrist condition. Thus differences in both motor noise and planning noise could possibly lead to more prior uncertainty, which are learned over trials. Admittedly, the movement variability among no-feedback trials can only serve as a proxy of the prior uncertainty. Future experiments should use more stringent psychophysics methods (e.g., two-alternative forced choice procedure) to precisely measure the prior uncertainty.

The distinct reliance on feedback between the wrist and the finger is unlikely to be caused by a difference in likelihood. First, the likelihood uncertainty in our task hinges on imposed visual perturbations, which are the same across two effector conditions. Second, if the larger weight of feedback for the wrist were caused by its distinct likelihood, we would need to assume that the wrist is associated with better sensory precision and thus less-variable posterior than the finger. Less variable posterior for the wrist is not supported by our data, because we observed larger motor variability for the wrist than for the finger in all our experimental conditions. Third, larger motor noise is unlikely to lead to less likelihood uncertainty in the wrist condition. To model motor noise, the natural way is to use a Kalman-filter approach to estimate the trial-by-trial changes. Given the large motor noise, the Kalman model would predict that the wrist condition is associated with larger visual noise and thus lead to smaller feedback weight. However, this is the opposite of what we find. Hence, we conclude that a difference in likelihood is unlikely to account for the difference in learning between effectors.

Another perspective to view the role of the prior is causal inference of movement errors (Kong et al. 2017; Wei and Körding 2009). In our task, the midpoint feedback is indicative of where the hand or rotated wrist is, and it might be discrepant from the expected location. This serves as a movement error, which leads to changes in the second half of the reach. The correction of the movement in reaction to the feedback has been associated with causal attribution of the error: if the error originates from poor motor control, such as incorrect motor commands due to muscle fatigue, the motor system should correct it in subsequent actions. If the error is regarded as caused by irrelevant causes such as random external perturbations, the motor system should ignore it entirely for subsequent actions (Berniker and Körding 2008; Kong et al. 2017; Wei and Körding 2009). Previous studies have found supporting evidence for this causal inference, which is probabilistic in nature given the uncertainty in various estimates (Berniker et al. 2010b; Shams and Beierholm 2010). For example, one of our previous studies showed that reaching toward a distant target, compared with a near target, is associated with larger movement variability (Wei and Körding 2009). Critically, when reaching to the distant target, people rely more on the feedback in trial-to-trial learning as manifested by more next-trial correction to the error caused by external perturbations. This is similar to our present finding that the wrist condition, with more motor noise, also exhibits more reliance on likelihood though for within-trial learning. In this previous research, the relation between movement variance and learning is explained by a probabilistic model based on causal inference of movement errors: errors will be corrected less if they are attributed to external, thus irrelevant, causes to the sensorimotor system, and errors will be corrected more if they are attributed to internal, thus relevant, causes to the sensorimotor system. During reaching to the distant target, the increased motor variability makes the errors more likely to be caused by the sensorimotor system, because the errors are within the normal range of unperturbed movements and thus lead to more error correction. From this perspective, if people have larger prior uncertainty of visuomotor mapping for the wrist condition, they may have more difficulty in differentiating whether the error is caused by poor motor control or by a random external perturbation. Particularly, even when they experience a large error, they tend to attribute the error to their own body and fully or even overly compensate for it. On the other hand, if people have less uncertainty of visuomotor mapping in the finger case, they find it easier to discern that the error is caused by an external perturbation and tend to ignore the current feedback information. This explanation can naturally explain the effector difference in weight of feedback. Furthermore, whereas the previous investigation suggests that movement amplitude affects prior uncertainty, the present study suggests that the choice of effectors can also affect prior uncertainty and its associated learning behaviors.

Will the prior difference between effectors be reduced or even eliminated if we give more learning trials to the subject? In our experiment, we analyzed people's learning behavior after they experienced a prior for ~1,000 trials. However, it is still possible that an extended period of learning can lead to a better acquisition of prior because such people can use a single prior for the two effectors. How much time is needed to acquire a prior is poorly understood, and researchers assume either a

prolonged timescale (i.e., life-long experience) or a short timescale (i.e., an experiment session) shapes a prior. The exact duration that is needed for building a prior warrants further investigation, and we postulate that it will be task or context dependent (Raviv et al. 2012).

