

# The Generalization of Prior Uncertainty during Reaching

Hugo L. Fernandes,<sup>1,2,3</sup> Ian H. Stevenson,<sup>4</sup> Iris Vilares,<sup>1</sup> and Konrad P. Kording<sup>1,5,6</sup>

<sup>1</sup>Department of Physical Medicine and Rehabilitation, Northwestern University and Rehabilitation Institute of Chicago, Chicago, Illinois 60611, <sup>2</sup>PDBC, Instituto Gulbenkian de Ciência, Oeiras 2780, Portugal, <sup>3</sup>Instituto de Tecnologia Química e Biológica, Universidade Nova de Lisboa, Oeiras 2780, Portugal, <sup>4</sup>Redwood Center for Theoretical Neuroscience, University of California, Berkeley, California 94720, <sup>5</sup>Department of Physiology, Northwestern University, Chicago, Illinois 60611, and <sup>6</sup>Department of Engineering Sciences and Applied Mathematics, Northwestern University, Evanston, Illinois 60208

Bayesian statistics defines how new information, given by a likelihood, should be combined with previously acquired information, given by a prior distribution. Many experiments have shown that humans make use of such priors in cognitive, perceptual, and motor tasks, but where do priors come from? As people never experience the same situation twice, they can only construct priors by generalizing from similar past experiences. Here we examine the generalization of priors over stochastic visuomotor perturbations in reaching experiments. In particular, we look into how the first two moments of the prior—the mean and variance (uncertainty)—generalize. We find that uncertainty appears to generalize differently from the mean of the prior, and an interesting asymmetry arises when the mean and the uncertainty are manipulated simultaneously.

**Key words:** Bayesian; generalization; prior; uncertainty; visuomotor rotation

## Introduction

In many sensorimotor tasks, people make use of prior information that allows perception and movement to be more accurate (Körding and Wolpert, 2004; Tassinari et al., 2006). In Bayesian statistics the prior reflects information accumulated from previous experience, which is then combined with incoming sensory feedback (the likelihood). As we interact with the world, we learn about its statistics (e.g., means and variances) and incorporate this information into our priors. However, since we are never in the same situation twice, we must use past information from different but similar situations to derive the right prior beliefs for a specific task (Shepard, 1987). Only by generalizing from past situations to our current one can we calculate what to expect.

In asking how humans generalize priors it is essential to understand how we represent uncertainty. There are a number of models of how the nervous system might represent uncertainty (Vilares and Kording, 2011). However, there is limited experimental evidence to constrain these models. Electrophysiological experiments have probed how single neurons represent movement-related variables (Georgopoulos et al., 1992; Moran and Schwartz, 1999), but little is known about the representation of uncertainty in sensorimotor tasks (Cisek and Kalaska, 2005; Rickert et al., 2009). Furthermore, to our knowledge, none of the

theoretical models for neural representations of uncertainty makes any predictions for the generalization of priors.

One way of characterizing the generalization of priors comes from previous generalization experiments in motor control (Shadmehr, 2004). During center-out reaching, training with a rotational perturbation (the imposed prior) in one learning direction biases movements to nearby targets, and this bias decreases with increasing distance from the learning direction (Krakauer et al., 2000). Previous studies showed that uncertainty has little effects on this generalization pattern (Fernandes et al., 2012). However, how uncertainty itself generalizes is unknown.

Here, we studied the generalization of prior uncertainty by imposing a noisy visuomotor rotation (the prior) during center-out reaching movements in one direction (learning direction), and manipulating the mean and variance of this rotation. After learning, we examined subjects' movements in other directions and measured subjects' uncertainty by probing their reliance on feedback (the likelihood). This paradigm allows us to assess how subjects adapt to the average visuomotor perturbation (prior mean), how subjects adapt to variability in the visuomotor perturbation (prior uncertainty), and how they generalize these quantities to new targets (prior generalization). In a first experiment we manipulated the variance without changing the mean. As observed with the mean in standard rotational generalization, we found a strong local effect where subjects' prior uncertainty peaks in the learning direction, i.e., subjects had a more complete generalization of the variance of the perturbation in movements close to the learning direction. However, unlike the mean, we found a global effect in the generalization of uncertainty. In subsequent experiments we manipulated the variance while introducing a nonzero mean perturbation and observed interesting nonlinear interactions between mean and variance—subjects had the highest prior uncertainty not in the learning direction but in a neighboring direction.

Received Sept. 10, 2013; revised July 16, 2014; accepted July 18, 2014.

Author contributions: H.L.F., I.H.S., and K.P.K. designed research; H.L.F. and I.V. performed research; H.L.F., I.H.S., I.V., and K.P.K. analyzed data; H.L.F., I.H.S., I.V., and K.P.K. wrote the paper.

H.L.F. was supported by Fundação para a Ciência e Tecnologia SFRH/BD/33525/2008. This work was supported by National Institutes of Health 1R01NS063399 and 2P01NS044393. We thank Daniel Acuna, Max Berniker, and Ben Walker for helpful discussions.

The authors declare not competing financial interests.

Correspondence should be addressed to Hugo L. Fernandes, Rehabilitation Institute of Chicago, Northwestern University, 345 East Superior Street, Attention: Kording Lab Room 1479, Chicago, IL 60611. E-mail: hugoguh@gmail.com.

DOI:10.1523/JNEUROSCI.3882-13.2014

Copyright © 2014 the authors 0270-6474/14/3411470-15\$15.00/0

## Materials and Methods

**Ethics statement.** The study and all experimental protocols were approved by the Northwestern University Institutional Review Board and are 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 participants.

**Experimental protocol: general.** Forty-eight right-handed, healthy subjects (17 male, 31 female; aged  $28.2 \pm 3.6$  years) participated in the experiments;  $n = 32$  in Experiment 1 and  $n = 8$  in each of Experiments 2 and 3. All were naive to the purpose of the experiments, and were paid according to their performance. Subjects made center-out reaches in a  $150 \times 150$  mm workspace. They controlled the position of a cursor with their right index finger, which was recorded using an Optotrak 3D Investigator Motion Capture System. A projector and mirror system was calibrated such that visual feedback was perceived as being in the movement plane (Fig. 1A; Fernandes et al., 2012), and the subject's view of his hand was blocked by the mirror.

The task was designed to measure how subjects generalize a learned variance (and the learned mean) of a noisy visuomotor rotation, that is, how the uncertainty related to a perturbation learned during movements into one direction (the learning direction) affects subsequent movements into other directions (the generalizing directions). The experiments combined two previously existing paradigms: one that allows measurement of generalization of the mean of perturbations (Krakauer et al., 2000; Fernandes et al., 2012) and another that allows measurement of how uncertain subjects are of a perturbation (the learned variance; Körding and Wolpert, 2004).

Subjects were instructed to make reaches to targets in certain directions (either the learning direction or a generalizing direction). Their task was to move the cursor, which they could control with their right index finger position, and make it reach the target. However, the position of the cursor was generally hidden and a visuomotor rotation (the perturbation) was imposed on the cursor (Figs. 1, 2A, 4A, 8B). When subjects were moving in the learning target direction, they were given endpoint feedback about the true cursor position, and were eventually able to correct endpoint errors in the learning direction, either by learning the mean of the perturbation (if there was a nonzero mean) or by learning the variance of the perturbation and how to use the midpoint feedback information (the likelihood, see below) to correct their movements. Afterward, they were instructed to make movements into other directions (the generalizing directions) to measure the generalization patterns of the learned mean (if there was a mean to adapt to) and of the learned variance of the perturbation. Generalization patterns were assessed by analyzing how subjects combine their previous knowledge about the distribution of perturbations (the prior) with the feedback that they receive midway through the movement (the likelihood) about the current true position of the hidden cursor (Körding and Wolpert, 2004). The ideal way to combine these two sources of information is to combine the means of the prior and likelihood, weighted by their relative precision (the inverse of the variance). Assuming that subjects combine this information optimally we can measure their relative uncertainty by computing the slope of a linear regression of the negative of final hand position (subjects' estimated perturbation) as a function of the perturbation (see Eqs. 1 and 2; Figure 2B). Analogously we are able to simultaneously measure the mean of the prior (see Eq. 3).

To probe uncertainty, we displayed noisy midpoint feedback (the likelihood), which consisted of five red circles identical to the cursor flashed midway through each trial reach. The position of these dots is sampled from an isotropic 2D Normal distribution centered on the true position of the cursor with variance  $\sim 5.1$  mm (chosen empirically to avoid complete reliance on either prior or likelihood, see below). Hence, they give uncertain information about the true position of the cursor. The midpoint feedback is shown already during the Familiarization sub-block. This way subjects get a better idea of how the dots relate to the position of the cursor. We use the final hand position to measure the level of uncertainty that the subject has on the hidden cursor position.

The intuition for Bayesian integration is that, if subjects are very certain about where the hidden cursor is, they will only weakly adjust their

reach following the noisy midpoint feedback. If, however, they are very uncertain about the position of the hidden cursor (for example, because the perturbation applied to it has high variance), then they will tend to make large adjustments following midpoint feedback display to correct their reach accordingly.

Using the nomenclature of the Bayesian framework, the perturbation (the visuomotor perturbation applied to the cursor) is sampled from a distribution with defined mean and variance. Throughout the experiment subjects learn the distribution of perturbations, and this learned distribution corresponds to their prior (note that this distribution is not necessarily the same as the distribution of imposed perturbations). The midway flashing dots that give uncertain information about the true position of the red cursor correspond to the likelihood and the subjects' estimated perturbation in a particular trial corresponds to the mean of the posterior. By looking at the slope of the negative of the final hand position (the proxy for subjects' estimated perturbation, see below, Data analysis) as a linear function of the perturbation (Fig. 2B), we can estimate the relative reliance on prior information—relative to midpoint feedback information (Körding and Wolpert, 2004). Hence we can compute a relative measure of subjects' learned and generalized uncertainty.

The learning direction was randomly sampled from one of four diagonal directions and generalization was measured in seven directions displayed at  $180$ ,  $\pm 90$ ,  $\pm 45$ , and  $\pm 25^\circ$  degrees from the learning direction. In the experiment, subjects control the position of a red circle, the cursor ( $\sim 3$  mm radius), with their right index finger. Except for the first familiarization trials the position of the cursor is hidden. Subjects were instructed to make radial reaches from a central blue circle, the starting circle ( $\sim 6$  mm radius), to one of eight yellow circles, the targets ( $\sim 6$  mm radius). Targets were all displayed at a distance of 72 mm from the central blue circle; 300 ms after positioning the cursor over the blue circle, the cursor disappeared, one of the eight targets appeared, and subjects had to reach it. On some of the trials (see below) the final position of the cursor was displayed for 500 ms (endpoint feedback). If the reach was successful, that is, if the center of the red cursor was inside the target, then the target turned white and subjects were rewarded by having a point added to their score. If a successful reach happened in those trials where no information was provided about the success of the reach (no endpoint feedback), then a point was added to a hidden score. To begin the next trial, subjects had to reposition the cursor in the starting blue circle. Except for the familiarization trials where the cursor was always visible, the nonrotated cursor was visible only within a distance of 10 mm from the center of the starting blue circle and disappeared as soon as the trial started, i.e., when the red cursor was inside the central blue circle. Since some subjects occasionally have difficulty finding their way back to the starting blue circle, 4 s after the previous trial was over the nonrotated cursor flashed every second for 50 ms to allow subjects to find the starting position.

The study consisted of Experiments 1, 2, and 3. The experiments differ mainly in that the mean of the imposed perturbations is zero in Experiment 1 and nonzero ( $\pm 30^\circ$ ) in the remaining experiments. Experiments 2 and 3 differ in the way the mean is introduced: in Experiment 2 the mean is  $\pm 30^\circ$  right from the start of each Learning sub-block, while in Experiment 3 the mean is introduced progressively during the first Learning sub-block.

### Experiment 1: generalization of uncertainty under zero mean rotation

The goal of Experiment 1 is to measure the generalization pattern of uncertainty. The experiment begins with an initial Familiarization block (40 trials, five movements to each target) where the cursor is always visible. No rotation was imposed during the Familiarization block. After that, the experiment is divided into two blocks of 720 trials, one for each level of variability (SD:  $4^\circ$  or  $12^\circ$ ). The two blocks differ only in level of variance and their order is pseudorandomized across subjects, so that half of the subjects started with the small prior uncertainty (SD:  $4^\circ$ ) and the other half of the subjects started with the high prior uncertainty block (SD:  $12^\circ$ ).

Each block of 720 trials consists of a Learning and a Testing sub-block. The learning direction is the same for both blocks, but selected randomly for each subject from four possible directions— $\pm 45$  and  $\pm 135^\circ$ . The maximum time to complete each trial is 4 s and there is a minimum time

of 400 ms to complete the second half of the reach. If any of these times is violated the trial is restarted. The minimum time threshold is to guarantee that subjects have enough time to integrate the midpoint feedback information.

