Uncertainty of feedback and state estimation determines the speed of motor adaptation

Kunlin Wei[1,2,3]* and Konrad Körding[2,3]

[1] Department of Psychology, Peking University, Beijing, China
[2] Departments of Physiology, Physical Medicine and Rehabilitation, and Applied Mathematics, Northwestern University, Chicago, IL, USA
[3] Bayesian Behavior Lab, Sensory Motor Performance Program, Rehabilitation Institute of Chicago, IL, USA

Humans can adapt their motor behaviors to deal with ongoing changes. To achieve this, the nervous system needs to estimate central variables for our movement based on past knowledge and new feedback, both of which are uncertain. In the Bayesian framework, rates of adaptation characterize how noisy feedback is in comparison to the uncertainty of the state estimate. The predictions of Bayesian models are intuitive: the nervous system should adapt slower when sensory feedback is more noisy and faster when its state estimate is more uncertain. Here we want to quantitatively understand how uncertainty in these two factors affects motor adaptation. In a hand reaching experiment we measured trial-by-trial adaptation to a randomly changing visual perturbation to characterize the way the nervous system handles uncertainty in state estimation and feedback. We found both qualitative predictions of Bayesian models confirmed. Our study provides evidence that the nervous system represents and uses uncertainty in state estimate and feedback during motor adaptation.

Keywords: motor learning, motor adaptation, uncertainty, Bayesian statistics

INTRODUCTION

To successfully perform sensorimotor tasks the nervous system needs to estimate the state of both the body and the environment; however, almost all real life estimation problems are plagued with uncertainty (von Helmholtz, 1863/1954). Two types of uncertainty will influence how we form our estimates about the state. First, each of our sensory modalities is noisy and thus offers only a certain level of precision which can vary depending on the situation. For example, vision is limited under dim-light or extra-foveal conditions, hearing becomes unreliable for weak sounds, and proprioception drifts without calibration from vision (Brown et al., 2003). This leads to *feedback uncertainty*. Second, the brain predicts changes of the state using knowledge of motor commands (efference copy) and the physics of the environment (McIntyre et al., 2001) possibly using an internal model of the task (Wolpert et al., 1995). This so-called forward prediction can be combined with sensory feedback to improve the accuracy of state estimation. However, this prediction is affected by the state of the body, which evolves in partially unpredictable ways and on many different time scales: neuromuscular noise contaminates muscle activity, muscles fatigue and recover frequently, and the body undergoes long-term changes such as diseases and development (Körding et al., 2007). These factors lead to *state estimation uncertainty* that directly affects the predictions the nervous system may make. Understanding the interaction between feedback uncertainty and state estimation uncertainty during sensorimotor tasks is one of the central problems in neural control of movement.

Bayesian statistics prescribes a systematic way of integrating multiple pieces of uncertain information and forms the basis of optimal estimation and control. Visuomotor behaviors exhibit a number of features predicted by optimal estimation and Bayesian statistics (van Beers et al., 1999; Korenberg and Ghahramani, 2002; Alais and Burr, 2004; Ernst and Bülthoff, 2004; Haith et al., 2008; Knill and Pouget, 2004; van Beers, 2009). In particular, Bayesian models make specific predictions about the role of feedback and state estimation uncertainty in estimating movement-related variables (Korenberg and Ghahramani, 2002; Burge et al., 2008; Izawa and Shadmehr, 2008). For example, when playing tennis we might observe that we return the ball short of an intended location. We will then adapt our next swing to this error. To optimally perform the task, we will adapt more if we clearly see the error as compared to if we only see it through peripheral vision (high feedback uncertainty) and that we will adapt more if we are still warming up for today's practice than if we have been playing for a while (low state estimation uncertainty). The Bayesian predictions for trial-by-trial adaptation are rather intuitive: across-trial adjustment is more pronounced when state estimates are more uncertain *and/or* when estimates in feedback are more certain to the nervous system.

The present study tests these predictions using a trial-by-trial adaptation task with visual perturbations. We compare the Bayesian model with state space models (e.g., Cheng and Sabes, 2006), both of which are frequently used for studying motor learning. We also compare to models that combine features of both state space models and Kalman models. In Experiment 1 we manipulate feedback uncertainty by blurring the visual feedback. In Experiment 2 we provide feedback of different quality for fixed periods of time to leave the subject with defined state estimation uncertainty. As qualitatively predicted by Bayesian models, adaptation is significantly faster with less feedback uncertainty and with more state estimation uncertainty. These results suggest that the nervous system represents feedback uncertainty and state estimation uncertainty and uses knowledge of uncertainty during motor adaptation.

MATERIALS AND METHODS

APPARATUS

The subjects were seated and held a handle at the endpoint of a 2-D robotic linkage with their right hand (**Figure 1A**, see elsewhere for detailed description of the manipulandum, Shadmehr and Mussa-Ivaldi, 1994). The hand could be moved freely and frictionless but restricted in the horizontal (transverse) plane. The seat height was adjusted for each subject to keep the right arm at the shoulder level. The right upper arm was also supported by a customized harness hung from the ceiling. The position of the hand was measured by the manipulandum at a frequency of 250 Hz. Visual feedback (detailed below) was projected in real time through an overhead projector (model NEC LT170) onto a white horizontal board with a refreshing rate of 75 Hz. Vision of the forearm and hand was obscured by the projection board.

EXPERIMENT 1

Visual representation of the hand position, the starting position, and the target were presented on a projection screen with a black background (**Figure 1B**). The starting position was displayed as a 0.5×0.5 cm white cross at the middle of the workspace aligned with the midline of the seated subject. The position of the hand was represented by a white cursor with a diameter of 0.3 cm. Each trial started when the subject placed the cursor at the starting position. After the hand remained stationary at that position for 100 ms, a target, represented as a yellow disc of 0.5-cm diameter, was displayed 15 cm straight to the right of the starting position. At the same time, the cross representing the starting position disappeared and a computer-generated beep was played to trigger the subject to make a horizontal movement to the target. We defined the movement direction from the starting position to the target position as the y axis of the workspace and the perpendicular direction as x axis.

As soon as the subject started moving their hand towards the target, the display of the cursor was extinguished. Visual feedback of the hand position was given again when the hand crossed the target in the y direction (**Figure 1B**). It remained visible for 100 ms at the same position. This feedback was perturbed in depth (the x direction in the figure), randomly and i.i.d. by one out of three possible values: 0, ±2 cm. This type of brief visual feedback at the end of the movement is similar to the terminal feedback utilized in other visuomotor adaptation studies (Cheng and Sabes, 2007; Wei and Kording, 2009). If this visual perturbation is effective, the subject will adapt and tend to move in the opposite direction of the visual disturbance in the subsequent trial. The compensatory movement would be evaluated as a x-direction deviation from the target when the hand crossed the target in the y direction (**Figure 1B**). Throughout the paper we use the word deviation to denote deviations of the actual hand position from the target position along this direction.

