

On-Line Processing of Uncertain Information in Visuomotor Control

Jun Izawa and Reza Shadmehr

Laboratory for Computational Motor Control, Department of Biomedical Engineering, Johns Hopkins School of Medicine, Baltimore, Maryland 21205

Our sensory observations represent a delayed, noisy estimate of the environment. Delay causes instability and noise causes uncertainty. To deal with these problems, theory suggests that the processing of sensory information by the brain should be probabilistic: to start a movement or to alter it midflight, our brain should make predictions about the near future of sensory states and then continuously integrate the delayed sensory measures with predictions to form an estimate of the current state. To test the predictions of this theory, we asked participants to reach to the center of a blurry target. With increased uncertainty about the target, reach reaction times increased. Occasionally, we changed the position of the target or its blurriness during the reach. We found that the motor response to a given second target was influenced by the uncertainty about the first target. The specific trajectories of motor responses were consistent with predictions of a “minimum variance” state estimator. That is, the motor output that the brain programmed to start a reaching movement or correct it midflight was a continuous combination of two streams of information: a stream that predicted the near future of the state of the environment and a stream that provided a delayed measurement of that state.

Key words: reaction time; uncertainty; autopilot; integration; computational model; motor control

Introduction

Time delay and noise are fundamental properties of information processing in the nervous system. When our sensory apparatus reports an event, that event actually took place some time ago. Forward models were initially proposed as a mechanism by which the nervous system could predict the sensory state (Miall et al., 1993; Wolpert and Miall, 1996). For example, for very brief actions like saccades in which the movement is complete before the arrival of sensory information, motor commands are corrected midflight via a process that predicts the state of the eyes during the movement (Chen-Harris et al., 2008). If we generalize this idea to movements that have a longer duration (beyond 100 ms), we face an interesting problem: the “top-down” stream of predicted sensory states always leads the “bottom-up” stream of delayed sensory observations. Presumably, the brain combines these two streams to form its estimates of sensory states (Vaziri et al., 2006). Because our observations are noisy and our predictions are uncertain, the process of combining the two streams of information should be weighted by our confidence about each source (Körding and Wolpert, 2004). The implication is that, during a single movement, motor output should depend on a continuous, weighted combination of the streams of internal predictions and sensory observations. Here, we report that the motor output during a simple reaching movement indeed carries a signature of this integration process.

Consider a task in which one is reaching to a target but the target jumps to a nearby location. One view is that, after some time delay Δ , the internal representation of the target by the brain will also jump to the new location. In this scenario, the latency of reaction and the rate of change in the motor output are only functions of the visual properties of the second target. However, in the continuous integration model, the brain has predicted the state of the target from time $t - \Delta$ to time t when the sensory system reports that, at time $t - \Delta$, the target jumped to a new location. Rather than discarding the predicted states, one combines the predictions with the observations with a weighting that depends on the uncertainty of each source of information. Because the process of combining the two streams is continuous, the belief about the state of the target gradually converges onto the second target. In this scenario, latency of reacting to the second target is constant, yet the rate of change in the motor output depends strongly on the uncertainties associated with the two targets. For example, the brain should be less willing to change the motor output if the delayed sensory information is less certain than its own predictions.

Here, we present a technique that, on a trial-by-trial basis, allowed us to control the subject’s uncertainties about the sensory feedback and its own predictions. We used Bayesian estimation and optimal feedback control theory (Liu and Todorov, 2007) to predict how the motor output should change as the properties of the reach target changed during a movement and then tested the predictions with behavioral experiments.

Materials and Methods

Twenty-eight subjects (7 males and 21 females; age range, 20–46 years old) held the handle of a robotic arm (Ariff et al., 2002) in a darkened room and reached to “blob”-like visual targets that were projected

Received June 30, 2008; revised Aug. 21, 2008; accepted Sept. 15, 2008.

This work was supported by National Institutes of Health Grants NS037422 and NS057814.

Correspondence should be addressed to Jun Izawa, Department of Biomedical Engineering, Johns Hopkins School of Medicine, 416 Traylor Building, 720 Rutland Avenue, Baltimore, MD 21205-2195. E-mail: jizawa@jhu.edu.