**Learning.** Subjects complete 240 trials of movements toward a single learning target direction with midpoint feedback (the cloud of circles flashed midway through the movement) and endpoint feedback (the final position of the cursor is displayed). During each trial, the cursor is hidden and rotated relative to hand position. The rotations applied within each block are sampled from the same normal distribution with mean  $0^\circ$  and SD pseudorandomly chosen to be either  $4^\circ$  or  $12^\circ$ .

**Testing uncertainty.** The Testing Uncertainty sub-block (480 trials) is composed by sequences of four trials. To prevent forgetting of the perturbation, the first two trials of these sequences are toward the learning direction and the other two trials are toward any two of the eight possible directions. Targets are chosen pseudorandomly so that exactly 20 reaches are made toward each generalizing direction (Fig. 2A). Endpoint feedback is provided only in trials toward the learning direction and midpoint feedback is provided in all directions.

### Experiments 2 and 3: generalization of uncertainty under nonzero mean rotation

**Experiment 2.** Experiment 2 is aimed at measuring the generalization pattern of mean and variance when both are perturbed simultaneously. The purpose of Experiment 2 is to see how changing the mean of a perturbation influences the generalization of uncertainty. Hence, the difference between Experiments 1 and 2 is essentially that in Experiment 2 the perturbations have a nonzero mean, specifically a mean rotation of  $\pm 30^\circ$ . The experiment starts with a Familiarization block (40 trials) just like the one in Experiment 1, and it is then divided into two blocks of 880 trials. Just like Experiment 1, in Experiment 2 the two main blocks differ in the magnitude of the SD of the applied perturbation— $4^\circ$  or  $12^\circ$  degrees. The reason for the larger number of trials is that there is an extra sub-block between the Learning (240 trials) and the Testing Uncertainty (480 trials) sub-blocks—the Testing Mean sub-block (160 trials).

**Testing mean.** The Testing Mean sub-block is for measuring the generalization of the mean (160 trials). This sub-block allows us to measure directly how each subject generalized the mean of the perturbation (Fig. 5C,D). Subjects make reaches toward all targets. Endpoint feedback and midpoint feedback are provided only in the learning direction trials. To prevent de-adaptation to the perturbation, the learning target direction is revisited at least twice every four trials; every sequence of four trials is composed of two reaches toward the learning target direction and two reaches toward any two of the eight targets. Targets are chosen pseudorandomly so that there are in total 10 reaches toward each of the generalizing directions. Even though midpoint feedback is not displayed in movements toward the generalizing directions, there is still a minimum amount of time to complete the second part of the movement in every trial. Hence, subjects still slow down halfway through the movement as in the trials where midpoint feedback is displayed.

We can then use these measurements of the generalization of the mean during the Testing Uncertainty sub-block; in each target direction, the perturbation will have a mean equal to how much the subject generalized the learned mean to that direction (as measured in the Testing Mean sub-block; Fig. 5C,D). For example, if subjects did not generalize the learned mean at all in a particular direction then the mean of the imposed perturbation in that direction (during the Testing Uncertainty sub-block) will be zero degrees. Notice that, even though there is no endpoint feedback during Testing Uncertainty, if the mean perturbation does not match the subject's generalized mean then the midpoint feedback could perturb the subject's learned mean and uncertainty, and consequently the measurement of generalization of uncertainty. Hence, by matching the mean of the probing perturbation in each of the generalizing directions with the generalized mean, we minimize the possibility of subjects learning from the midpoint feedback information (see description of the experimental design below for further details regarding this issue).

**Experiment 3.** Experiment 3 aims at understanding how behavior changes when the average of the imposed distribution of perturbations is introduced progressively. The task was thus the same as Experiment 2

except that the mean of the distribution of perturbations is increased progressively from  $0$  to  $30^\circ$  during the first half (linearly over trials 1–120) of the first Learning sub-block (for a comparison of learning curves, see Fig. 8A).

This experiment was motivated by the results observed in Experiment 2 (an asymmetry in the generalization of uncertainty, see Results), which could suggest that generalization of uncertainty has a visual feedback reference frame (see below, Models for generalization). By introducing the mean of the perturbations progressively rather than suddenly we expect visual feedback to generally appear closer to the learning direction during learning. If the hypothesis that generalization of uncertainty has a visual feedback-centered reference frame is true, then we expect the asymmetry to disappear or at least be significantly smaller compared with Experiment 2.

**Experimental protocol: details.** There are two important details to be noticed regarding the experimental design of the experiments. It is not possible to measure a baseline for uncertainty using this protocol. It is not possible to measure a baseline for relative reliance on midpoint feedback due to the fact that we need to introduce a perturbation to measure the slope (Eqs. 1 and 2, Figure 2B). For that reason we measured the generalization of two different levels of variability—SD of 4 and  $12^\circ$ . These SD values were chosen empirically based on several constraints that the task imposes: at the same time that both values need to be sufficiently different, the smaller variance cannot be too small, otherwise the range of the perturbations is not large enough to measure the relative reliance on midpoint feedback (the slope of a linear regression) with a reasonable confidence interval. The higher variance condition cannot be too large otherwise it could introduce nonlinearities (Körding et al., 2007, Wei and Körding, 2009) and because of the constraints inherent to working in a circular support. The SD of the likelihood was chosen empirically so that the slopes would be close to 0.5. This is the range where behavior is influenced equally by prior and likelihood, and thus where fluctuations in uncertainty will have the highest effect. Several values were tested while designing the experiment, starting with the theoretical value that would produce the desired slope, and changing it until values obtained for the slope were  $\sim 0.5$ .

In the generalizing directions, the SD of the perturbation used to probe uncertainty is the same regardless of the SD of the imposed perturbation in the learning direction. Since midpoint feedback is necessary to measure subject's relative uncertainty, special care is needed to ensure that this feedback does not bias measurements of generalization. Differences in learning and sensorimotor integration could both lead to spurious differences in the patterns of relative uncertainty. In all experiments we do not provide endpoint feedback in the generalizing directions and the spread and timing of the midpoint feedback was the same across the two variance conditions. Additionally, we set the variance of the perturbation in the generalizing directions to the geometric mean of the two SDs used in the learning directions, namely  $\sqrt{4 \times 12^\circ}$ . This guarantees that the only difference is the distribution of perturbations between the blocks of trials during movements in the learning direction. The important consequence is that, even if the midpoint feedback allowed subjects to learn during generalization trials, learning would only act to bring the two generalizations curves closer together. The methods used here, thus, set a lower bound on the distance between the generalization patterns for the two variance conditions.

### Data analysis: general

**Final hand position and estimated perturbation.** In this paradigm, the final hand position angle,  $\theta_{\text{h}}$ , gives us a measure of subjects estimated perturbation  $\hat{\theta}$ , specifically:  $\hat{\theta} = -\theta_{\text{h}}$ . Indeed, if the subject estimates a  $\hat{\theta}$  degree rotation then his/her best attempt at hitting the target is by having the final hand position angle be  $-\hat{\theta}$  degrees. We compute the final hand position for each trial by averaging the last data point before the hand goes beyond a distance of 72 mm—the target distance—from the center of the starting circle and the first data point after that. Trials were restarted if subjects did not go beyond the target distance, thus  $\theta_{\text{h}}$  and  $\hat{\theta}$  are defined for every trial.

**Measuring generalization of uncertainty: relative reliance on midpoint feedback (slope).** We assume that the estimated perturbation corresponds

to the mean of the posterior (Körding and Wolpert, 2004). Assuming Gaussian distributions, an ideal observer/actor would combine information from their prior over cursor perturbations ( $\theta_p$ ) and the perturbation angle sensed from the midpoint feedback information ( $\theta_f$ ) weighting their values by their relative precisions, according to the following:

$$\hat{\theta} = \frac{\sigma_f^2}{\sigma_p^2 + \sigma_f^2} \theta_p + \frac{\sigma_p^2}{\sigma_p^2 + \sigma_f^2} \theta_f. \quad (1)$$

Where  $\sigma_p^2$  and  $\sigma_f^2$  denote subjects' uncertainty in prior and midpoint feedback, respectively. Subjects' estimated angle of the perturbation ( $\hat{\theta}$ ) applied to the cursor is reflected in the angle of their final hand position. This equation is an approximation of Bayesian integration with circular variables, but works well when the perturbation and feedback variances are small as they are here. As a proxy for the sensed cursor perturbation angle ( $\theta_f$ ), we use the centroid of the sampled flashing dots of the likelihood. That is, instead of considering the real perturbation angle we considered the angle that a vector from the center of the central blue circle to the centroid of the flashing dots would do with the target direction if the subject had moved straight to the target. Importantly, this equation (Eq. 1) allows us to estimate the relative learned/generalized variance of the prior for each learning/generalizing direction, and to compute a relative generalization function for uncertainty. Indeed, for the reaches data in each variance condition and direction we can perform a linear regression (Fig. 2B) of the estimated perturbation ( $\hat{\theta}$ ) as a function of the corrected perturbation angle ( $\theta_f$ ). The value of the slope ( $s_p$ ) of that linear regression gives us a measure of relative reliance on midpoint feedback/likelihood. Importantly, if subjects are behaving in a Bayes-optimal way (see Eq. 1), then the slope should be equal to the following:

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

Hence, the slope gives us a relative measure of a subject's prior uncertainty for each variance condition and direction. If a subject's prior uncertainty is much higher than the likelihood uncertainty, then this slope value will approach one, i.e., the subject will rely almost exclusively on midpoint feedback. If, on the other hand, the prior is much less uncertain than the midpoint feedback information, then this slope will approach zero, and subjects should almost not rely on the midpoint feedback information. This slope, the relative reliance on midpoint feedback, serves as the basis for most of our analysis. We use the median of 1000 bootstrap samples to reduce the influence of outliers when computing the slopes. Specifically, if for a given target direction we have data from  $n$  trials in the Testing Uncertainty sub-block, we compute the slope on 1000 sets of size  $n$  randomly drawn with replacement from the original  $n$  trials. The median of these 1000 values will be our estimate of the slope, and by computing the SD of these 1000 values we obtain an estimate of the SEM. The centroid adjustment and bootstrapped slope estimates provide more robust measures of behavior, but using unadjusted perturbations and maximum likelihood slope estimates produce qualitatively very similar results.

### Measuring generalization of the mean

*Inferred mean.* We are able to infer the mean of the prior in both experiments using the data from the Testing Uncertainty sub-block. We do this by computing the intercept of a linear regression; we can rearrange Equation 1 to obtain the following:

$$\theta_f = \theta_p + \frac{\sigma_f^2 + \sigma_p^2}{\sigma_f^2} (\theta_f - \hat{\theta}). \quad (3)$$

That is, for each target direction, the subjects' prior mean,  $\theta_p$ , is estimated using the intercept of the linear regression of  $\theta_f$  as a function of  $\theta_f - \hat{\theta}$ .

*Direct measurement of the prior's mean.* In Experiments 2 and 3, during the Testing Mean sub-block, we were able to directly measure generalization of the mean in each of the generalizing directions; during the trials in the generalizing directions of the Testing Mean sub-block, subjects were not shown midpoint feedback (the likelihood) and hence their estimate—as inferred by final hand position—is assumed to be the mean of the prior distribution in that direction. Using this information we were

able to compute, during the experiment, the means of the perturbations used to probe uncertainty in the generalizing directions during the Testing Uncertainty sub-block (see above, Experimental protocol). During trials in the learning direction subjects were still shown midpoint feedback. Thus, their average final hand position during trials toward the learning direction is an estimate of the mean of the posterior and not of the prior. The generalization patterns obtained during the Testing Mean sub-block match very well with the ones inferred using the data from the Testing Uncertainty sub-blocks ( $F_{(1,7)} = 3.27$ ,  $p = 0.11$ , two-way (testing block, direction) repeated measures (subject) ANOVA; Figure 5C, first and second rows for individual subject data,  $D$ ). The higher complexity of this task, relative to previous studies that measured generalization of means (Fernandes et al., 2012), may have led to smaller variability (possibly due to smaller variability in cognitive strategies) across subjects.

*Absolute mean and percentage adaptation to the mean.* In Experiments 2 and 3, since for half of the subjects the mean of the perturbation was  $-30^\circ$ , we normalized the estimated perturbation (as measured by the negative of the angle of final hand position) according to the sign of the mean of the perturbation; we multiplied by  $-1$  the angle of the estimated perturbation if the mean of the perturbation was negative ( $-30^\circ$ ). Hence, a positive absolute mean corresponds to a movement that counteracts the perturbation. We call the measurements of the mean using this normalization the inferred absolute mean and the measured absolute mean, depending on whether they were determined using data from the Testing Uncertainty or from the Testing Mean sub-block, respectively. Using the inferred absolute mean we compute the inferred percentage adaptation to the mean, which is defined as the amount of learned/generalized mean relative to the learned mean in the learning direction (Figure 4E). The inferred percentage adaptation to mean in the learning direction is hence, by definition, 100%.