In addition to the magnitude of the perturbation, the visual feedback for each trial was also randomly assigned one of three different levels of blurring to manipulate its uncertainty (**Figure 1C**). The first level had no blurring and the terminal feedback was a single cursor that was designed to have the least feedback uncertainty (*NoBlur*). The second level presented the terminal feedback as a cloud of five cursors whose locations were randomly drawn from

an isotropic two-dimensional Gaussian distribution with a mean on the perturbed (or unperturbed) location and a standard deviation of 2 cm (*SmallBlur*). The third level was similar to the second, but the standard deviation of the distribution was 4 cm (*LargeBlur*). By showing the hand position as a cloud of cursors, we introduced additional feedback uncertainty relative to one cursor case. Larger standard deviations result in higher feedback uncertainty.
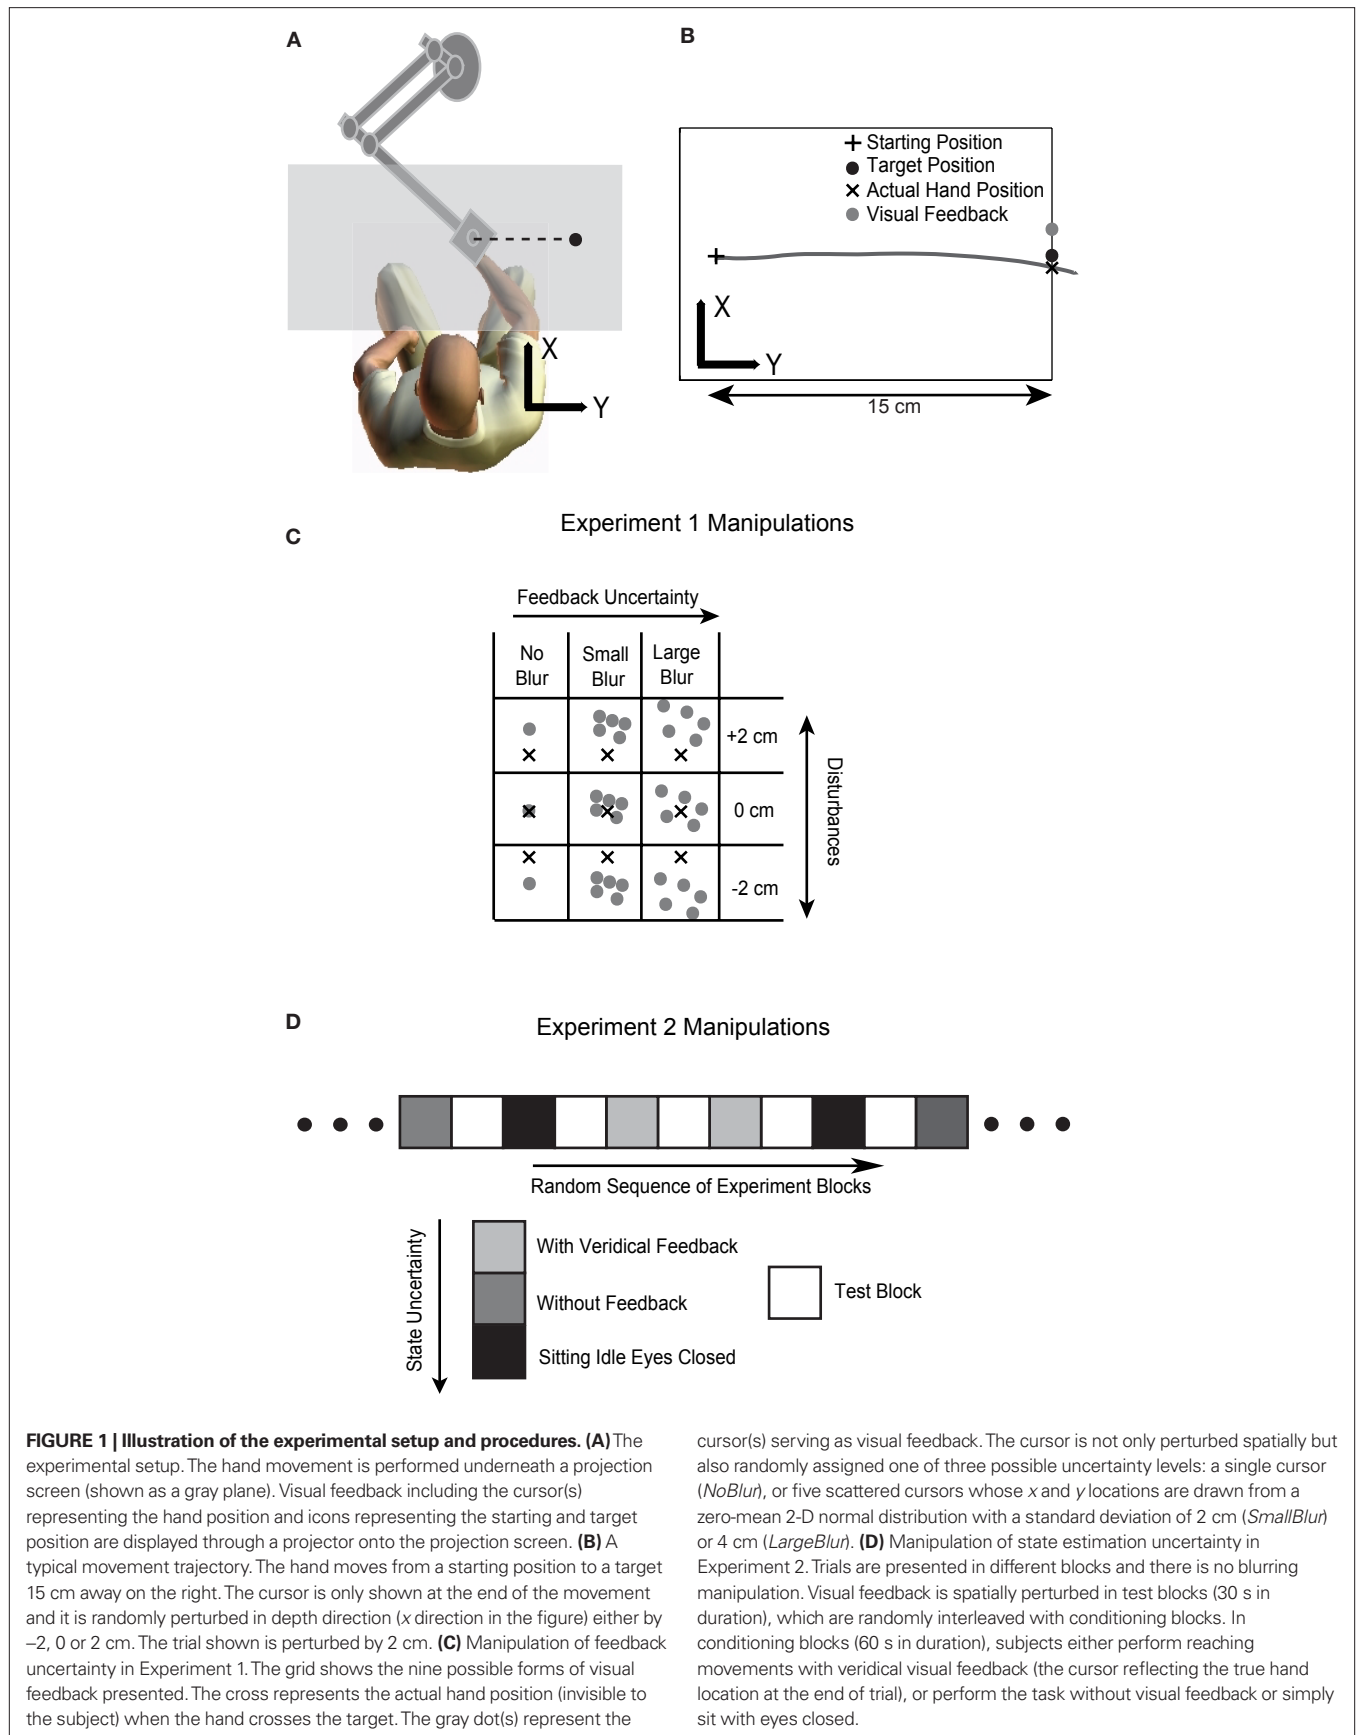
This approach, blurring visual feedback by displaying a cloud of cursors, has been used previously in manipulating uncertainty of visual target in studies of sensorimotor control (Trommershäuser et al., 2003; Körding and Wolpert, 2004; Tassinari et al., 2006). This blurring is also similar to so-called "Gaussian blob" that have been extensively used in psychophysics studies (Solomon et al., 1997; Schofield and Georgeson, 1999; Solomon, 2002) where the visual target was shown as a blob whose luminance was drawn from Gaussian distributions. In the present study we combined the features of these visual stimuli to generate a blurred and randomized visual perturbation.