Bayes-optimal behaviors have been viewed as evidence that humans act as ideal observers in sensorimotor tasks. Several recent studies suggest that prior learning is not purely informational but largely depends on the learning context (Chambers et al. 2017; Fernandes et al. 2014). Particularly, one time perception study showed that priors with different distributions could be shared across effectors (Roach et al. 2017). However, we have taken a step further: our study shows a single distribution of prior could be learned differently across different effectors. Our study highlights that investigations of information-theoretical process of sensorimotor tasks should take seemingly trivial aspects of motor execution, such as the effector to be used, into consideration. Even when experimentally imposed feedback parameters are set to be the same between similar motor tasks, the actual learning behavior and sensorimotor prior is effector dependent. This result challenges the naive Bayesian estimation of priors, which suggests the priors should not change with the effector that executes the motor plan. Though not explicitly discussed, Bayesians generally assume that there is a single prior that incorporates what we “know” about the world. Our study suggests that this is oversimplifying the case. Our finding also calls for caution in using the experimentally imposed prior as the actual prior, especially for sensorimotor tasks that critically depend on movement contexts, and may suggest the need for caution in painting an overarching Bayesian view of human behavior.

GRANTS

This study was funded by National Natural Science Foundation of China Grants 31371020, 31622029, and 61533001, Beijing Municipal Science & Technology Commission Grant Z171100000117015, Beijing Natural Science Foundation Grant 5194024, and National Institutes of Health Grant 5R01NS063399. The funding bodies had no role in the design of the study or in the collection, analysis, and interpretation of data or in writing the manuscript.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

C.Y. and K.P.K. conceived and designed research; C.Y. and K.P.K. performed experiments; C.Y., H.W., and K.W. analyzed data; C.Y., H.W., K.W., and K.P.K. interpreted results of experiments; C.Y. and H.W. prepared figures; C.Y. and K.W. drafted manuscript; C.Y., H.W., K.W., and K.P.K. edited and revised manuscript; C.Y. and K.P.K. approved final version of manuscript.

REFERENCES

- Bejjanki VR, Knill DC, Aslin RN. Learning and inference using complex generative models in a spatial localization task. *J Vis* 16: 9, 2016. doi:[10.1167/16.5.9](https://doi.org/10.1167/16.5.9).
- Berniker M, Körding K. Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci* 11: 1454–1461, 2008. doi:[10.1038/nn.2229](https://doi.org/10.1038/nn.2229).
- Berniker M, Voss M, Körding K. Learning priors for Bayesian computations in the nervous system. *PLoS One* 5: e12686, 2010a. doi:[10.1371/journal.pone.0012686](https://doi.org/10.1371/journal.pone.0012686).