### Models for generalization

Here we consider on-line learning models that use gradient descent to learn the mean and SD of the imposed prior. For each trial the mean and SD are updated to minimize the expected squared error between the target angle and the final cursor position angle.

Consider the subjects' original prior assuming no visuomotor rotation:

$$p_0 = \mathcal{N}(\theta_p^0, \sigma_p^0) = \mathcal{N}(0^\circ, \sigma_p^0). \quad (4)$$

Where  $\sigma_p^0$  denotes the subject's original prior SD and we assume that there is no bias in the subject's original prior mean, i.e., we assume that  $\theta_p^0 = 0^\circ$ . We assume also that the original prior is the same for all directions. Throughout this section we generally use  $\theta$  to denote perturbation related angles and  $\phi$  to denote target angles. Note that the perturbation angles  $\theta$  are always given relative to that trial's target direction and target angles are given relative to the learning target direction.

We hypothesize that subjects start with their original prior mean ( $\theta_p^0 = 0^\circ$ ) and SD ( $\sigma_p^0$ ), and that in each learning trial (during the Learning sub-blocks) they update their prior mean and SD (using Eqs. 7–10 below). We assume that this learning of the perturbation can generalize to other directions, and that the amount of generalization depends on the similarity between the two trials/contexts (the context in which the perturbation was learned and the new context/trial). The perceived similarity between two particular trials/contexts depends on the reference frame that the subject uses for learning that particular variable (mean or SD): this reference frame can be target based (i.e., trials with similar context are trials toward similar target directions; see Model T and mean update in Model VF below), or it can be a visual-feedback/cursor-based reference frame (i.e., trials with similar contexts are trials that have midpoint visual feedback in a similar position; see SD update in Model VF below).

How similar two trials are perceived by a subject can be defined by a context similarity function. Here we define the context similarity function,  $W_{\phi_r}(\phi_g)$ , for a given generalizing direction  $\phi_g$  relative to the learning reference frame direction  $\phi_r$ , as a scaled von Mises function:

$$W_{\phi_r}(\phi_g) = (b_0 + \exp(b_1 \cos(\phi_g - \phi_r))) / \alpha, \quad (5)$$

where  $b_0$  is a baseline for context similarity,  $b_1$  defines the width of the context similarity, and  $\alpha = b_0 + \exp(b_1)$  is a normalization factor so that context similarity is 1 in the reference frame direction, that is,  $W_{\phi_c}(\phi_c) = 1$ . For instance, if the reference frame is target centered then learning, and thus the peak of generalization, happens in the learning target direction,  $\phi_p$ , i.e.,  $W_{\phi_c}(\phi_p) = 1$  (see Model T below). This is the same as saying that generalization is complete in the learning target direction. If the reference frame is visual feedback centered then learning happens/generalization is maximal in the direction of the centroid of the midpoint visual feedback in that particular trial,  $\phi_c$  (see Eq. 7, the gradient update of SD in Model VF below). Notice that, unlike the target-based reference frame, in the visual feedback-centered reference frame the direction of learning generally changes from trial to trial.

Both models proposed below have the following parameters:  $b_0^{\theta_p}$ ,  $b_0^{\sigma_p}$ ,  $b_1$ , which are the three context similarity function parameters (see Eq. 5), allowing for different context similarity baselines for mean ( $b_0^{\theta_p}$ ) and variance ( $b_0^{\sigma_p}$ ); the initial prior uncertainty  $\sigma_p^0$  and likelihood uncertainty  $\sigma_f$ ; and the learning rates of mean and variance,  $\eta_{\theta_p}$  and  $\eta_{\sigma_p}$ . These seven parameters were enough to produce good fits to the data from the first block. However, we observed that none of the models managed to capture decreases in the variance of the prior during the second block. For this reason, and to account for possible differences between the Learning and the Testing sub-blocks, we added an extra parameter that scales the model's output of prior uncertainty at the end of learning before fitting it to the testing data. Hence both Models T and VF have a total of eight parameters. While Model T is target centered, Model VF tests the hypothesis that generalization of variance has a nontarget-centered reference frame—the visual feedback information.

**Model T: target-centered reference frame for both mean and variance.** On each  $i$ th trial of the Learning sub-block, the subject is trying to minimize the squared error between the target angle and the final cursor position angle, that is, trying to learn the mean  $\theta_p$  and variance  $\sigma_p^2$  of the prior imposed in the learning direction such that

$$(\hat{\theta}_p^{\phi_i}, \hat{\sigma}_p^{\phi_i}) = \arg \min_{\theta_p, \sigma_p} e(\theta^i, \theta_p, \sigma_p), \quad (6)$$

where  $\theta^i$  is the perturbation imposed during the  $i$ th trial and  $e(\theta^i, \theta_p, \sigma_p) = (\theta^i - \hat{\theta}(\theta_p, \sigma_p))^2$  where  $\hat{\theta}(\theta_p, \sigma_p)$  is defined in Equation 1. The model takes as input the learning trials and assumes that the SD and mean of the subject's prior evolve according to the following:

$$\sigma_p^{\phi_i, i} = \sigma_p^{\phi_i, i-1} - \eta_{\sigma_p} W_{\phi_i}(\phi_g) \frac{\partial e}{\partial \sigma_p} \Big|_{\theta_p^{\phi_i, i-1}, \sigma_p^{\phi_i, i-1}} \quad (7)$$

$$\theta_p^{\phi_i, i} = \theta_p^{\phi_i, i-1} - \eta_{\theta_p} W_{\phi_i}(\phi_g) \frac{\partial e}{\partial \theta_p} \Big|_{\theta_p^{\phi_i, i-1}, \sigma_p^{\phi_i, i-1}} \quad (8)$$

Notice in Equations 7 and 8 that the reference frame is the learning target direction,  $\phi_i$ ; the context similarity function,  $W_{\phi}(\cdot)$ , is computed relative to the learning direction and the gradient is evaluated in both equations at the current mean ( $\theta_p^{\phi_i, i-1}$ ) and SD ( $\sigma_p^{\phi_i, i-1}$ ) of the prior in the learning direction.

**Model VF: different reference frames, target centered for mean but visual feedback centered for variance.** The only difference between the Models T and VF is that in Model VF the reference frame for uncertainty in the prior (SD,  $\sigma_p$ ) is centered on the angle of the centroid of the displayed cloud of dots,  $\phi_c$ , while the context function for the mean remains centered on the learning target direction,  $\phi_i$ :

$$\sigma_p^{\phi_i, i} = \sigma_p^{\phi_i, i-1} - \eta_{\sigma_p} W_{\phi_c}(\phi_g) \frac{\partial e}{\partial \sigma_p} \Big|_{\theta_p^{\phi_i, i-1}, \sigma_p^{\phi_i, i-1}} \quad (9)$$

$$\theta_p^{\phi_i, i} = \theta_p^{\phi_i, i-1} - \eta_{\theta_p} W_{\phi_i}(\phi_g) \frac{\partial e}{\partial \theta_p} \Big|_{\theta_p^{\phi_i, i-1}, \sigma_p^{\phi_i, i-1}} \quad (10)$$

Notice that in the equation for the gradient update of the SD of the prior (Eq. 9), the context similarity function,  $W_{\phi}(\cdot)$ , is computed relative to  $\phi_c$ . Notice also that in both equations, even though, as in Model T, the gradient is evaluated relative to the current mean of the prior in the

learning direction,  $\theta_p^{\phi_i, i-1}$ , the SD is evaluated in the direction where the cloud of dots appeared,  $\theta_p^{\phi_c, i-1}$ .

To avoid using behavioral data obtained during the Learning sub-block, the model uses the predicted position of the cloud of dots as a proxy for  $\phi_c$ . This predicted position is obtained by computing, given the trial perturbation, where the cloud of dots would appear if the subject performed a straight center-out movement corrected by the current mean of the prior. Equivalent results were obtained when data from the Learning sub-block, the actual angle of centroid of the cloud of dots, was used. However, using only testing data allows for a fair comparison of all models, and allows us to simulate the models even in the absence of behavioral data.

### Computing the gradient

To compute the partial derivative of the error function,  $e(\theta^i, \theta_p, \sigma_p)$ , we observe that

$$e(\theta^i, \theta_p, \sigma_p) = (\theta^i - \hat{\theta}(\theta_p, \sigma_p))^2 = (\theta^i - s_p \theta_f - (1 - s_p) \theta_p)^2,$$

where  $s_p = \sigma_p^2 / (\sigma_p^2 + \sigma_f^2)$  and  $\theta$  is defined in Equation 1.

Applying the chain rule we obtain the partial derivatives:

$$\frac{\partial e}{\partial \sigma_p} = 2 \left( \theta^i - \theta_p + \frac{\sigma_p^2}{\sigma_p^2 + \sigma_f^2} (\theta_p - \theta) \right) (\theta_p - \theta) \frac{\sigma_f^2}{(\sigma_p^2 + \sigma_f^2)^2} 2 \sigma_p$$

$$\frac{\partial e}{\partial \theta_p} = -2 (\theta_p - (1 - s_p) \theta_p - s_p \theta_f) (1 - s_p),$$


which are then used in the gradient update Equations 7–10.

### Model fitting

We fitted the models to the slope and mean data of each subject by minimizing the squared distance to the subject's slope and mean in each direction weighted by the precision [inverse variance, obtained using bootstrapping, see above, Measuring generalization of uncertainty: relative reliance on midpoint feedback (slope)] of each slope and mean data point. To account for discrepancies between the Learning and Testing sub-blocks, both models have an additional scaling parameter that allows us to fit the output of the learning model to subject's prior uncertainty during testing. To compare models (Fig. 7), we perform a Wilcoxon signed rank test on the weighted root mean square errors (RMSEs).

## Results

Here we ask how a noisy perturbation in one learning direction affects reaches into other directions. In particular, we aim to extend movement generalization studies by understanding how both the mean and the variance of a perturbation imposed during reaches into the learning direction affect other movements. Subjects controlled the position of a hidden cursor with their right index finger while their true hand position was occluded by a projector mirror system (Fig. 1). They made reaches from the workspace center to one of eight concentric targets with a visuomotor rotation applied to the hidden cursor position. The visuomotor rotation was drawn randomly each trial from a Gaussian distribution with fixed mean and variance. During learning subjects were incentivized to make reaches to one of the targets (learning target) and received endpoint feedback that allowed them to adapt to the distribution of perturbations (see Materials and Methods). During testing subjects also made reaches to the other targets (generalizing targets), without endpoint error feedback, allowing us to probe generalization. All subjects went through two sub-blocks of learning, each with a different variance ( $\sigma_p$ : 4° or 12°). We measured how the learned variance generalizes, first without perturbing the mean (mean of 0° in Experiment 1) and then while also perturbing the mean (mean of  $\pm 30^\circ$  in Experiments 2 and 3).



**Figure 1.** Experimental setup and typical trajectory data. **A**, Subjects move a hidden cursor from a starting position to a target (yellow) by moving their occluded right index finger. We measure the generalization of the learned variance of a perturbation using the response to a noisy midpoint cursor feedback (red dots). **B**, Experiment 1, with zero mean perturbation, and a SD ( $\sigma_p$ ) of either  $4^\circ$  or  $12^\circ$  (shown example trajectories are for  $\sigma_p = 12^\circ$ ). Subject's hand and cursor position during learning trials. Each thin red line represents the trajectory of the real hand position (hand, left) or of the hidden cursor position (cursor, right) during a trial toward the learning direction (yellow target). The black solid line represents the average trajectory toward the learning direction. Average trajectories for generalizing directions are shown as black dashed lines (corresponding targets are black dots). The position where midpoint feedback is triggered is denoted by the red circle. Left, Hand trajectories are also superimposed on the workspace shown in **A**. **C**, Experiments 2 and 3, where the absolute mean perturbation is  $30^\circ$  (introduced abruptly in Experiment 2 and progressively in Experiment 3) and the SD is either  $4^\circ$  or  $12^\circ$  (shown example trajectories are from Experiment 2 with  $\sigma_p = 12^\circ$ ). Same notation as **B**. Note that, because the cursor is rotated  $30^\circ$  on average, the subjects' hand has to move on average in a  $-30^\circ$  direction (relative to the target) so that the cursor hits the target.