After the visual feedback was extinguished, the hand was pulled back by the robot linkage to the starting position for the next trial without showing the cursor. The cursor appeared again once it was placed in the vicinity of the starting position (<0.5 cm). Restoring the initial hand position without vision minimizes visual calibration that might reduce adaptation in the next trial. Nevertheless, visual calibration around the starting position still exists. However, since this calibration is present in all conditions, its effect can be readily assumed to not interact with the main effects we are after.

Subjects were instructed to "hit" the target as accurately as possible with a single reach without corrective movement near the target. They were also informed that the cursor(s) were indicative of their hand position. Before starting the task, subjects were familiarized with the setup by practicing without visual disturbances for 3 min. Formal data collection consisted of 3 (perturbation amplitude) × 3 (blurring condition) × 50 (repetition) = 450 trials. The experiment lasted about 40 min with a 1-min mandatory break after the 225th trial. One subject had only finished 360 trials due to technical difficulties during data collection.

EXPERIMENT 2

Subjects performed the same reaching task as in the Experiment 1. However, the terminal feedback of the hand was only manipulated with spatial perturbations but without blurring, i.e., it was only shown as a single cursor. Adaptation to trial-by-trial perturbations was assessed in blocks of 30-s period (test block). Before each test block, one of three possible conditioning blocks was randomly assigned when subjects either performed the task with unperturbed terminal feedback, or without terminal feedback, or simply sat in the dark without movement (**Figure 1D**). These conditioning blocks lasted for 1 min each. They were designed to condition state estimation uncertainty to different levels by providing different quality of sensory feedback for certain duration. In the first condition subjects performed the task with the terminal feedback that truly reflected the actual hand position. In this condition, both visual and proprioceptive feedback were available for estimating the hand position and thus the nervous system had low uncertainty in estimating its state. In the second condition subjects had no visual error feedback to inform their performance and they solely relied

FIGURE 1 | Illustration of the experimental setup and procedures. (A) The experimental setup. The hand movement is performed underneath a projection screen (shown as a gray plane). Visual feedback including the cursor(s) representing the hand position and icons representing the starting and target position are displayed through a projector onto the projection screen. **(B)** A typical movement trajectory. The hand moves from a starting position to a target 15 cm away on the right. The cursor is only shown at the end of the movement and it is randomly perturbed in depth direction (x direction in the figure) either by −2, 0 or 2 cm. The trial shown is perturbed by 2 cm. **(C)** Manipulation of feedback uncertainty in Experiment 1. The grid shows the nine possible forms of visual feedback presented. The cross represents the actual hand position (invisible to the subject) when the hand crosses the target. The gray dot(s) represent the cursor(s) serving as visual feedback. The cursor is not only perturbed spatially but also randomly assigned one of three possible uncertainty levels: a single cursor (*NoBlur*), or five scattered cursors whose x and y locations are drawn from a zero-mean 2-D normal distribution with a standard deviation of 2 cm (*SmallBlur*) or 4 cm (*LargeBlur*). **(D)** Manipulation of state estimation uncertainty in Experiment 2. Trials are presented in different blocks and there is no blurring manipulation. Visual feedback is spatially perturbed in test blocks (30 s in duration), which are randomly interleaved with conditioning blocks. In conditioning blocks (60 s in duration), subjects either perform reaching movements with veridical visual feedback (the cursor reflecting the true hand location at the end of trial), or perform the task without visual feedback or simply sit with eyes closed.

on proprioception to estimate the hand position. This would potentially result in higher uncertainty in state estimates than in the first condition. In the third condition subjects sat idle without implementing motor plan and we expected this would lead to the highest uncertainty about the state of the motor system compared to the other two cases. Each type of conditioning blocks was presented eight times in a pseudo-random order. Data analysis was focused on the adaptation trials in test blocks following each conditioning block. The experiment had a total of 24 conditioning blocks and 24 corresponding test blocks, which led to about 36 min of data collection. There was no resting break as eight conditioning blocks without movements prevented subjects from fatiguing. Depending on subjects' movement speed, the number of trials finished within the given experiment time varied with an average count of 678.

PARTICIPANTS

Both experiments had seven volunteer subjects. Three subjects in Experiment 1 subsequently participated in Experiment 2. All subjects provided informed consent before experiments and were naïve to the purpose of the experiments. Experimental procedures were approved by the institutional review board of the Northwestern University. As our experiment required the subjects to move the hand from the left to the right, we only recruited right-handed subjects to eliminate the effect of handedness. All subjects had normal or corrected-to-normal vision.

KALMAN FILTER MODEL

Given the uncertainty in feedback and in state estimation and the assumption of a random walk of the relevant variables, the optimal solution for estimating the state can be obtained by a Kalman filter. The Kalman filter can optimally combine the feedforward estimates and the newly received feedback, as we illustrate in detail below.

Subjects make a reach based on their estimated hand position at the end of last movement and observe the movement outcome shown as a displayed cursor. This observation leads to an update of the estimate of hand position and thus influences the next reaching movement. This trial-by-trial adjustment is the basis of the Kalman filter:

$$\hat{x}_{k|k-1} = F\hat{x}_{k-1|k-1} \tag{1}$$

where $\hat{x}_{k-1|k-1}$ is the estimated hand position after the sensory feedback of the movement $k-1$ is integrated with the feedforward estimate. $\hat{x}_{k|k-1}$ is the feedforward estimate for movement k before the new feedback is taken into account. We assume the actual location of the hand at the end of movement k is a direct readout of this feedforward estimate. The feedforward calculation is governed by its forward control model F, which is also called the state transition model that captures the state change from one trial to the next. The uncertainty in the state estimate is obtained by the forward model and characterized by the standard deviation about the state σ_k which is updated after every new observation. Note the state estimation uncertainty varies from trial to trial and can only be inferred from the data.

The observed hand position z_k is noisy and thus may not be identical to the actual hand position:

$$z_k = H\hat{x}_{k|k-1} + \upsilon, \ \upsilon \sim N(0, R^2) \tag{2}$$

where H is the observation model, $\hat{x}_{k|k-1}$ is the actual location of the hand and υ is the observation noise that is assumed to be drawn from a Gaussian distribution with zero mean and a standard deviation of R, which captures feedback uncertainty. The Kalman filter will compute the difference between the observed state z_k and the predicted state $\hat{x}_{k|k-1}$ and use this error to update the estimation as well as state estimation uncertainty σ_k. How much the estimate is updated is specified by the Kalman gain, which is a function of the ratio between state estimation uncertainty and feedback uncertainty. All the code for analysis will be made available online.