DOI:10.1523/JNEUROSCI.3063-08.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/2811360-09\$15.00/0

(EP739; Optoma) (refresh rate is 70 Hz) onto a horizontal screen directly above the plane of motion. All subjects were healthy and naive about the objective of the experiment. The measured maximum brightness of the projector on the screen was $3500 \text{ cd}\cdot\text{sr}/\text{m}^{-2}$, and the minimum was $0.05 \text{ cd}\cdot\text{sr}/\text{m}^{-2}$. We used Matlab (MathWorks) and Cogent Graphics (University College London, London, UK) to control the visual stimuli. A light-emitting diode (LED) positioned in the handheld handle of the robot provided continuous, real-time feedback of hand position at all times during the reaching movements. Protocols were approved by the Johns Hopkins School of Medicine Institutional Review Board, and all subjects signed a consent form.

Typical trials are shown in Figure 1, *A* and *B*. A trial started when the robot positioned the hand inside a start box ($5 \times 5 \text{ mm}$). A fixation cross was projected at 10 cm with respect to the start box with maximum luminance. This fixation was maintained throughout the trial, and head position was controlled using a chin rest. A blob representing the target of the reach appeared at a random position (distance of 20 cm with respect to start position, at an angle of -15° to 15° with respect to the center line) at a random time (0.5–2.5 s after presentation of the fixation cross). In some trials (see below), the characteristics of the target stimulus, i.e., its center and/or its luminance, changed just as the reach started (hand speed exceeding 3 cm/s). We were interested in quantifying the reaction to this change.

The visual stimulus. Gaussian blobs have been used extensively to control uncertainty about the location of a visual target, particularly in the field of visual psychophysics (Solomon et al., 1997; Schofield and Georgeson, 1999; Solomon, 2002; Baldassi and Burr, 2004; Heron et al., 2004; Tassinari et al., 2006). To produce a blob, we set $[R,G,B] = L'[1,1,1]$, where L' is associated with luminance of the visual signal. We then drew a Gaussian in which luminance was as follows:

$$L'(x,y) = A \frac{1}{\sigma\sqrt{2\pi}} \exp(-((T_x - x)^2 + (T_y - y)^2)/\sigma^2). \quad (1)$$

In this equation, A is the amplitude of the blob, σ is its SD, and (T_x, T_y) is the center of the blob. This parameter A was constant at 27.6 so that the maximum luminance of the smallest ($\sigma = 1.26 \text{ cm}$) blob was 25% of the maximum luminance of the projector. We manipulated the properties of the blob in two ways, by changing its center (T_x, T_y) or by changing its luminance through the parameter σ . For example, we found that, as σ of the blob increased, so did the variance of reach endpoints (see Results), suggesting that the blob was an effective method to vary subject's uncertainty about target position.

Target jump. In one-third of the trials (selected randomly), the blob was erased and a new blob was drawn to either the right or left (with equal probability) of the original target. The amplitude of the shift was $\pm 5^\circ$ ($\pm 1.75 \text{ cm}$ perpendicular to the target direction). The timing of the jump was aligned to reach onset (hand speed exceeding 3 cm/s). Projectors have an intrinsic and sometimes variable delay in responding to commands. Therefore, we used a photo sensor (S1223; Hamamatsu Photonics) to detect the precise moment when the visual information was updated on the screen ($\sim 50 \text{ ms}$ after movement start). Hand position data were aligned by this temporal marker.

Main experiment. Subjects ($n = 11$) performed 720 trials composed of 10 blocks on 2 d. Each block had 72 trials. On 32 trials, the SD of the target changed without a change in its center. On eight trials, the SD did not change but the target center changed. On 16 trials, both the SD and the target center changed. On the remaining 16 trials, the target did not change. In trials in which the SD changed, we set the SD of the first and second targets as follows: [1st SD, 2nd SD] = $[\sigma_S, \sigma_M], [\sigma_L, \sigma_M], [\sigma_M, \sigma_S], [\sigma_M, \sigma_L]$, where $\sigma_S = 3.80 \text{ cm}$, $\sigma_M = 6.32 \text{ cm}$, and $\sigma_L = 7.58 \text{ cm}$ (S, small; M, medium; L, large). For trials in which the SD did not change, we set the SD of the first and second targets as follows: [1st SD, 2nd SD] = $[\sigma_M, \sigma_M]$. For trials in which the target center changed, the blob center was displaced by $\pm 5^\circ$ ($\pm 1.75 \text{ cm}$ with equal probability).