As subjects adapt to the noisy visuomotor rotations they update their knowledge of both the mean and variance of the perturbations. We can probe subjects' prior uncertainty (i.e., subjects' estimate of the variance of the perturbation in a particular direction) by providing noisy feedback about the cursor position midway through the movement in the form of a cloud of dots (Körding and Wolpert, 2004). Subjects ( $n = 32$  in Experiment 1 and  $n = 8$  per experiment for Experiments 2 and 3) use this midpoint feedback information to correct their movements during each reach (Figs. 1, 2A) and they rely more on feedback the more uncertain they are about the cursor position. Computing the slope (Fig. 2B) of the negative of final hand position angle (proxy for estimated perturbation) as a function of the perturbation angle (proxy for perturbation sensed via midpoint feedback) provides a measure of the uncertainty that subjects have about the expected perturbations (prior uncertainty) relative to the uncertainty about the midpoint feedback (likelihood uncertainty; Körding and Wolpert, 2004). Intuitively we can see that, if subjects are very certain about the perturbation (low prior uncertainty) then they will tend to ignore the noisy midpoint feedback information and the slope will have a value closer to zero. If on the other hand they have a high prior uncertainty relative to the uncertainty in the midpoint feedback, they will tend to rely only on midpoint feedback and hence the slope will have a value of one. For standard Bayesian integration using Gaussian distributions (Körding and Wolpert, 2004), the slope  $s_p$  is given by

$$s_p = \frac{\sigma_p^2}{\sigma_p^2 + \sigma_f^2},$$

where  $\sigma_p^2$  and  $\sigma_f^2$  are the variances of the prior and likelihood distributions, respectively (see Materials and Methods for details). Hence, larger slopes indicate a higher reliance on sensory feedback and higher uncertainty about the perturbations (Fig. 2B). Whether subjects behave according to the predictions of Bayesian integration or not, the slope is a measure of how uncertain they are about the hidden perturbation.

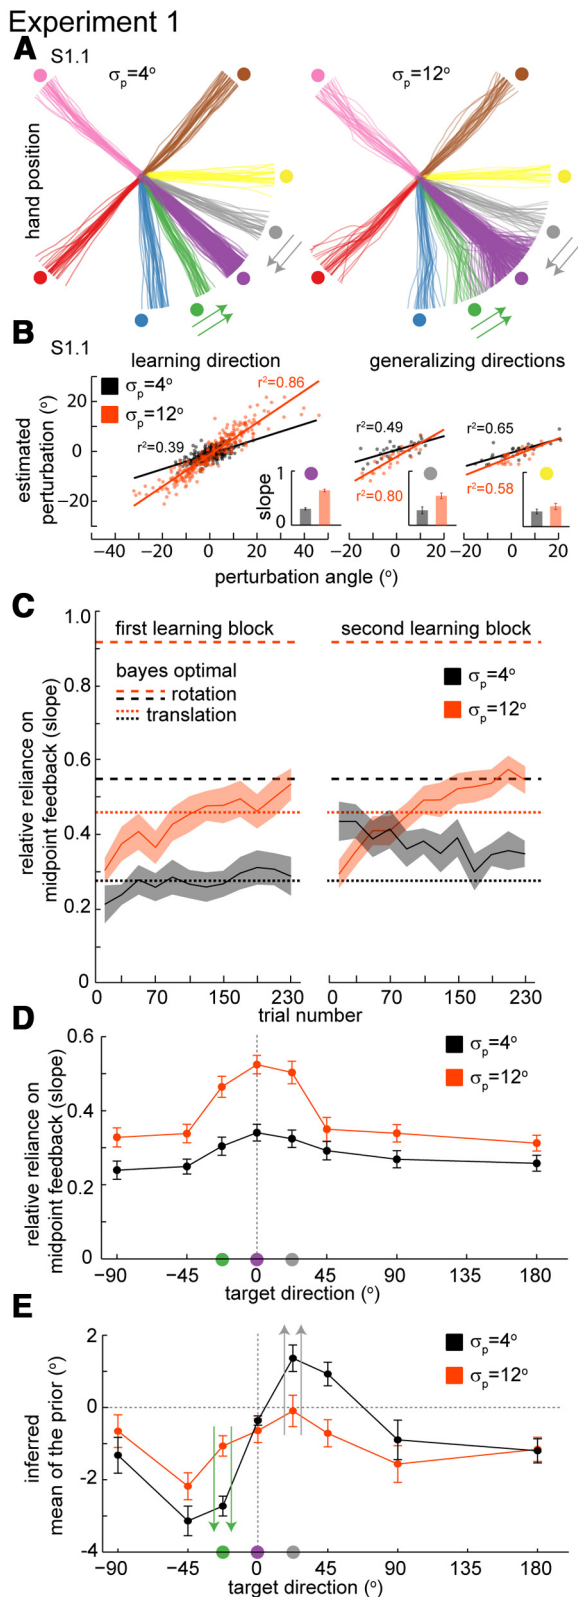
### Experiment 1

We first wanted to know how uncertainty generalizes with a zero mean perturbation. We find that subjects learn about the vari-

ance and exhibit smaller slopes for the small variance condition than for the high-variance condition (Fig. 2C). This is what we should expect since a smaller slope implies that the subject relies less on the midpoint feedback and, hence, that the subject is more certain a priori about the hidden cursor position. Furthermore, learning curves under the same variance condition converge to the same value during the Learning sub-block, no matter which condition subjects started in, and learning appears to asymptote before we assess generalization. Indeed, these two groups of subjects (the ones that started with the high-variance condition and the ones that started with the low-variance condition) did not show a significant difference between their slope values after training ( $F_{(1,472)} = 0.95, p = 0.33$  for main effect of group, four-way nested ANOVA over subject, variance, group, target direction where subject is nested in group). By the end of each Learning sub-block subjects have adapted to the new variance condition. In the following analysis we thus combine data across groups to ask how subjects generalize this learned variance.

Subject exhibiting smaller slopes for the small prior variance condition (Fig. 2C) is qualitatively consistent with "Bayes optimal" behavior. But which convergence values would be optimal? Under the unrealistic assumption that subjects have perfect acuity, attention, and centroid estimation (Tassinari et al., 2006), the Bayesian optimal values for slope are given by Equation 2. In reality, however, if we relax this assumption we expect them to be smaller than the optimal predictions. Furthermore, the slope also depends on subjects' interpretation of the task. If the subjects, instead of perceiving the task as a visuomotor rotation (as it is), perceive it as a visuomotor displacement/translation then, after observing inferring the displacement midway through the movement, they would correct half as much as they should and thus we would measure a slope that is half the theoretical prediction. We find that subjects' slope values at the learning block generally lie between these two (Fig. 2C). This finding may suggest that subjects are uncertain about the kind of applied perturbation (Taylor and Ivry, 2013) or that the subject's actual likelihood (sensory) uncertainty differs from the task uncertainty due to acuity, attention, or bias.

To examine generalization of uncertainty we quantify subjects' reliance on midpoint feedback (as measured with the slope)



**Figure 2.** Experiment 1. *A*, Hand trajectories during the Testing Uncertainty sub-blocks for the two levels of variance (SD of 4 and 12°) in the learning direction (purple) and the generalizing directions (other colors) for one subject (Subject 1.1). Arrows indicate the angular direction of an eventual use-dependent effect (*E*). *B*, Probing uncertainty in the prior by computing the relative reliance on midpoint feedback. Estimated perturbation as a function of the perturbation angle, for one subject (same as in *A*), in the learning direction (left) and in two generalizing directions (+22.5° and +45°; right) during the Testing Uncertainty sub-block. Solid lines denote linear fits to the data. Insets, Slope or relative reliance on midpoint feedback ( $\pm$  SE) during

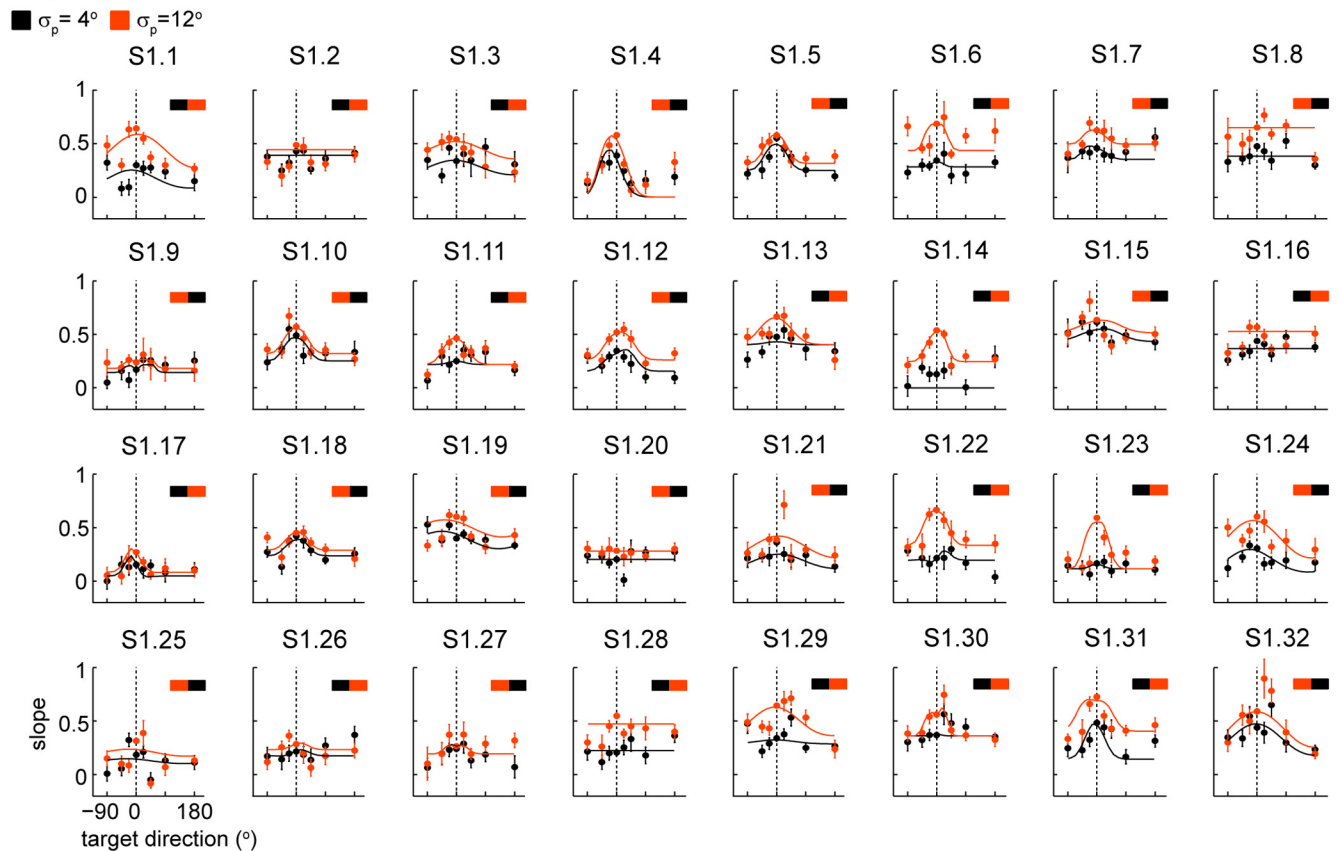
as a function of the direction of movement. We find that the reliance on midpoint feedback (Figs. 2*D*, 3 for individual subject data) is significantly different between the two variance conditions ( $F_{(1,31)} = 65.13, p < 10E-8$ , two-way repeated-measures ANOVA on slope for eight directions and two variance conditions) and slopes for the high-variance condition are higher than those for the low-variance condition for movements into all directions ( $p \leq 0.006$ , for every direction, paired *t* tests,  $n = 32$ ). Uncertainty in the prior (i.e., a subject’s learned variance of the prior, as measured by the relative reliance on midpoint feedback) increases in all directions and decays with increasing distance from the learning direction. Unlike the mean (see Experiments 2 and 3 below; but see Fernandes et al., 2012), uncertainty appears to have a strong global component.

Even though the perturbation had zero mean, we can infer the mean of the subject’s prior by analyzing the intercept of a linear regression using data from the Testing sub-block (see Materials and Methods for details). As with the slopes, we do not find a significant difference between the inferred means in the two groups of subjects ( $F_{(1,472)} = 0, p = 0.97$ , four-way nested ANOVA), and hence we pool the data from both groups in the remaining analyses. We find an interesting asymmetry consistent with use-dependent learning/adaptation theory (Diedrichsen et al., 2010; Huang et al., 2011; Verstynen and Sabes, 2011); in the reaches toward targets that neighbor the learning target and this bias decays with distance from the learning target (Fig. 2*A*, arrows, *E*). Furthermore the bias is stronger in the low-variance condition ( $p < 0.001$  for both 22.5° target directions, paired *t* tests) when movements tend to be less variable and hand position covers a narrower region. We observed signs of a similar effect in a previous study exploring generalization of the mean (Fernandes et al., 2012). Here, in the absence of mean adaptation, there is a clear asymmetry where the large number of subjects and the increased complexity of the task (reduced cognitive strategies and across-subject variability) may make the effect more observable. These results suggest a weak use-dependent learning effect in this experiment.

Experiment 1 is an extension of previous studies that measured the generalization of fixed visuomotor perturbations, but instead of manipulating the mean we have manipulated the variance (Krakauer et al., 2000). In a previous study we showed that the generalization of the mean seems to be unaffected by changes in prior uncertainty (Fernandes et al., 2012). Fully understanding the generalization of uncertainty requires some understanding of how simultaneously perturbing the mean affects the generalization of uncertainty. Experiments 2 and 3 aim to characterize these interactions and differences between generalization of the mean and variance.