Our specific model assumes that the system is linear with one latent state and Gaussian noise and its goal is to minimize the squared error of its estimates. The assumptions of linearity and Gaussian noise sources have been proven to sufficiently capture dynamics in simple sensorimotor tasks, especially in repetitive reaching tasks (e.g., Baddeley et al., 2003). The assumption of minimizing squared error is also reasonable since in the current study subjects are required to reach to the target as accurately as possible without directional bias.

We investigated the adaptation of hand placement in depth (the direction orthogonal to the movement direction), since the visual perturbation was only applied in that direction. As a result our model can be simplified greatly and all the variables in the model are scalars. The random trial-by-trial visual perturbations serve as the observations for the Kalman filter. In Experiment 1, the feedback uncertainty was manipulated by blurring cursor feedback. We thus leave the variances of the observation model R^2 as free parameters, one for each blurring condition. The other two free parameters are the transition model F and the variance of the process noise Q (a standard parameter in the Kalman filter). We also assume that the visual feedback is accurately perceived by subjects thus the observation model H is set to 1.

In essence, the Kalman filter model optimally estimates the hand position based on its feedforward estimates (outgoing motor commands, predicted state) and sensory feedback (incoming visual feedback). Thus, the model is about the sensory component of the task. We assume that subjects use the feedforward estimate (predicted state) to guide the hand and this estimate exhibits itself as the actual hand position. In other word, we assume the motor execution is a direct reflection of the motor planning. This is a reasonable assumption given that our study focus is on the role of uncertainty during the estimation process and that the task is a simple reaching task with a straight trajectory (see a similar treatment, Wei and Kording, 2009). We did not apply the Kalman filter model to Experiment 2 since test blocks were too short for that kind of analysis; as such Experiment 2 is a qualitative study for Bayesian predictions on state estimation uncertainty.

STATE SPACE MODEL

As an alternative to the Kalman filter model, the state space model has been extensively used in studying trial-by-trial adaptation in the literature (Baddeley et al., 2003; Donchin et al., 2003; Cheng and Sabes, 2006; Fine and Thoroughman, 2007). We fit our data with a state space model as follows:

$$\hat{x}_{k|k-1} = A\hat{x}_{k-1|k-1} + B\left(y_k - A\,\hat{x}_{k-1|k-1}\right) \tag{3}$$

where $\hat{x}_{k|k-1}$ is the estimated hand position after perceiving the visual feedback y_k at the movement k. It is linearly dependent on the previous estimated hand position $\hat{x}_{k-1|k-1}$ as well as the perceived error feedback $y_k - A\hat{x}_{k-1|k-1}$. The parameter A captures how the previous estimate influences the current estimate. The parameter B is the learning rate, which captures how much correction the system makes upon perceiving error information. To make the state space model comparable to the Kalman filter model, we make B vary as a function of blurring conditions as such it is a vector of size 3. Compared to the Kalman filter model, A is qualitatively similar to the transition model F (see Eq. 1) and B to the Kalman gain (see a general description of the Kalman model and our codes published online). Despite its similarities to the Kalman filter model, the state space model has no means to incorporate changing uncertainty in state estimate. The model fitting is done in the same way as for the Kalman filter model.

LINEAR INTERPOLATION BETWEEN KALMAN FILTERING AND THE STATE SPACE MODEL

Individual subjects may use a strategy that is in between Kalman filtering (estimation of state estimation uncertainty on trial-by-trial basis) and state space models (fixed estimation of state estimation uncertainty for the duration of the experiment). We thus construct a new model that is a hybrid of both models above, by simply replacing point estimates with a linear combination of the predictions of the Kalman model and the state space model ($\hat{x}_{k,\text{Kalman}}$ and $\hat{x}_{k,\text{SSM}}$):

$$\hat{x}_k = \alpha\,\hat{x}_{k,\text{Kalman}} + (1-\alpha)\,\hat{x}_{k,\text{SSM}} \qquad (4)$$

We fit this hybrid model to the data to find the best α for explaining the data and this α quantifies the contribution from the Kalman model.
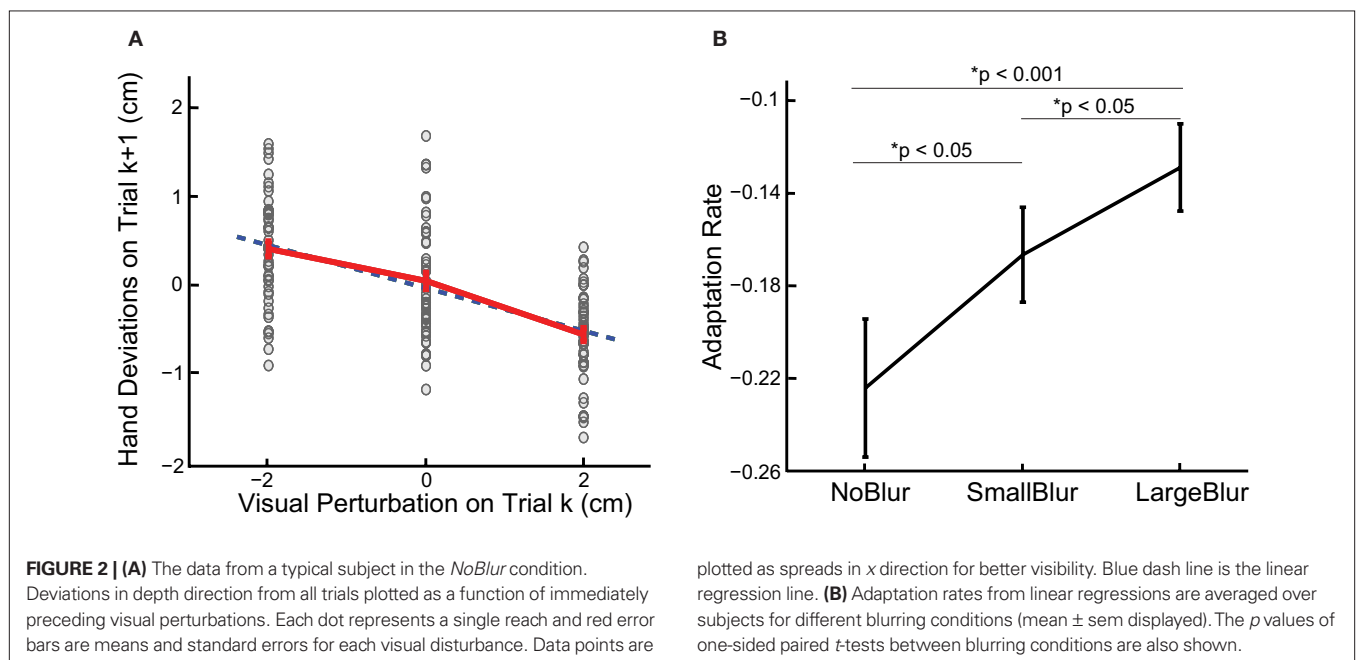
RESULTS

The present study manipulated feedback uncertainty (Experiment 1) and state estimation uncertainty (Experiment 2) during adaptation to trial-by-trial visual perturbations. The subject moved the hand to a target and the feedback about hand position was only shown briefly at the end of each movement. This visual feedback was perturbed spatially and subjects made corrections in the opposite direction of the perturbation during the next trial. The size of this adaptation was calculated as a function of the uncertainty of the visual feedback (Experiment 1) or as a function of state estimation uncertainty (Experiment 2). With these two experiments we can systematically characterize the influence of uncertainty on learning speed.