Task instructions. Subjects were instructed that, once the blob appeared, they should react as soon as possible and position the handheld LED at its center. They were told that, occasionally after movement initiation, the blob will move and they should react as soon as possible. At

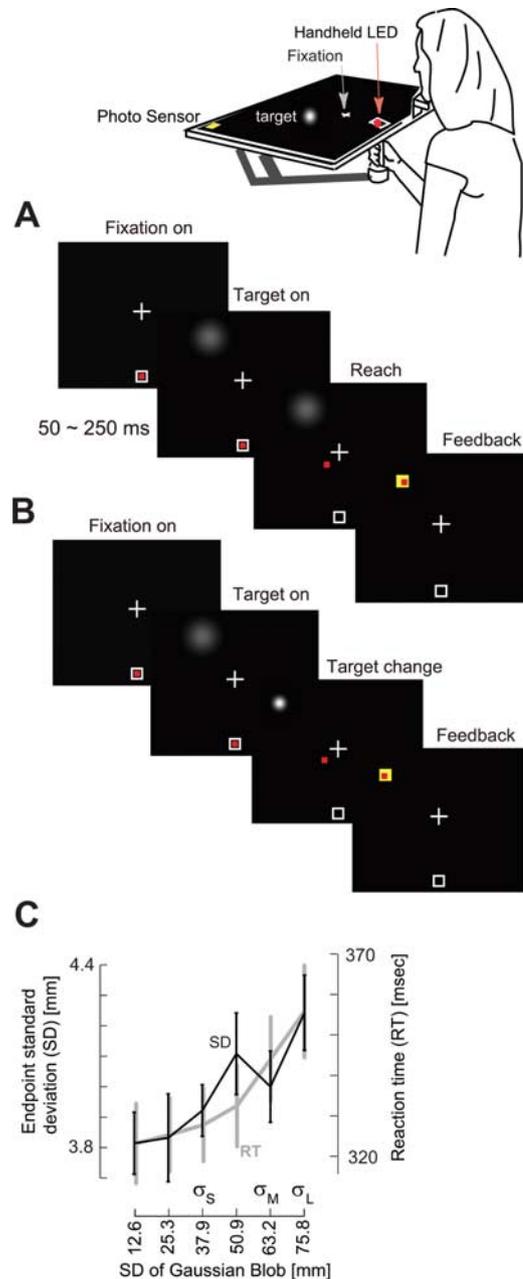


Figure 1. Experimental procedures and results from the control 1 experiment. **A**, Subjects held a handle that housed an LED (red mark) and fixated a cross at 10 cm with respect to the start box. The reach target was a fuzzy blob that appeared at a random position and at a random time (0.5–2.5 s) after the presentation of the fixation cross. Subjects were instructed to start their reach as soon as they saw the target and terminate their reach by placing the LED in the center of the blob. The LED was always on. After reach completion, feedback was provided to indicate distance to the center of the blob. **B**, On a fraction of trials, at $\sim 100 \text{ ms}$ after reach initiation, the center of the blob and/or SD changed. **C**, Endpoint SDs and reaction times (RT) are shown as a function of blob SD in the task shown in **A**. As the target uncertainty increased, reaction times increased, and movements ended with greater spatial variance with respect to target center. Data points are mean and SEM across subjects.

the end of each trial (when the hand speed fell below 3 cm/s), a box of size $5 \times 5 \text{ mm}$ appeared at the center of the blob. If the handheld LED was within that box and if the maximum reach speed was between 0.7 and 1.0 m/s, they were rewarded with a positive score. Otherwise, feedback was provided to increase or decrease speed and the score did not change. Note that the handheld LED was visible throughout the entire experiment. Therefore, the objective was to make a movement that accurately placed the handheld LED at the center of the visual target. The blocks were

separated by a brief break. Before start of data collection, subjects participated in a familiarization block in which target centers occasionally jumped.

Control 1. The objective of this control experiment was to determine whether varying the σ of the blob effectively varied the subject's uncertainty about the center of the target. In the experiment, the target remained stationary and its σ was held constant throughout the trial but changed from trial to trial. We tested six σ values (the values are indicated in Fig. 1C). Subjects ($n = 28$, including the 11 subjects from the main experiment, 11 subjects from control 2 experiment, and 6 new subjects) performed four blocks of 72 trials. In each block, the six blobs were presented 12 times in a randomized order. Our dependent measures were reaction time and endpoint accuracy.

Control 2. In experiment 1, the center of the blob changed at the moment its SD changed. Changes in SD of the blob affected its visual appearance (e.g., its luminance). In this control experiment ($n = 11$ subjects), first, the SD of the blob changed and then, after 100 ms, its center changed. Therefore, unlike the main experiment, in control 2, the luminance values of the first and second targets were identical.