←  
movements in that direction for this subject. Colors of circles in the inset indicate which target direction the data correspond to in *A*. *C*, Learning of uncertainty for two groups of subjects: subjects that started with the low-variance condition and subjects that started with the high-variance condition. Solid lines are average slopes ( $\pm$  SEM) across subjects considering bins of 20 trials. Dashed lines correspond to the optimal theoretical values assuming perfect visual acuity and complete adaptation of prior uncertainty. Dotted lines correspond to the optimal behavior if the subjects assume the perturbation is a translation instead of a rotation. *D*, Relative reliance on feedback for the two levels of prior uncertainty as a function of target direction relative to learning direction (0°). Mean ( $\pm$  SEM) of slopes across all subjects. Colors of the circles on the *x*-axis indicate the correspondent directions in *A*. *E*, Inferred mean of prior ( $\pm$  SEM) across all subjects. Arrows indicate the expected direction of an eventual bias caused by use-dependence. Arrow’s colors match the colors in *A*.

## Experiment 1



**Figure 3.** Individual subject slope data for Experiment 1. Error bars  $\pm$  SEM (bootstrap). Lines are Model VF (see above, Models) fits to individual subjects. Black and orange bar (inset) indicates the order in which the different uncertainty blocks were presented to each subject (black denotes the  $4^\circ$  condition and orange denotes the  $12^\circ$  condition).

## Experiment 2

In Experiment 2 our aim is to characterize how the mean of a perturbation affects the generalization of uncertainty. As in Experiment 1, it is important to quantify the effects of the different perturbation variances (low uncertainty:  $\sigma_p = 4^\circ$  or high uncertainty:  $\sigma_p = 12^\circ$ ; Fig. 4A) and to determine whether training order (i.e., high-to-low vs low-to-high variance) matters. The average learning curve for slope across subjects ( $n = 8$ ) seems to asymptote to values similar to the ones observed in Experiment 1 (Fig. 4B), even with the addition of a nonzero mean perturbation. We found that, in block 1, the reliance on midpoint feedback (slope; Fig. 4C) is significantly different between the two variance conditions ( $F_{(1,8)} = 19.72$ ,  $p = 0.02$ , two-way repeated-measures ANOVA on slope for eight directions and two variance conditions). However, in Experiment 2 we found a significant difference in reliance on midpoint feedback between the two groups of subjects after training ( $F_{(1,112)} = 8.7$ ,  $p = 0.004$ , four-way nested ANOVA—subject, variance, group, direction where subject is nested in group; Fig. 4C,D). Because the order of training now matters we will present and analyze the data from the two blocks separately.

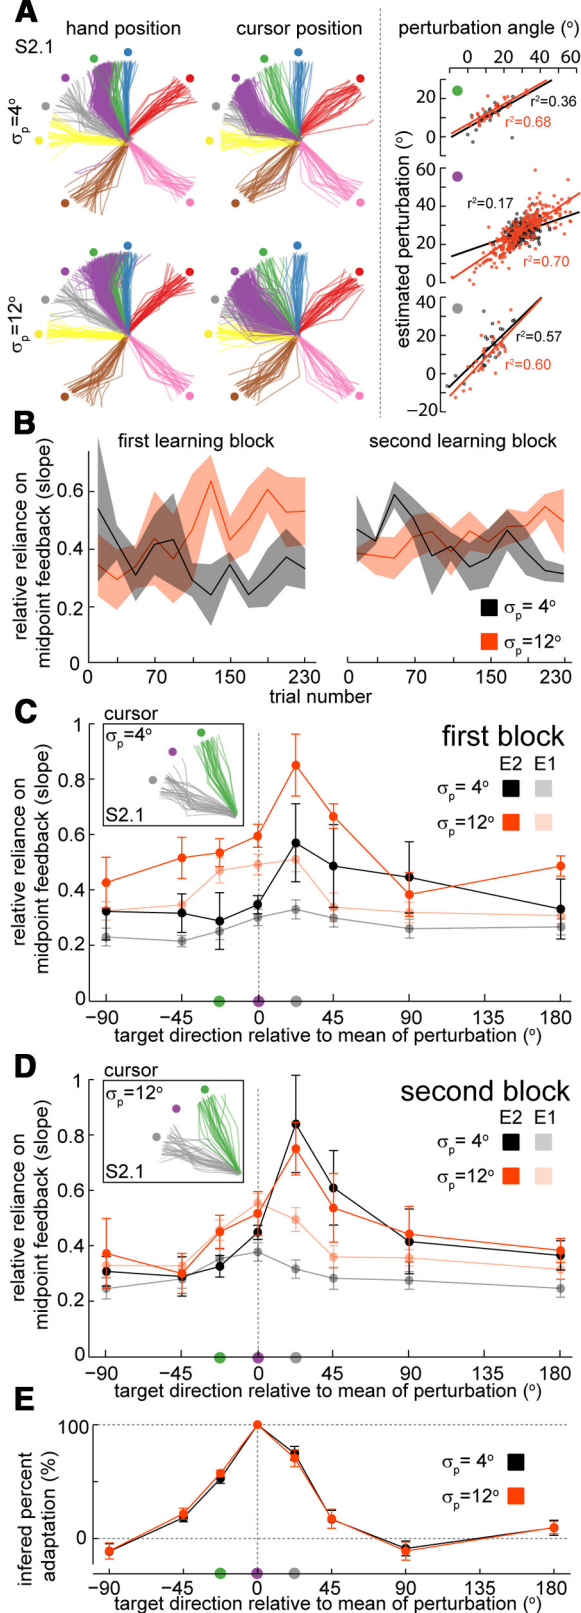
In contrast to Experiment 1, in Experiment 2 we find a strong asymmetry in the generalization of uncertainty (Fig. 4C,D; for individual subject data, see Fig. 5C). The generalized prior uncertainty as measured by the relative reliance on midpoint feedback is higher than expected for the neighboring targets on one of the sides of the learning direction, even higher than the learned uncertainty in the learning direction. These directions of higher

uncertainty correspond to the opposite side (relative to the target direction) to where the hand has to move to correct for the perturbation; that is, they correspond to the direction of the mean of the perturbation (Figs. 4A, 5A). Furthermore, these are the directions where the midpoint visual feedback is more often displayed during early learning (before subjects learned the mean) and correspond to trials with bigger final cursor position errors (Fig. 5B). Furthermore, this asymmetry is observed consistently across subjects and sign of the perturbation (Fig. 5C,E), and seems to be robust to any subject-specific cognitive strategies (Taylor and Ivry, 2013). We find that, when changed simultaneously, the mean and variance of perturbations have asymmetric effects in the generalization of the variance of the prior over those perturbations.

Since we find a surprising asymmetry in the generalization of variance in Experiment 2, we can ask whether manipulating mean and variance simultaneously has a similar effect in the generalization of the mean. We find that the generalization of the mean angular perturbation is local, with a width of about  $30^\circ$ , similar to what has been reported in previous studies (Krakauer et al., 2000; Fernandes et al., 2012; Figs. 4E, 5D). As in previous studies, with similar center-out reaching designs (Fernandes et al., 2012), generalization to targets at a  $\pm 90^\circ$  angular distance from the learning target is not significantly different from zero ( $p > 0.13$  for the  $\pm 90^\circ$  targets in both uncertainty conditions,  $t$  tests). Furthermore, in agreement with Experiment 1, in the directions that neighbor the learning target we observe an asymmetry consistent with use-dependent learning. Note that the use-



**Experiment 2**



**Figure 4.** Experiment 2. **A**, Left, Hand and cursor trajectories of a particular subject (Subject 2.1) during the Testing Uncertainty sub-blocks for the two levels of variance (SD of 4 and 12 $^{\circ}$ ) in the learning direction (purple) and the generalizing directions (other colors). Right, Probing uncertainty in the prior by computing the relative reliance on midpoint feedback. Estimated perturbation as a function of the perturbation angle, for the same subject (Subject 2.1), in the learning direction and in two generalizing directions ( $-22.5$  and  $+22.5^{\circ}$ ) during the Testing

dependent asymmetry, although reflected as an asymmetric generalization pattern in the mean, can be interpreted as movements close to the training direction being attracted by the direction where training occurred (Fig. 2A,E, arrows). It seems that the asymmetry observed in the generalization pattern of the mean is be hand movement related while the asymmetry observed in the generalization of uncertainty seems to be related to overlapping visual feedback.

**Models**

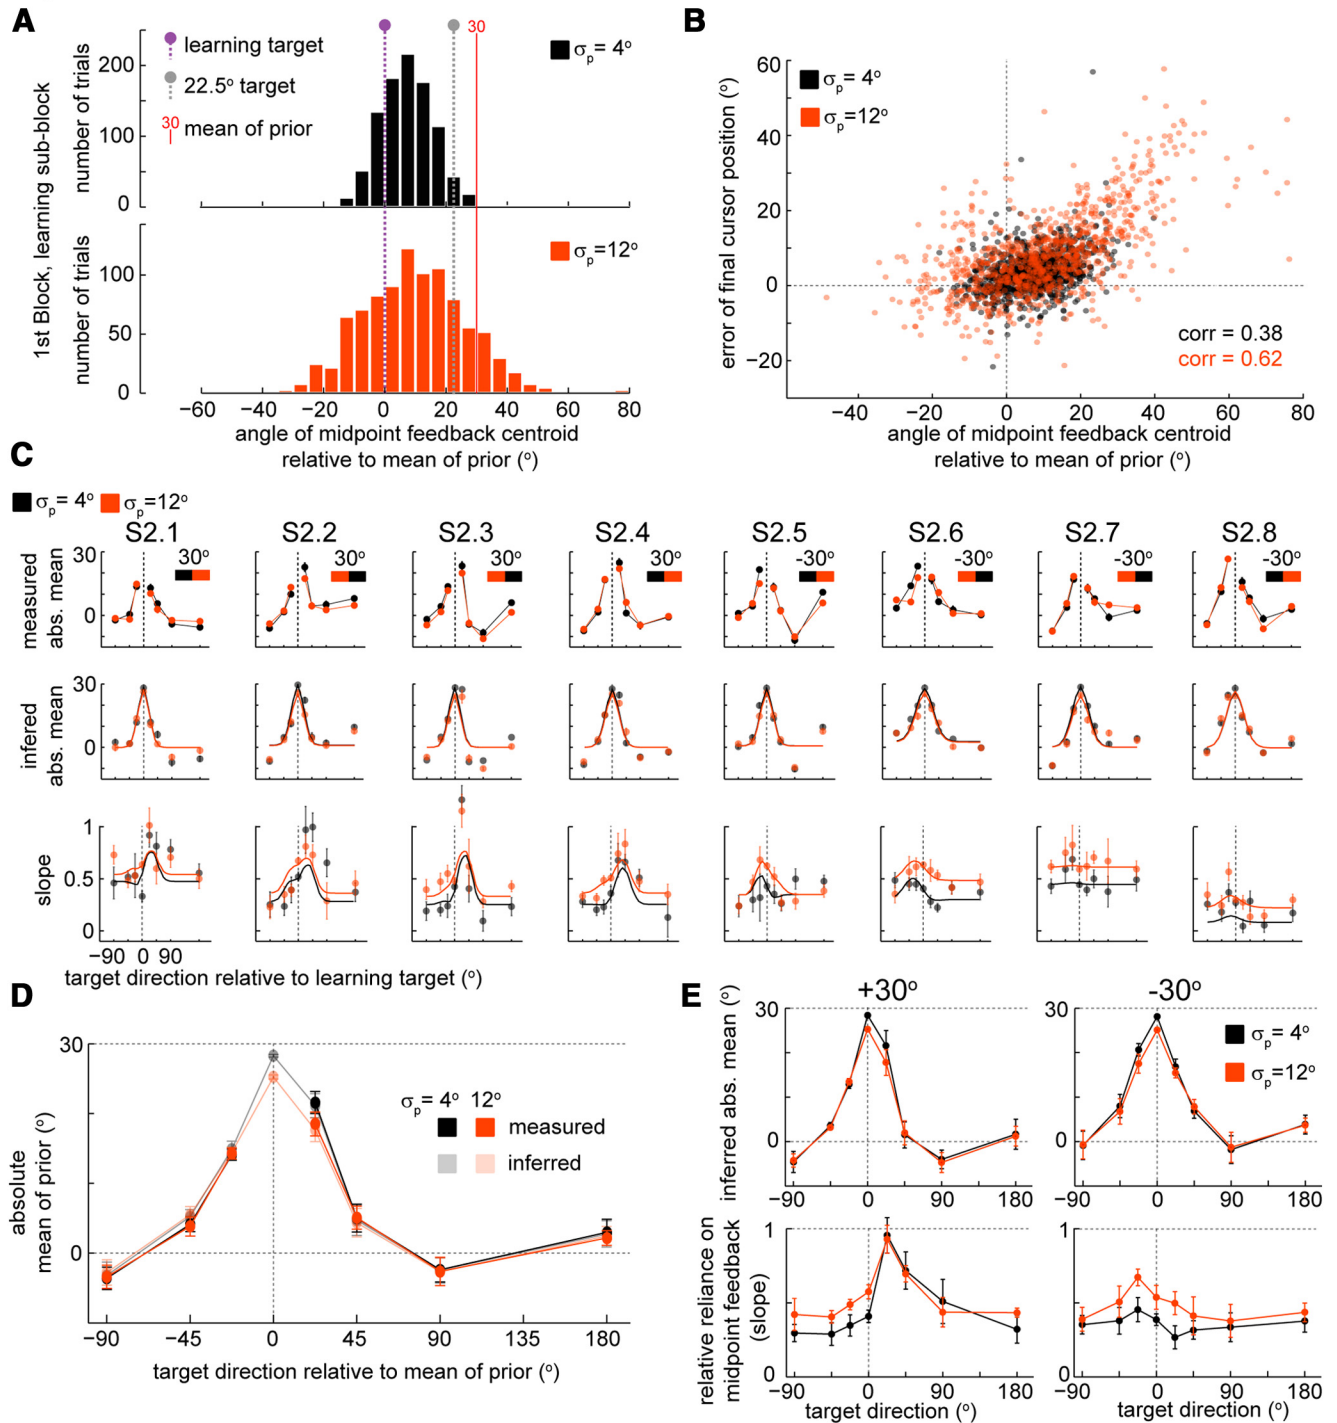
If the amount of generalization depends only on similarity between contexts and context is symmetric around target, then we would not expect to see an asymmetric pattern in the generalization of variance. In practice, however, the coordinate systems in which subjects try to solve the problem can have an influence on the generalization patterns (Gonzalez Castro et al., 2011). To allow for this possibility we hypothesized that the asymmetry could have arisen from a context that is not target centered. Do subjects learn about visual feedback position (Taylor et al., 2013) when generalizing uncertainty?