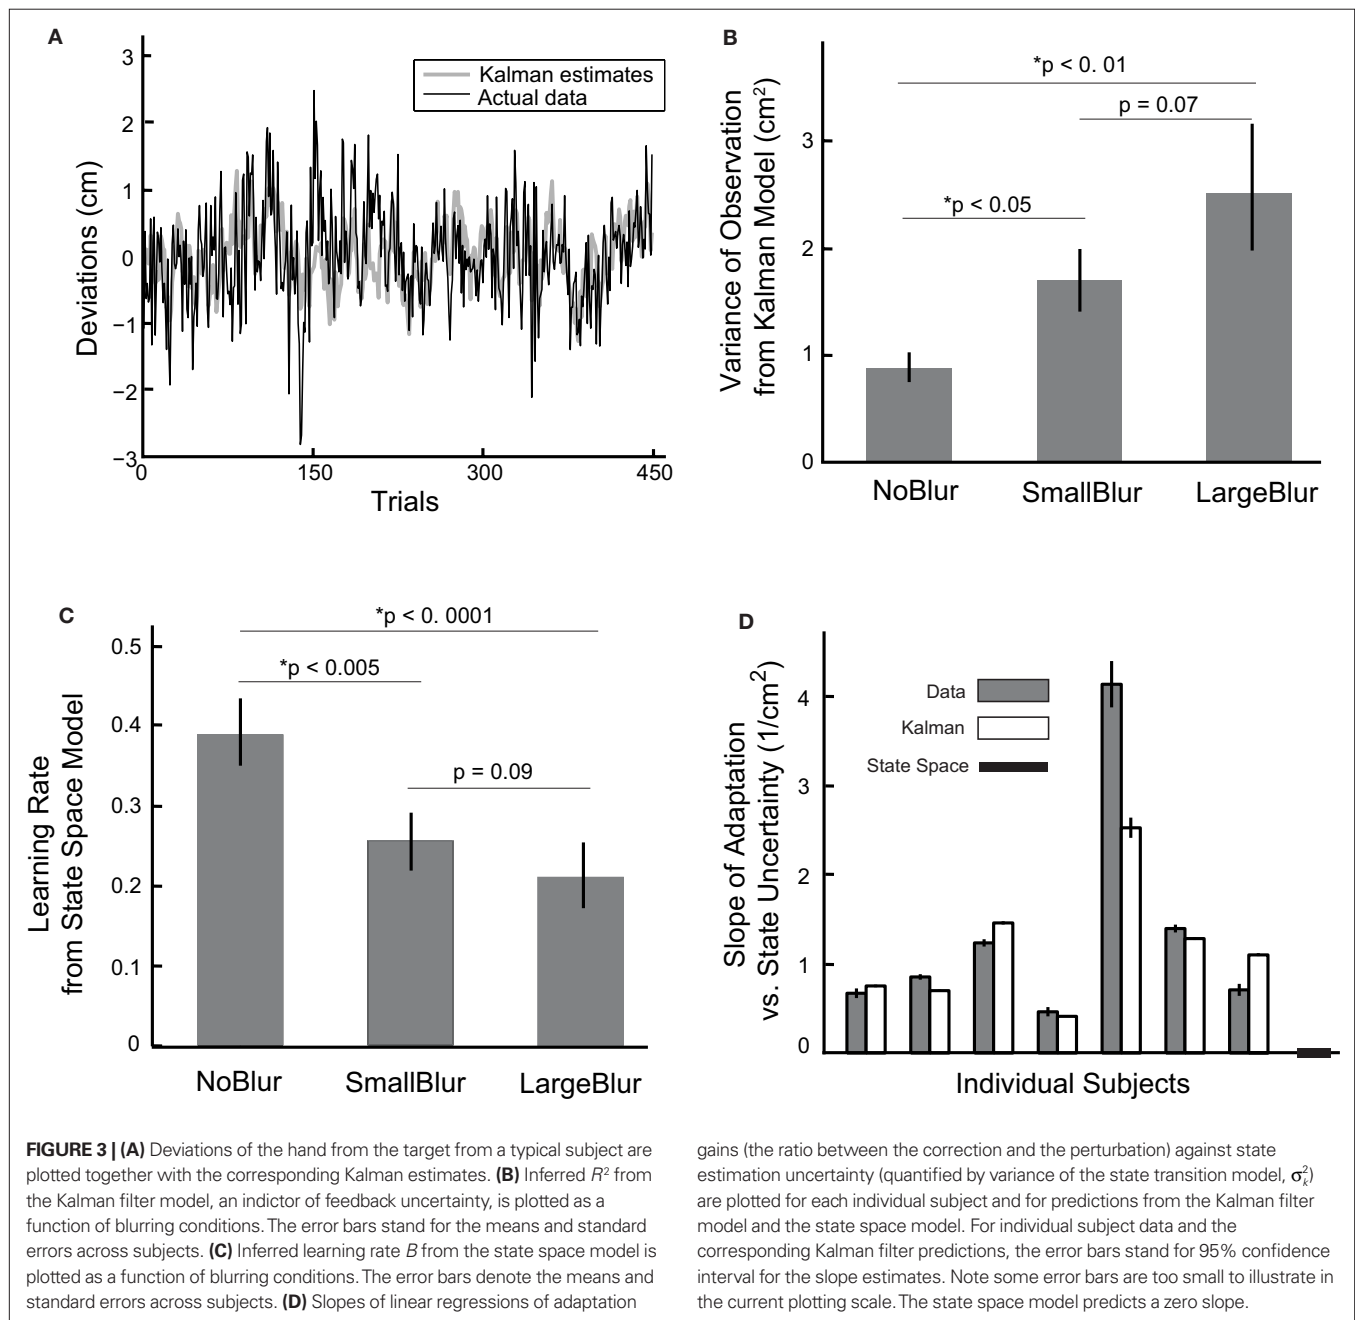### EXPERIMENT 1: FEEDBACK UNCERTAINTY SLOWS DOWN ADAPTATION

The position of visual feedback affects the hand position in the next trial (**Figure 2A**). On average, the deviation of the hand from the target is upward (positive) when the preceding visual perturbation is downward (negative) and vice versa. Thus, the slope of linear regression between the deviation and the perturbation is negative. This is clear evidence of trial-by-trial adaptation to random visual disturbances (Wei and Kording, 2009). The slope can serve as a measure of the rate of adaptation: the more negative the slope, the faster the adaptation to visual perturbations. This slope is the primary measure for adaptation in our analysis and we will call it as *adaptation rate* throughout the paper.