- Berniker M, Wei K, Körding K. *Bayesian Approaches to Modeling Action Selection*. Cambridge: Cambridge University Press, 2010b.
- Chambers C, Sokhey T, Gaebler-Spira D, Körding KP. The integration of probabilistic information during sensorimotor estimation is unimpaired in children with cerebral palsy. *PLoS One* 12: e0188741, 2017. doi:10.1371/journal.pone.0188741.
- Fernandes HL, Stevenson IH, Körding KP. Generalization of stochastic visuomotor rotations. *PLoS One* 7: e43016, 2012. doi:10.1371/journal.pone.0043016.
- Fernandes HL, Stevenson IH, Vilares I, Körding KP. The generalization of prior uncertainty during reaching. *J Neurosci* 34: 11470–11484, 2014. doi:10.1523/JNEUROSCI.3882-13.2014.
- Geisler WS, Kersten D. Illusions, perception and Bayes. *Nat Neurosci* 5: 508–510, 2002. doi:10.1038/nn0602-508.
- Jazayeri M, Shadlen MN. Temporal context calibrates interval timing. *Nat Neurosci* 13: 1020–1026, 2010. doi:10.1038/nn.2590.
- Knill DC, Richards W (Editors). *Perception as Bayesian Inference*. Cambridge: Cambridge University Press, 1996. doi:10.1017/CBO9780511984037.
- Kong G, Zhou Z, Wang Q, Körding K, Wei K. Credit assignment between body and object probed by an object transportation task. *Sci Rep* 7: 13415, 2017. doi:10.1038/s41598-017-13889-w.
- Körding K. Decision theory: what “should” the nervous system do? *Science* 318: 606–610, 2007. doi:10.1126/science.1142998.
- Körding KP, Wolpert DM. Bayesian integration in sensorimotor learning. *Nature* 427: 244–247, 2004. doi:10.1038/nature02169.
- Maloney LT, Mamassian P. Bayesian decision theory as a model of human visual perception: testing Bayesian transfer. *Vis Neurosci* 26: 147–155, 2009. doi:10.1017/S0952523808080905.
- Mamassian P, Knill DC, Kersten D. The perception of cast shadows. *Trends Cogn Sci* 2: 288–295, 1998. doi:10.1016/S1364-6613(98)01204-2.
- Miyazaki M, Nozaki D, Nakajima Y. Testing Bayesian models of human coincidence timing. *J Neurophysiol* 94: 395–399, 2005. doi:10.1152/jn.01168.2004.
- O’Reilly JX, Jbabdi S, Rushworth MF, Behrens TE. Brain systems for probabilistic and dynamic prediction: computational specificity and integration. *PLoS Biol* 11: e1001662, 2013. doi:10.1371/journal.pbio.1001662.
- Raviv O, Ahissar M, Loewenstein Y. How recent history affects perception: the normative approach and its heuristic approximation. *PLoS Comput Biol* 8: e1002731, 2012. doi:10.1371/journal.pcbi.1002731.
- Roach NW, McGraw PV, Whitaker DJ, Heron J. Generalization of prior information for rapid Bayesian time estimation. *Proc Natl Acad Sci USA* 114: 412–417, 2017. doi:10.1073/pnas.1610706114.
- Russell S, Wefald E. Principles of metareasoning. *Artif Intell* 49: 361–395, 1991. doi:10.1016/0004-3702(91)90015-C.
- Sato Y, Körding KP. How much to trust the senses: likelihood learning. *J Vis* 14: 13, 2014. doi:10.1167/14.13.13.
- Schlicht EJ, Schrater PR. Effects of visual uncertainty on grasping movements. *Exp Brain Res* 182: 47–57, 2007. doi:10.1007/s00221-007-0970-8.
- Shams L, Beierholm UR. Causal inference in perception. *Trends Cogn Sci* 14: 425–432, 2010. doi:10.1016/j.tics.2010.07.001.
- Stevenson IH, Fernandes HL, Vilares I, Wei K, Körding KP. Bayesian integration and non-linear feedback control in a full-body motor task. *PLOS Comput Biol* 5: e1000629, 2009. doi:10.1371/journal.pcbi.1000629.
- Stocker AA, Simoncelli EP. Noise characteristics and prior expectations in human visual speed perception. *Nat Neurosci* 9: 578–585, 2006. doi:10.1038/nn1669.
- Tassinari H, Hudson TE, Landy MS. Combining priors and noisy visual cues in a rapid pointing task. *J Neurosci* 26: 10154–10163, 2006. doi:10.1523/JNEUROSCI.2779-06.2006.
- Trommershäuser J, Körding K, Landy MS (Editors). *Sensory Cue Integration*. New York: Oxford University Press, 2011. doi:10.1093/acprof:oso/9780195387247.001.0001.
- Vilares I, Howard JD, Fernandes HL, Gottfried JA, Körding KP. Differential representations of prior and likelihood uncertainty in the human brain. *Curr Biol* 22: 1641–1648, 2012. doi:10.1016/j.cub.2012.07.010.
- Wei K, Körding K. Relevance of error: what drives motor adaptation? *J Neurophysiol* 101: 655–664, 2009. doi:10.1152/jn.90545.2008.
- Wing AM, Kristofferson AB. Response delays and the timing of discrete motor responses. *Percept Psychophys* 14: 5–12, 1973. doi:10.3758/BF03198607.