To see if the data are consistent with this hypothesis and to implement models where feedback position is relevant, we need to consider the distribution of learning data. One natural way of implementing such a model is in terms of on-line gradient descent. Every trial, one goal of the movement system may be to update certain parameters so that future movements will be better—we want to go down the gradient of errors (Thoroughman and Shadmehr, 2000; Taylor et al., 2013). We thus implemented two on-line learning models that take the perturbations imposed during the learning trials. These models implement gradient descent on the value of (assumed direction-dependent) mean and variance of the before minimize the squared error between target angle and final cursor position angle of each trial (see Materials and Methods for details). While for both models the reference frame for the generalization of the mean is target centered, Model VF (Visual Feedback) uses a coordinate system for generalization of variance related to the position of visual feedback while Model T (Target) uses only target-centered coordinates. This way of phrasing the problem allows us to consider the effect of the candidate coordinate systems on learning and generalization.

We find that Model VF captures the generalization patterns of both experiments (Figs. 3, 5C, second and third rows for individual subject fits, 6). Importantly, Model VF was able to capture the asymmetric generalization of uncertainty of Experiment 2 (Fig. 6A) and, simultaneously, explain the data in Experiment 1 (Fig. 6C), except for the use-dependent effect. We find that, while none

←  
 Uncertainty sub-block. Solid lines denote linear fits to the data. Colored circles indicate to which target direction (left) the data corresponds. **B**, Learning of uncertainty for the two groups of subjects: subjects that started with the low-variance condition and subjects that started with the high-variance condition. Solid lines are average slopes ( $\pm$  SEM) across subjects considering bins of 20 trials. **C, D**, Relative reliance on midpoint feedback for the two levels of prior uncertainty as a function of target direction relative to mean of perturbation ( $\pm 30^{\circ}$ ) following the first (**C**) and second (**D**) Learning sub-blocks. Mean ( $\pm$  SEM) of slopes across all subjects ( $n = 8$ ) in Experiment 2 (opaque solid lines). The transparent lines are the results of Experiment 1 (same data as in Fig. 2D separated by blocks). Insets show trajectories to the two neighboring directions of the learning target for Subject 2.1 (same data as **A**). **E**, Inferred percentage adaptation ( $\pm$  SEM) for the mean in the generalizing directions relative to the learning direction across all subjects. Colors of circles on the x-axis in **C–E** indicate the correspondent target directions in **A** and insets in **C** and **D**.

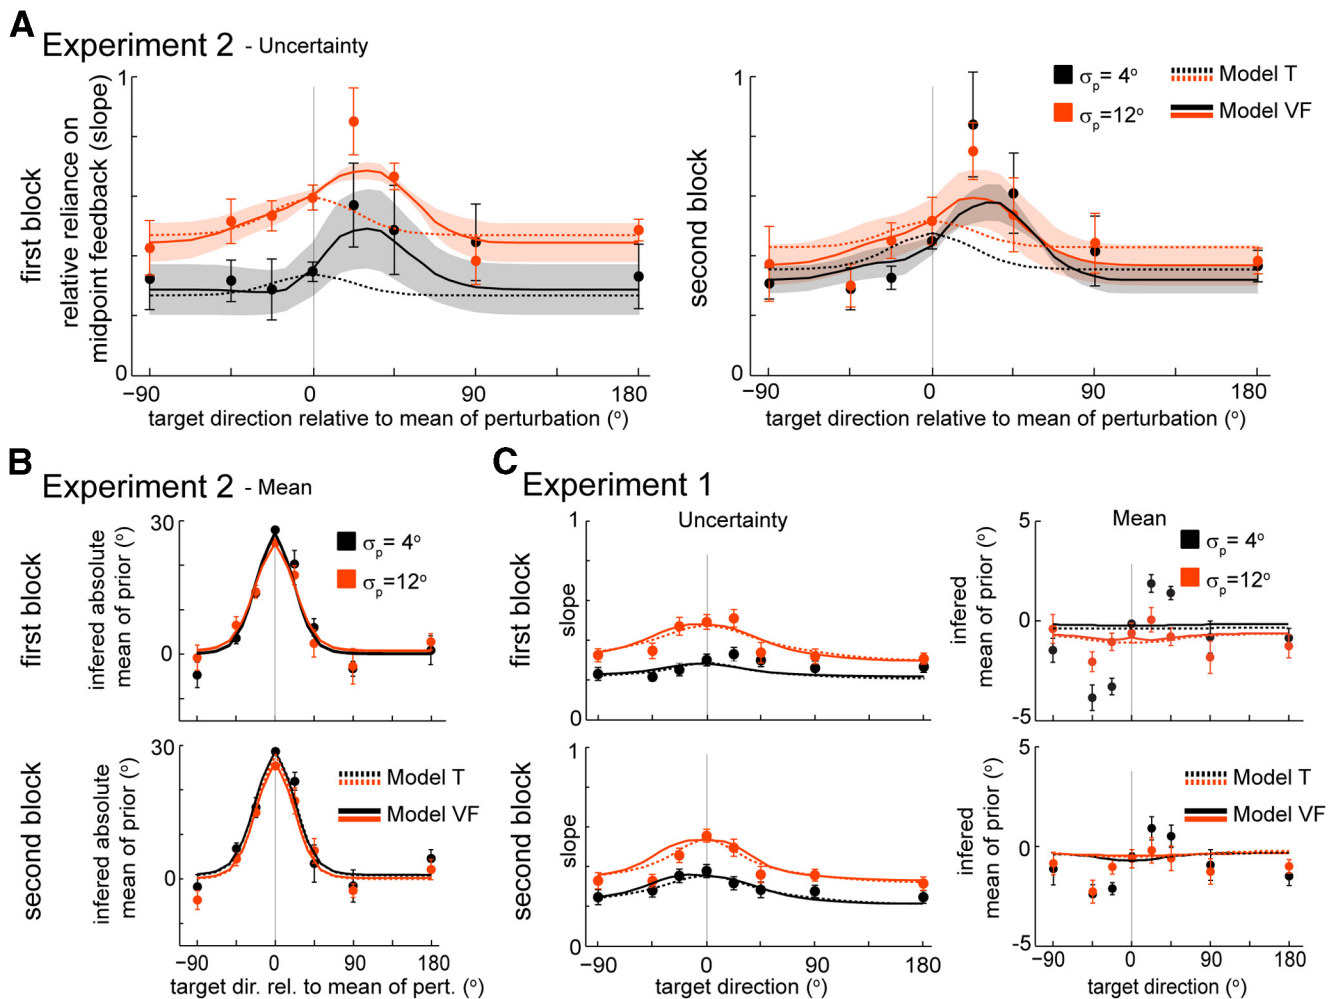
### Experiment 2



**Figure 5.** Individual subject data for Experiment 2 and further data analysis. **A**, Distribution of directions of the centroid of midpoint feedback (relative to the sign of the mean of the prior) for all subjects, for movements toward the learning direction during the first Learning sub-block. Target colors correspond to the circle colors on the x-axis of Figure 4C; this is the Learning sub-block that preceded the Testing sub-block used to make that figure. **B**, Trial error as a function of the angle of midpoint feedback centroid during the first Learning sub-block for all subjects in Experiment 2. **C**, Individual subject data for Experiment 2. Error bars  $\pm$  SEM (bootstrap). Black and orange bar (first row, inset) indicates the order in which the different uncertainty blocks were presented to each subject. Lines in second and third row are Model VF (see above, Models) fits to individual subjects. **D**, Generalization of mean ( $\pm$  SEM) measured during the Testing Mean sub-block (opaque), compared with the inferred generalization of the mean (transparent, same data as Fig. 4E) during Testing Uncertainty sub-block. **E**, Generalization of mean and of relative reliance on midpoint feedback (slope) separated by sign of the mean of the perturbation bars  $\pm$  SEM. The asymmetry in generalization of uncertainty (slope) was stronger for the positive mean (+30°, left bottom).

of the models is significantly better for Experiment 1 ( $p = 0.16$  for uncertainty and  $p = 0.61$  for mean; Wilcoxon signed rank test; Fig. 7A, C), Model VF is better than Model T for Experiment 2: notice that for all subjects the RMSE for uncertainty in Model VF

is smaller or equal than the RMSE obtained with Model T ( $p = 0.008$  for uncertainty and  $p = 0.055$  for the mean; Wilcoxon signed rank test; Fig. 7B, D). Although lacking a normative interpretation/justification, using different reference frames for mean



**Figure 6.** On-line learning models for different reference frames. **A**, Fits of the models (shaded area is  $\pm$  SEM for Model VF fits) for the slope data of Experiment 2 (same as Fig. 4C,D, opaque). Error bars  $\pm$  SEM. **B**, Models' fits for the mean data of Experiment 2. **C**, Fits of the models for mean and slope data (same as Fig. 4C,D, transparent) of Experiment 1. Error bars  $\pm$  SEM. Lines are average across subjects of individual fits. Error bars  $\pm$  SEM.

and variance and using gradient descent learning accurately captures the generalization patterns across experiments.

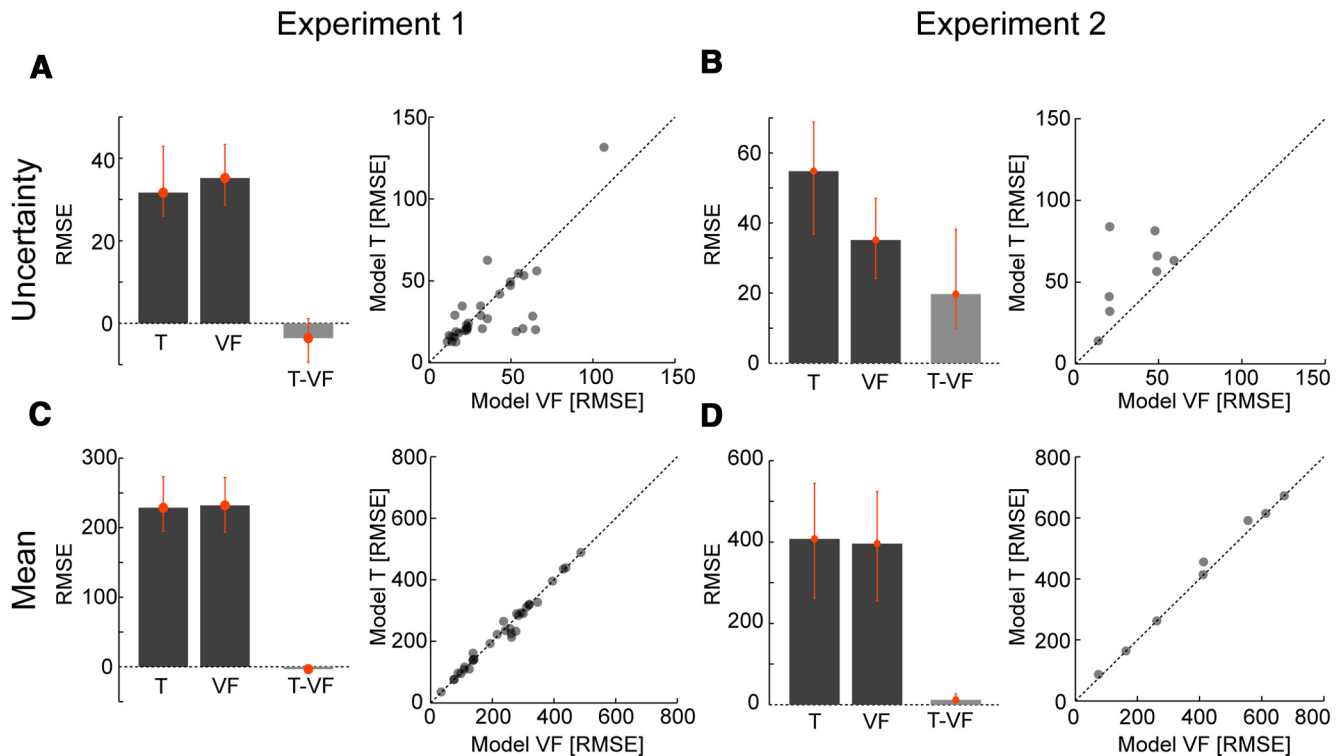
### Experiment 3

Previous studies have examined the effects of sudden and gradually introduced perturbations (Kagerer et al., 1997; Turnham et al., 2012) and have found different behaviors when perturbations are introduced gradually; for instance, gradually introduced visuomotor perturbations can lead to more complete adaptation and larger aftereffects (Kagerer et al., 1997). If Model VF is correct then we expect gradually introduced perturbations (i.e., when the mean of the distribution of perturbations is increased slowly) to produce a less asymmetric generalization curve for uncertainty compared with Experiment 2. In this case, during the learning trials, we expect errors to be small and the midpoint feedback to appear consistently closer to learning target.