To evaluate adaptation rate as a function of feedback uncertainty, we perform linear regression for each blurring condition separately. All the slopes are negative and significantly different from zero (t-test, $p < 0.0001$ for all conditions), indicating learning for all conditions. Moreover, increased blurring leads to less adaptation: -0.233 ± 0.023, -0.178 ± 0.015 and -0.133 ± 0.017 (mean and sem across subjects, same below) for *NoBlur*, *SmallBlur* and *LargeBlur* conditions, respectively (**Figure 2B**). Adaptation rate is significantly more negative with less blurring (*NoBlur* vs. *SmallBlur*: $p < 0.05$; *NoBlur* vs. *LargeBlur*: $p < 0.001$; *SmallBlur* vs. *LargeBlur*: $p < 0.05$, all tests one-sided paired t-tests). This result serves as direct evidence that feedback uncertainty slows down adaptation.

FIGURE 2 | (A) The data from a typical subject in the *NoBlur* condition. Deviations in depth direction from all trials plotted as a function of immediately preceding visual perturbations. Each dot represents a single reach and red error bars are means and standard errors for each visual disturbance. Data points are plotted as spreads in *x* direction for better visibility. Blue dash line is the linear regression line. **(B)** Adaptation rates from linear regressions are averaged over subjects for different blurring conditions (mean ± sem displayed). The *p* values of one-sided paired *t*-tests between blurring conditions are also shown.

Both of the Kalman model and the state space model provide fairly good fits to the data (see **Figure 3A** for fitting of the Kalman model to the time-series data from a typical subject). More importantly, both models support the prediction that larger feedback uncertainty leads to slower learning. The variance of observation (R^2), which is an indicator of feedback uncertainty and obtained from the fitting of the Kalman model, increases with increasing blurring of visual feedback (**Figure 3B**). This variance is estimated to be 0.88 ± 0.13, 1.71 ± 0.28 and 2.53 ± 0.65 cm^2 for *NoBlur, SmallBlur* and *LargeBlur* conditions, respectively. One-sided paired *t*-tests show that variance of observation is significantly larger with larger feedback uncertainty for two condition comparisons (*NoBlur*

vs. *SmallBlur*: $p < 0.05$; *NoBlur* vs. *LargeBlur*: $p < 0.01$), with the last comparison yielding marginally significance (*SmallBlur* vs. *LargeBlur*: $p = 0.07$). The state space model gives a similar picture as depicted by the adaptation rate results (**Figure 3C**): the learning rate *B* decreases with more blurring. The *B* is estimated to be 0.39 ± 0.04, 0.25 ± 0.04 and 0.21 ± 0.04 for *NoBlur, SmallBlur* and *LargeBlur*, respectively. One-sided paired *t*-tests yield significant differences between *NoBlur* vs. *SmallBlur* ($p < 0.005$) and *NoBlur* vs. *LargeBlur* ($p < 0.0001$). The comparison between *SmallBlur* vs. *LargeBlur* is marginally significant ($p = 0.09$). Regardless the choice of model – feedback uncertainty has a strong effect on the fitted model behavior.

FIGURE 3 | (A) Deviations of the hand from the target from a typical subject are plotted together with the corresponding Kalman estimates. **(B)** Inferred R^2 from the Kalman filter model, an indictor of feedback uncertainty, is plotted as a function of blurring conditions. The error bars stand for the means and standard errors across subjects. **(C)** Inferred learning rate *B* from the state space model is plotted as a function of blurring conditions. The error bars denote the means and standard errors across subjects. **(D)** Slopes of linear regressions of adaptation gains (the ratio between the correction and the perturbation) against state estimation uncertainty (quantified by variance of the state transition model, σ_k^2) are plotted for each individual subject and for predictions from the Kalman filter model and the state space model. For individual subject data and the corresponding Kalman filter predictions, the error bars stand for 95% confidence interval for the slope estimates. Note some error bars are too small to illustrate in the current plotting scale. The state space model predicts a zero slope.

From the model fitting, we can also obtain indirect support of the prediction that larger state estimation uncertainty leads to faster adaptation. State estimation uncertainty is estimated by the Kalman filter, which gives us a trial-by-trial estimate (σ_k). On the other hand, we can calculate the magnitude of adaptation for each trial as a gain factor, taking the ratio between the hand deviation (or the model predicted hand deviation) in one trial and the actual perturbation in the immediately preceding trial. This ratio is also called adaptation gain as it indicates how much adjustment people make relative to a certain size of perturbation. When we regress adaptation gain against state estimation uncertainty we find a significant positive correlation for all the subjects (1.35 ± 0.48 cm^{-2} across subjects; **Figure 3D**, $p < 0.001$). The predicted average slope from the Kalman filter model gives a close match of 1.17 ± 0.26 cm^{-2}. In comparison, the state space model does not incorporate uncertainty and it thus predicts a slope of zero. We thus provide evidence that state estimation uncertainty affects learning speed – an effect that is not predicted by state space models that only take the most recent feedback into account.

We wanted to know if our data is sufficient for a model comparison between the Kalman filter (with 5 free parameters) and the state space model (with 4 free parameters). We compare two models by using cross-validation (10-folds), which does not unjustly benefit models with more free parameters. The Kalman filter model captures $26.6 \pm 4.0\%$ of variance in the data and performs marginally better than the state space model which captures $25.4 \pm 3.9\%$ of variance. However, the difference is not significant ($p = 0.13$, one-sided paired t-test). The data thus does not allow us to directly distinguish these two model classes. We also test if the data can be understood assuming hybrids of Kalman estimates and state space estimates (see Materials and Methods). Indeed, two subjects that are not well predicted by the Kalman model (with small variance explained in cross-validation tests) yield small weights on the Kalman filter. The weight α equals 0.04 and 0.41, respectively, indicating small contribution of state estimation uncertainty. However, across subjects the contribution from the Kalman model is substantial, with α averaging 0.63 ± 0.13. This result suggests that the used strategy varies widely across the population. Subjects tend to estimate state estimation uncertainty based on recent trials as in the Kalman filter and very long durations as implicitly predicted by the state space model.
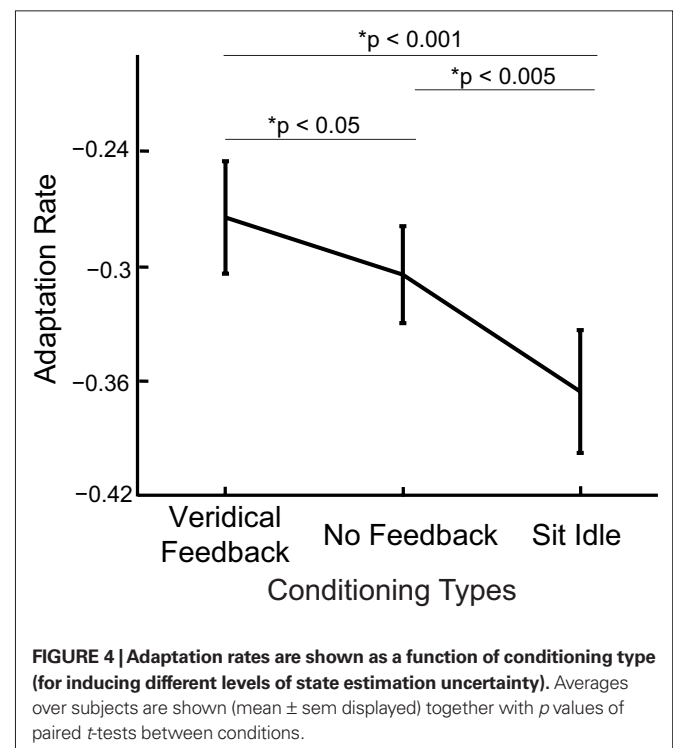
Our regression analysis of state estimation uncertainty with adaptation gain indicates that state estimation uncertainty affects the learning. However, we are unable to distinguish two models as the variance explained by the Kalman model is only marginally better than the state space model. The lack of statistical power for model comparison is not surprising as Experiment 1 is not designed to address this issue of state estimation uncertainty. Because feedback conditions and visual perturbations vary from trial to trial randomly, the state estimation uncertainty only fluctuates in a small range over trials. This results in small observed difference. To make the effect of state estimation uncertainty more pronounced, we employ a novel method to specifically manipulate state estimation uncertainty in Experiment 2.