To test this hypothesis we ran a third experiment (Experiment 3) with the same structure as Experiment 2 except for the fact that, during the first Learning sub-block, the mean of the perturbation is introduced slowly. Specifically, the mean of the distribution of perturbations changes linearly from 0 to 30° during the first half (120 trials) of the first Learning sub-block (see Materials and Methods for details). The progressively introduced mean is reflected in the learning curves (Fig. 8A) and the behavior during

learning seems to be less variable across subjects. Notice that, as in Experiment 2, during the Testing Uncertainty sub-block, hand trajectories in the learning direction (Fig. 8B, left, purple traces) overlap with hand trajectories in trials toward the neighboring target direction of lower uncertainty (Fig. 8B, left, green traces). Conversely, if there is incomplete learning of the mean in the learning target direction then the hidden cursor trajectories in the learning direction will be closer to the hidden cursor trajectories in the direction of maximum uncertainty (Fig. 8B, right, gray traces). The asymmetry in the generalization of uncertainty is still present (Fig. 8C,D), but it is smaller than in Experiment 2 ( $p = 0.009$ , two-way nested ANOVA over subject, variance, and experiment where subject is nested in experiment). As predicted, when introducing the mean of the perturbation progressively, we observe a smaller asymmetry in the generalization of uncertainty.

We could have expected the asymmetry to disappear completely; however, despite the mean of the perturbation being introduced progressively, the centroid of the midpoint feedback still appears closer to the direction of higher uncertainty (Fig. 9A,B). Even though, unlike Experiment 2, in Experiment 3 we do not find a significant difference between Models T and VF ( $p > 0.05$  for uncertainty and mean, Wilcoxon signed rank test), we still observe an asymmetric generalization pattern consistent with



**Figure 7.** Model comparison. **A–D**, Model comparison for Models T and VF for Uncertainty (**A, B**) and for Mean (**C, D**). Left, Weighted RMSE across subjects (95% confidence intervals, bootstrap) of each model and of the difference between models for each subject. Right, Scatterplot of the RMSE for each subject for Models T and VF.

a visual-centered reference frame, and Model VF is able to qualitatively capture the asymmetry in most of the subjects (Fig. 9C).

In a first analysis it seems that the asymmetry in the generalization of the mean and the asymmetry in the generalization of uncertainty have different causes. The former is hypothesized to be caused by use-dependence and the latter by a visual-feedback reference frame. We cannot exclude, however, that they might have a common explanation. In fact, the amount of asymmetry (as measured by the difference in generalized mean or slope between the two target directions that neighbor the learning direction) in the mean generalization and in the uncertainty generalization are correlated ( $r = 0.58$ ,  $p < 0.001$  for Experiments 2 and 3 combined,  $r = 0.58$ ,  $p = 0.033$  for Experiment 2, and  $r = 0.49$ ,  $p = 0.053$  for Experiment 3). Here we have found an asymmetry in the generalization pattern of uncertainty that is consistent with the use of a visual-feedback reference frame and excludes the possibility of a pure target-based reference frame for the generalization of uncertainty.

## Discussion

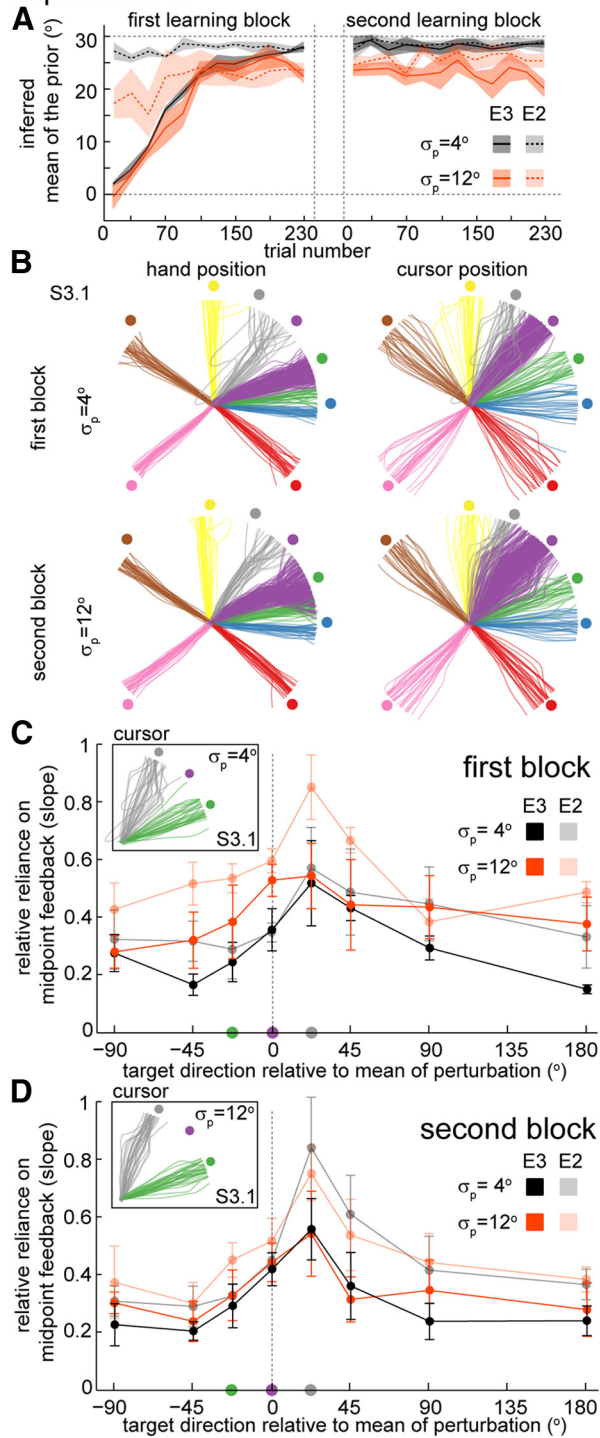
Here we examined how priors over a stochastic visuomotor perturbation generalize. We examined, in particular, how prior uncertainty generalizes, that is, how knowledge of the trial-by-trial variability of an event (in our experiment, the perturbation) generalizes to similar but distinct contexts (here, reaches to other directions). We first tested generalization in a paradigm in which only the variance of the distribution of rotation perturbations was changed, while the mean remained zero. We found that, similarly to standard generalization of constant visuomotor rotations, generalization of uncertainty has a local component. However, unlike the mean, it affects movements into all directions. We then tested how uncertainty generalizes when we introduce a stochastic perturbation with nonzero mean, i.e., when

both the mean and variance are perturbed. We observed asymmetric generalization that is qualitatively consistent with a descriptive, on-line learning model that assumes that mean and variance generalize according to different reference frames.

In movement research, generalization experiments are usually interpreted as being directly related to neuronal tuning properties (Krakauer et al., 2000; Thoroughman and Shadmehr, 2000; but see Pearson et al., 2010; Taylor and Ivry, 2013). Under this interpretation they constrain our conceptualization of neural computation and reveal a great deal about the neural basis of sensorimotor integration. We had seen evidence for some independence in representation of mean and variance of priors in previous studies when we showed that uncertainty does not affect the width of generalization of the mean (Fernandes et al., 2012). The results of this study indicate that knowledge of the variance of external perturbations might be represented in a way that is distinct from the knowledge about the mean—both the extent of generalization and reference frames appear to differ.

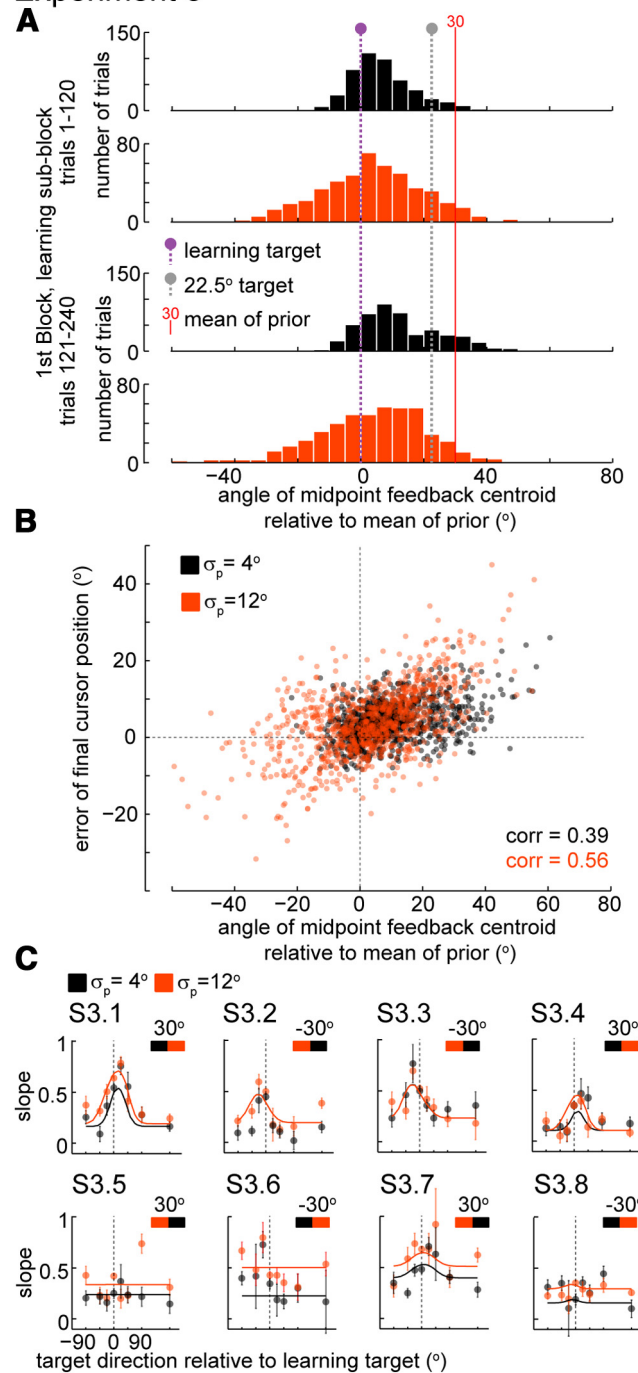
The degree to which the brain is “Bayesian” has been extensively debated over the last decade (Doya, 2007). Many studies have shown that the brain achieves Bayes-like behavior for familiar tasks (such as reaching) and that this behavior stems from ongoing learning (Berniker et al., 2010). Such general-purpose Bayesian behavior may result from a variety of non-Bayesian/heuristic neural representations. Alternatively, Bayesian ideas may be far more fundamental to the organization of the brain in the sense that there is something Bayesian about the neural code itself. For example, spikes in populations of neurons might directly represent probability distributions, including their means and variances (Hinton and Sejnowski, 1983; Zemel et al., 1998; Hoyer and Hyvärinen, 2003; Sahani and Dayan, 2003; Wu et al., 2003; Ma et al., 2006; Deneve, 2008; Fiser et al., 2010; Ma, 2010;

Experiment 3



**Figure 8.** Experiment 3. **A**, Learning of mean for Experiments 2 (dashed lines) and 3 (solid lines) for both groups of subjects: subjects that started with the low-variance condition ( $\sigma_p = 4^\circ$ , black) and subjects that started with the high-variance condition ( $\sigma_p = 12^\circ$ , orange). Lines are average slopes ( $\pm$  SEM) across subjects considering bins of 20 trials. **B**, Hand and cursor trajectories of a particular subject (Subject 3.1) during the Testing Uncertainty sub-blocks for the two levels of variance (SD of 4 and  $12^\circ$ ) in the learning direction (purple) and the generalizing directions (other colors). **C, D**, Relative reliance on midpoint feedback for the two levels of prior uncertainty as a function of target direction relative to the mean of the perturbation ( $\pm 30^\circ$ ) following the first (**C**) and second (**D**) Learning sub-blocks. Mean ( $\pm$  SEM) of slopes across all subjects ( $n = 8$ ) in Experiment 3 (opaque solid lines). The transparent lines are the results of Experiment 2 (same data as in Fig. 4C,D). Insets show trajectories to the two neighboring directions of the learning target for Subject 3.1 (same data as **B**). Colors of circles on the x-axis in **C** and **D** indicate the correspondent target directions in the insets and in **A**.