EXPERIMENT 2: STATE ESTIMATION UNCERTAINTY ACCELERATES ADAPTATION

We directly manipulated state estimation uncertainty by providing feedback of different quality for an extended period of time. During the period of this conditioning, subjects either performed the reaching movement with veridical visual feedback of the hand location, with no visual feedback at all, or simply sat idle with eyes closed. In the first condition, the reaching movement was performed with both visual and proprioceptive feedback of the hand, leading to small state estimation uncertainty. In the second condition, only proprioceptive feedback about the hand location was available. In the condition of sitting idle, the nervous system had no chance to execute related motor commands for reaching movements and received neither visual nor proprioceptive feedback and would thus lead to high state estimation uncertainty. How the resulting state estimation uncertainty influenced adaptation was evaluated by asking subjects to perform the same reaching task with trial-by-trial visual perturbations (as in Experiment 1) in subsequent test blocks.

We assess the adaptation rate in test blocks in the same way as in Experiment 1, i.e., linearly regressing the hand deviations against preceding perturbations. The data from test blocks following each type of conditioning blocks are pooled together for this regression. The adaptation rate is indeed ranked with degree of uncertainty in state estimation (**Figure 4**). The fastest adaptation happens after subjects sit in the dark for a minute (adaptation rate of -0.36 ± 0.03), the second fastest after performing the task without visual feedback (adaptation rate of -0.30 ± 0.02) and the slowest after performing the task with veridical visual feedback (adaptation rate -0.27 ± 0.03). One-sided paired t-tests show that the adaptation rate is significantly faster (more negative) with larger state estimation uncertainty (with veridical feedback vs. without: $p < 0.05$; with veridical feedback vs. sit idle: $p < 0.001$; without feedback vs. sit idle: $p < 0.005$). We can not fit the Kalman filter model to the data of Experiment 2 since there are not enough trials (less than 10) within a test block to allow the model to

FIGURE 4 | Adaptation rates are shown as a function of conditioning type (for inducing different levels of state estimation uncertainty). Averages over subjects are shown (mean ± sem displayed) together with p values of paired t-tests between conditions.

converge. Adaptation rates calculated from the behavioral data show that subjects adapt more when state estimation uncertainty is increased.

DISCUSSION

Our sensorimotor system and the environment we interact with are inherently noisy and our knowledge about them will thus be uncertain. To move precisely, the system not only needs to continuously estimate its state – as investigated by previous studies – but also need to deal with uncertainty in feedback and in estimation. If human sensorimotor control is close to optimal (in a Bayesian sense) as suggested in many previous studies (cf., Körding, 2007), Bayesian theory predicts that less feedback uncertainty and more state estimation uncertainty should make the system rely more on the new sensory feedback and thus exhibit faster learning (Korenberg and Ghahramani, 2002; Körding et al., 2007; Burge et al., 2008; Izawa and Shadmehr, 2008). The present study provides solid evidence to support these two predictions. We tested trial-by-trial adaptation in a simple reaching task and systematically varied these two types of uncertainties. Taken together, the behavioral findings and model comparisons provide support for the qualitative Bayesian predictions about the effect of uncertainty in motor adaptation and sensorimotor control.

Feedback uncertainty in motor adaptation has been addressed in previous studies with Kalman filter models. Korenberg and Ghahramani (2002) proposed that uncertain feedback could lead to slow adaptation in theory but they did not proceed to experimental tests. In a groundbreaking study, Burge et al. (2008) asked about the effect of uncertainty on learning rates using a reaching adaptation task where subjects learn about the environment. The used block design can, however, not readily distinguish between a strategy in which the system learns how to learn (meta learning), and actual use of uncertainty in a trial-by-trial fashion. Our findings complement these studies and produces solid experimental support to the Bayesian view on motor adaptation.

Feedback uncertainty has also been addressed by studies of within-trial behavior. Shadmehr and his colleagues have found that uncertainty in visual feedback influences online control of human reaching movements (Izawa and Shadmehr, 2008). The hand movement was found to be adjusted to a "jumped" target as a function of uncertainty associated with the target representation. This result indicates that uncertainty influences movement control on an even faster time scale as identified in our trial-by-trial adaptation. We expect the dependence of estimation on statistical properties of feedback should operate on multiple time scales, though its effect on longer time scales, such as motor recovery from diseases, motor development or aging, is still in need of experimental verification.

Besides investigating feedback uncertainty, we also studied the effect of state estimation uncertainty. Effect of such uncertainty was first observed in Experiment 1. However, this experiment was originally designed to manipulate feedback uncertainty only and state estimation uncertainty had to be derived from time series data. It thus represents only indirect support of the Bayesian prediction. The more direct test of the hypothesis has been achieved in Experiment 2 where state estimation uncertainty was specified by a novel manipulation. We have found that increased state estimation uncertainty in the nervous system leads to faster adaptation – as predicted by the Bayesian framework.

Our findings provide a novel focus for motor learning studies. Many studies have documented learning speeds for different motor tasks of varying complexity and with varying sensory and motor components (cf., Schmidt and Lee, 2005). Besides these descriptive studies, people have also pursued practical questions such as how to facilitate learning by designing appropriate feedback. On the theoretical side, there are several ideas: the dynamical systems approach emphasizes the process of learning as a change in the degrees of freedom being used (eg. Mitra et al., 1997; Newell et al., 2001), or as exploration and optimization of the stability properties of the task (Müller and Sternad, 2004), or as parameterization of an internal representation of limb dynamics with the interacting environmental forces (Shadmehr and Mussa-Ivaldi, 1994; Thoroughman and Shadmehr, 2000). Our study complements existing theories and indicates that the statistical properties of sensory feedback and state estimation are also determining factors for motor learning (see also: Korenberg and Ghahramani, 2002; Körding et al., 2007). The uncertainty in feedback and in the actor should be taken into account for applied questions in learning and for theoretical developments of motor learning.

The findings from our study might serve as guiding principles for practicing motor learning and rehabilitation. For instance, robotic rehabilitation, where robots assist patients to relearn impaired movement abilities, have recently been gaining popularity (Burgar et al., 2000; Hogan and Krebs, 2004). These rehabilitation programs are usually performed in a virtual reality environment where visual and haptic feedback can be easily manipulated. Our study suggests that by reducing feedback uncertainty and increasing patients' state estimation uncertainty the learning process can be accelerated.