Experiment 3



**Figure 9.** Individual subject data for Experiment 3 and further data analysis. **A**, Distribution of directions of the centroid of midpoint feedback (relative to the sign of the mean of the prior) for all subjects in Experiment 3 ( $n = 8$ ) for movements toward the learning direction during the first Learning sub-block; first half (trials 1–120, when the mean of the distribution of perturbations is being progressively increased from 0 to  $30^\circ$ ) is separated from second half (trials 121–240) of that Learning sub-block. Target colors correspond to the circle colors on the x-axis of Figure 8C; this is the Learning sub-block that preceded the Testing sub-block used to make that figure. **B**, Trial error as a function of the angle of midpoint feedback centroid during the first Learning sub-block for all subjects in Experiment 3. **C**, Individual subjects slope data for Experiment 3. Error bars  $\pm$  SEM (bootstrap). Lines are Model VF fits to individual subjects. Black and orange bar (inset) indicates the order in which the different uncertainty blocks were presented to each subject (black denotes the  $4^\circ$  condition and orange denotes the  $12^\circ$  condition).

Soltani and Wang, 2010). However, none of these “Bayesian brain” theories explicitly predicts generalization of uncertainty. Furthermore, generalization is probably related to underlying neural representations in a more complex way than generally assumed in motor control research. Dissociation between generalization of mean and variance emerges immediately from our results and produces an important challenge to extensions of “Bayesian brain” theories to generalization.

The lack of computational predictions for generalization of priors, and of uncertainty in particular, is mirrored in experimental work where the focus both in behavioral as well as in electrophysiological studies in motor control has been on the generalization and representation of fixed perturbations without any trial-by-trial variability (but see Fernandes et al., 2012; Verstynen and Sabes, 2011). To our knowledge, previous work had not indicated that means and variances could generalize differently.

Previous studies have suggested that the reference frames for generalization might be based on actual rather than planned (target-based) trajectories (Gonzalez Castro et al., 2011). Other studies have shown that generalization depends on context, and that we can expect different generalization patterns if different contexts are imposed (Berniker and Kording, 2008; Brayanov et al., 2012; Taylor and Ivry, 2013). Studies that focus on the adaptation of the mean suggested that feedback plays an important role in adaptation (Huang et al., 2011) and generalization (Taylor et al., 2013). Differences in visual error information lead to changes in generalization that can be explained by a neural network that assumes error feedback processing on a set of homogeneous and invariant tuning functions (Taylor et al., 2013). The use of the nontarget-based reference frame for the learning of the mean has thus been shown previously. This study suggests the use of a visual reference frame for the generalization of uncertainty.

We have characterized the differences between the generalization of mean and uncertainty of a visuomotor rotation. We have shown, in particular, that unlike the mean, uncertainty generalizes globally. The mean, however, has been shown to generalize globally in visuomotor gain experiments (Krakauer et al., 2000). A useful follow-up experiment would be to extend the visuomotor gain paradigm to also include uncertainty, and characterize the asymmetries, differences in width and reference frames of the generalization patterns under that paradigm.

We find clear signs that movements are biased toward typical directions of previous hand movements, which is consistent with the use-dependent learning/adaptation hypothesis (Diedrichsen et al., 2010; Huang et al., 2011; Verstynen and Sabes, 2011). We find this in both experiments and it is particularly evident in Experiment 1 where the mean of the distribution of stochastic perturbations was zero; this use-dependent asymmetry scales with the uncertainty level and exists even when there is zero mean perturbation. Although our model captures features of generalization patterns for both mean and uncertainty, it does not capture this use-dependent aspect of the generalization. Future modeling work is needed to account for these effects by incorporating a hand-centered reference frame or other “model-free” learning processes (Huang et al., 2011). In fact, even though Model VF is consistent with the observed symmetry we cannot exclude that it might be caused by other mechanisms. It is not an unreasonable hypothesis that the same mechanism responsible for the use-dependent asymmetry in the generalization of the mean is responsible for the asymmetry in the generalization of uncertainty. If this is true it happens in a way that is not obvious to us and future research could try to address it.

Where priors come from and how they are represented are fundamental questions in learning and behavior. As we never experience the same situation twice, constructing priors depends crucially on our ability to generalize. However, generalization in both perception and action is a result of how the brain represents the external world. In perception research, studies that hypothesize priors based on the statistics of natural scenes (Geisler et al., 2001; Roth and Black, 2005; Burge et al., 2010; DiMattina et al., 2012) generally assume certain invariances where global generalization occurs along many dimensions of the stimulus. When calculating orientation priors, for instance, color and contrast are assumed to be irrelevant and only the statistics over orientation are considered important (Girshick et al., 2011). In movement research, it is generally assumed that the system is invariant to the content of the visual scene and that generalization only depends on (angular) distance (Krakauer et al., 2000), velocity (Goodbody and Wolpert, 1998), or the way an object is held (Ingram et al., 2010; but see Taylor et al., 2013). For both perception and action, the nature of the underlying representations determines the shape of generalization, and experiments like the ones presented here are a way at psychophysically addressing these.

## References

- Berniker M, Kording K (2008) Estimating the sources of motor errors for adaptation and generalization. *Nat Neurosci* 11:1454–1461. [CrossRef Medline](#)
- Berniker M, Voss M, Kording K (2010) Learning priors for Bayesian computations in the nervous system. *PLoS One* 5:e12686. [CrossRef Medline](#)
- Brayanov JB, Press DZ, Smith MA (2012) Motor memory is encoded as a gain-field combination of intrinsic and extrinsic action representations. *J Neurosci* 32:14951–14965. [CrossRef Medline](#)
- Burge J, Fowlkes CC, Banks MS (2010) Natural-scene statistics predict how the figure-ground cue of convexity affects human depth perception. *J Neurosci* 30:7269–7280. [CrossRef Medline](#)
- Cisek P, Kalaska JF (2005) Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45:801–814. [CrossRef Medline](#)
- Deneve S (2008) Bayesian spiking neurons I: inference. *Neural Comput* 20:91–117. [CrossRef Medline](#)
- Diedrichsen J, White O, Newman D, Lally N (2010) Use-dependent and error-based learning of motor behaviors. *J Neurosci* 30:5159–5166. [CrossRef Medline](#)
- DiMattina C, Fox SA, Lewicki MS (2012) Detecting natural occlusion boundaries using local cues. *J Vis* 12(13):15. [CrossRef Medline](#)
- Doya K, Ishii S, Pouget A, Rao RPN (2007) Bayesian brain: probabilistic approaches to neural coding. Cambridge, MA: MIT.
- Fernandes HL, Stevenson IH, Kording KP (2012) Generalization of stochastic visuomotor rotations. *PLoS One* 7:e43016. [CrossRef Medline](#)
- Fiser J, Berkes P, Orbán G, Lengyel M (2010) Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn Sci* 14:119–130. [CrossRef Medline](#)
- Geisler WS, Perry JS, Super BJ, Gallogly DP (2001) Edge co-occurrence in natural images predicts contour grouping performance. *Vision Res* 41:711–724. [CrossRef Medline](#)
- Georgopoulos AP, Ashe J, Smyrnis N, Taira M (1992) The motor cortex and the coding of force. *Science* 256:1692–1695. [CrossRef Medline](#)
- Girshick AR, Landy MS, Simoncelli EP (2011) Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nat Neurosci* 14:926–932. [CrossRef Medline](#)
- Gonzalez Castro LNG, Monsen CB, Smith MA (2011) The binding of learning to action in motor adaptation. *PLoS Comput Biol* 7:e1002052. [CrossRef Medline](#)
- Goodbody SJ, Wolpert DM (1998) Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79:1825–1838. [Medline](#)
- Hinton GE, Sejnowski TJ (1983) Optimal perceptual inference. In: *Proceedings of the IEEE computer society conference on computer vision and pattern recognition*, pp 448–453.
- Hoyer P, Hyvärinen A (2003) Interpreting neural response variability as Monte Carlo sampling of the posterior. In: *Advances in neural informa-*

- tion processing systems (Becker S, Thrun S, Obermayer K, eds), pp 293–300. Cambridge, MA: MIT.
- Huang VS, Haith A, Mazzoni P, Krakauer JW (2011) Rethinking motor learning and savings in adaptation paradigms: model-free memory for successful actions combines with internal models. *Neuron* 70:787–801. [CrossRef Medline](#)
- Ingram JN, Howard IS, Flanagan JR, Wolpert DM (2010) Multiple grasp-specific representations of tool dynamics mediate skillful manipulation. *Curr Biol* 20:618–623. [CrossRef Medline](#)
- Kagerer FA, Contreras-Vidal JL, Stelmach GE (1997) Adaptation to gradual as compared with sudden visuo-motor distortions. *Exp Brain Res* 115:557–561. [CrossRef Medline](#)
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247. [CrossRef Medline](#)
- Körding KP, Beierholm U, Ma WJ, Quartz S, Tenenbaum JB, Shams L (2007) Causal inference in multisensory perception. *PLoS One* 2:e943. [CrossRef Medline](#)
- Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. *J Neurosci* 20:8916–8924. [Medline](#)
- Ma WJ (2010) Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Res* 50:2308–2319. [CrossRef Medline](#)
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438. [CrossRef Medline](#)
- Moran DW, Schwartz AB (1999) Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82:2676–2692. [Medline](#)
- Pearson TS, Krakauer JW, Mazzoni P (2010) Learning not to generalize: modular adaptation of visuomotor gain. *J Neurophysiol* 103:2938–2952. [CrossRef Medline](#)
- Rickert J, Riehle A, Aertsen A, Rotter S, Nawrot MP (2009) Dynamic encoding of movement direction in motor cortical neurons. *J Neurosci* 29:13870–13882. [CrossRef Medline](#)
- Roth S, Black MJ (2005) On the spatial statistics of optical flow. In: *IEEE International Conference on Comput Vision*, Vol. 1, pp 42–49.
- Sahani M, Dayan P (2003) Doubly distributional population codes: simultaneous representation of uncertainty and multiplicity. *Neural Comput* 15:2255–2279. [CrossRef Medline](#)
- Shadmehr R (2004) Generalization as a behavioral window to the neural mechanisms of learning internal models. *Hum Mov Sci* 23:543–568. [CrossRef Medline](#)
- Shepard RN (1987) Toward a universal law of generalization for psychological science. *Science* 237:1317–1323. [CrossRef Medline](#)
- Soltani A, Wang XJ (2010) Synaptic computation underlying probabilistic inference. *Nat Neurosci* 13:112–119. [Medline](#)
- Tassinari H, Hudson TE, Landy MS (2006) Combining priors and noisy visual cues in a rapid pointing task. *J Neurosci* 26:10154–10163. [CrossRef Medline](#)
- Taylor JA, Ivry RB (2013) Context-dependent generalization. *Front Hum Neurosci* 7:171. [CrossRef Medline](#)
- Taylor JA, Hieber LL, Ivry RB (2013) Feedback-dependent generalization. *J Neurophysiol* 109:202–215. [Medline](#)
- Thoroughman KA, Shadmehr R (2000) Learning of action through adaptive combination of motor primitives. *Nature* 407:742–747. [CrossRef Medline](#)
- Turnham EJ, Braun DA, Wolpert DM (2012) Facilitation of learning induced by both random and gradual visuomotor task variation. *J Neurophysiol* 107:1111–1122. [CrossRef Medline](#)
- Verstynen T, Sabes PN (2011) How each movement changes the next: an experimental and theoretical study of fast adaptive priors in reaching. *J Neurosci* 31:10050–10059. [CrossRef Medline](#)
- Vilares I, Kording K (2011) Bayesian models: the structure of the world, uncertainty, behavior, and the brain. *Ann NY Acad Sci* 1224:22–39. [CrossRef Medline](#)
- Wei K, Körding K (2009) Relevance of error: what drives motor adaptation? *J Neurophysiol* 101:655–664. [Medline](#)
- Wu S, Chen D, Niranjana M, Amari S (2003) Sequential Bayesian decoding with a population of neurons. *Neural Comput* 15:993–1012. [CrossRef Medline](#)
- Zemel RS, Dayan P, Pouget A (1998) Probabilistic interpretation of population codes. *Neural Comput* 10:403–430. [CrossRef Medline](#)