Our findings about the effects of feedback and state estimation uncertainty have direct implications for neurophysiological studies. In recent years researchers have made significant progresses at investigating neural correlates of uncertainty and its computation in the brain (Zemel et al., 1998; Fiorillo et al., 2003; Ma et al., 2006; Preuschoff et al., 2006; Behrens et al., 2007; Deneve, 2008; Rushworth and Behrens, 2008). Some studies examined how sensory stimuli influence decision making in sensorimotor tasks such as hand reaching and saccades (Gold and Shadlen, 2001, 2003; Cisek and Kalaska, 2005; Kiani and Shadlen, 2009). These studies show clear evidence that the nervous system represents and uses feedback uncertainty during sensorimotor tasks. Our reaching task can also be viewed as a decision task. Instead of a categorical choice of decision making, it demands a movement outcome that is defined in a continuous space. Variants of our experiment could be used to probe the representation of feedback uncertainty in continuous tasks in the nervous system. Furthermore, state estimation uncertainty has not been addressed systematically in neurophysiological studies, possibly due to difficulty in manipulating it. Our approach allows defining this uncertainty, by providing sensory feedback of different quality and quantity for certain durations. It can thus serve as a means to assess the representation of state estimation uncertainty. How state estimation uncertainty is represented and how it affects neural computations is central for deepening our understanding of how the nervous system integrates information for sensorimotor control and decision making in dynamical environments.

REFERENCES

Alais, D., and Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Curr. Biol.* 14, 257–262.

Baddeley, R. J., Ingram, H. A., and Miall, R. C. (2003). System identification applied to a visuomotor task: near-optimal human performance in a noisy changing task. *J. Neurosci.* 23, 3066–3075.

Behrens, T., Woolrich, M., Walton, M., and Rushworth, M. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221.

Brown, L., Rosenbaum, D., and Sainburg, R. (2003). Limb position drift: implications for control of posture and movement. *J. Neurophysiol.* 90, 3105–3118.

Burgar, C., Lum, P., Shor, P., and Van der Loos, H. (2000). Development of robots for rehabilitation therapy: the Palo Alto VA/Stanford experience. *J. Rehabil. Res. Dev.* 376, 663–673.

Burge, J., Ernst, M. O., and Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *J. Vis.* 8, 1–19.

Cheng, S., and Sabes, P. N. (2006). Modeling sensorimotor learning with linear dynamical systems. *Neural. Comput.* 18, 760–793.

Cheng, S., and Sabes, P. N. (2007). Calibration of visually guided reaching is driven by error-corrective learning and internal dynamics. *J. Neurophysiol.* 97, 3057–3069.

Cisek, P., and Kalaska, J. F. (2005). Neural correlates of reaching decisions in dorsal premotor cortex: specification of multiple direction choices and final selection of action. *Neuron* 45, 801–814.

Deneve, S. (2008). Bayesian spiking neurons I: inference. *Neural. Comput.* 20, 91–117.

Donchin, O., Francis, J. T., and Shadmehr, R. (2003). Quantifying generalization from trial-by-trial behavior of adaptive systems that learn with basis functions: theory and experiments in human motor control. *J. Neurosci.* 23, 9032–9045.

Ernst, M., and Bülthoff, H. (2004). Merging the senses into a robust percept. *Trends Cogn. Sci. (Regul. Ed.)* 8, 162–169.

Fine, M. S., and Thoroughman, K. A. (2007). Trial-by-trial transformation of error into sensorimotor adaptation changes with environmental dynamics. *J. Neurophysiol.* 98, 1392–1404.

Fiorillo, C. D., Tobler, P. N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. *Science* 299, 1898–1902.

Gold, J. I., and Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci. (Regul. Ed.)* 5, 10–16.

Gold, J. I., and Shadlen, M. N. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *J. Neurosci.* 23, 632.

Haith, A., Jackson, C., Miall, C., and Vijayakumar, S. (2008). Unifying the sensory and motor components of sensorimotor adaptation. **Neural** Information Processing Systems; Vancouver.

Hogan, N., and Krebs, H. (2004). Interactive robots for neuro-rehabilitation. *Restor. Neurol. Neurosci.* 22, 349–358.

Izawa, J., and Shadmehr, R. (2008). On-line processing of uncertain information in visuomotor control. *J. Neurosci.* 28, 11360.

Körding, K. (2007). Decision theory: what "should" the nervous system do? *Science* 318, 606–610.

Körding, K., Tenenbaum, J., and Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nat. Neurosci.* 10, 779–786.

Körding, K. P., and Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature* 427, 244–247.

Kiani, R., and Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science* 324, 759.

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

Korenberg, A. T., and Ghahramani, Z. (2002). A Bayesian view of motor adaptation. *Curr. Psychol. Cogn.* 21, 537–564.

Ma, W. J., Beck, J. M., Latham, P. E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nat. Neurosci.* 9, 1432–1438.

McIntyre, J., Zago, M., Berthoz, A., and Lacquaniti, F. (2001). Does the brain model Newton's laws? *Nat. Neurosci.* 4, 693–694.

Mitra, S., Riley, M. A., and Turvey, M. T. (1997). Chaos in human rhythmic movement. *J. Mot. Behav.* 29, 195–198.

Müller, H., and Sternad, D. (2004). Decomposition of variability in the execution of goal-oriented tasks: three components of skill improvement. *J. Exp. Psychol. Hum. Percept. Perform.* 30, 212–233.

Newell, K. M., Liu, Y. T., and Mayer-Kress, G. (2001). Time scales in motor learning and development. *Psychol. Rev.* 108, 57–82.

Preuschoff, K., Bossaerts, P., and Quartz, S. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51, 381.

Rushworth, M., and Behrens, T. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* 11, 389–397.

Schmidt, R. A., and Lee, T. D. (2005). *Motor Control and Learning: a Behavioral Emphasis.* Champaign, IL: Human Kinetics.

Schofield, A. J., and Georgeson, M. A. (1999). Sensitivity to modulations of luminance and contrast in visual white noise: separate mechanisms with similar behaviour. *Vision Res.* 39, 2697–2716.

Shadmehr, R., and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 14(Pt 2), 3208–3224.

Solomon, J. A. (2002). Noise reveals visual mechanisms of detection and discrimination. *J. Vis.* 2, 105–120.

Solomon, J. A., Lavie, N., and Morgan, M. J. (1997). Contrast discrimination function: spatial cuing effects. *J. Opt. Soc. Am. A* 14, 2443–2448.

Tassinari, H., Hudson, T. E., and Landy, M. S. (2006). Combining priors and noisy visual cues in a rapid pointing task. *J. Neurosci.* 26, 10154–10163.

Thoroughman, K. A., and Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature* 407, 742–747.

Trommershäuser, J., Maloney, L. T., and Landy, M. S. (2003). Statistical decision theory and the selection of rapid, goal-directed movements. *J. Opt. Soc. Am. A* 20, 1419–1433.

van Beers, R. J. (2009). Motor learning is optimally tuned to the properties of motor noise. *Neuron* 63, 406–417.

van Beers, R. J., Sittig, A. C., and Gon, J. J. (1999). Integration of proprioceptive and visual position-information: an experimentally supported model. *J. Neurophysiol.* 81, 1355–1364.

von Helmholtz, H. F. (1863/1954). *On the Sensations of Tone as a Physiological Basis for the Theory of Music*, 2nd edn. New York: Dover Publics.

Wei, K., and Kording, K. (2009). Relevance of error: what drives motor adaptation? *J. Neurophysiol.* 101, 655–664.

Wolpert, D. M., Ghahramani, Z., and Jordan, M. I. (1995). An internal model for sensorimotor integration. *Science* 269, 1880–1882.

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