Data analysis. We defined the y -axis as the line that connected the start point and the fixation cross (direction of initial motion of the hand) and the x -axis as a line perpendicular to this axis. To compare reach trajectories with various targets, we rotated each reach so that the line connecting the start position with the initial target was along the y -axis. We excluded reach trials in which hand speed at time of target jump was >2 SDs outside of the mean for that subject.

We categorized data for each subject into 11 sets: no target jump trials (called baseline) and five combinations of the change of SD of the blob for the right and left target jumps. For each subject, we calculated the mean hand path of each category and then subtracted the baseline trajectory. The response to a target jump was quantified as half the difference between right and left jump in the mean hand acceleration along the x -axis (equivalent to mirroring the response to the left jump along the y -axis before averaging it with the right-jump trials). The response onset was defined as the time at which the acceleration along the x -axis exceeded 10 cm/s^2 for at least 100 ms. The slope of the response (its rate of change) was calculated as the difference of acceleration amplitudes at 160 and 140 ms with respect to the onset of the target jump. The maximum response was calculated as the peak value of x acceleration.

Simulations. The change in target position produced a response in the trajectory of the hand. The hypothesis that we wanted to test was that, when the target jumped, the subject's representation of the target position behaved Kalman-like (Kalman, 1960) and gradually moved from its initial position to the observed position. We performed a set of simulations to determine what the hypothesis predicted regarding hand trajectories.

We used the approach suggested by Liu and Todorov (2007) to produce reach trajectories in the context of target jumps. In this framework, the controller is represented as a time sequence of feedback gains. At each instant of time, this gain is applied to the estimate of state, producing a motor command that moves the limb. This state estimate includes the state of the controlled object (the handheld LED) and the target. A critical component of our simulations was the feature that we could manipulate the uncertainty of the state estimates via the blobs that represented reach targets. We wanted to predict how changes in uncertainty should affect the motor reaction to the change in stimulus properties.

System dynamics. We modeled the dynamics of the arm as a point mass in Cartesian coordinates. The state vector is given by $[p_x, \dot{p}_x, p_y, \dot{p}_y, f_x, f_y, T_x, T_y]$, where p was hand position, f was a muscle-like force, and T was target position along the x - and y -axes. We discretized the system dynamics with time step $\delta = 10$ ms. Thus, the system dynamics were as follows:

$$\mathbf{x}^{(k+1)} = A\mathbf{x}^{(k)} + B\mathbf{u}^{(k)} + \boldsymbol{\varepsilon}^{(k)},$$

where

$$A_d = \begin{bmatrix} 1 & \delta & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & \delta/m & 0 \\ 0 & 0 & 1 & \delta & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & \delta/m \\ 0 & 0 & 0 & 0 & 1 - \delta/\tau & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 - \delta/\tau \end{bmatrix} \quad (2)$$

$$A = \begin{bmatrix} A_d & 0 \\ 0 & I_{2 \times 2} \end{bmatrix}$$

$$B = \begin{bmatrix} 0_{4 \times 2} \\ \delta/\tau & 0 \\ 0 & \delta/\tau \\ 0_{2 \times 2} \end{bmatrix}.$$

m is a mass (0.3 kg), τ is a time constant reflecting first-order dynamics of muscles (40 ms), and $\boldsymbol{\varepsilon}$ was Gaussian with zero mean and covariance $\Omega \boldsymbol{\varepsilon} = \text{diag}(0, 0, 0, 0, \sigma_w^2, \sigma_w^2, \sigma_{T_w}^2, \sigma_{T_w}^2, \sigma_{T_w}^2, \sigma_{T_w}^2)$.

To represent time delay in transmission of information, we extended the state vector. An event at time step k was sensed by the first stage of sensory processing and then transmitted to the next stage at time $k + \delta$, where $\delta = 10$ ms. This information continued to propagate and became "observable" with a delay of 100 ms. Thus, we modeled delay by first extending the state vector $\mathbf{x}_e^{(k)} = [\mathbf{x}_0^{(k)}, \mathbf{x}_\delta^{(k)}, \mathbf{x}_{2\delta}^{(k)}, \mathbf{x}_{3\delta}^{(k)}, \dots, \mathbf{x}_\Delta^{(k)}]^T$ and then allowed the system to observe only the most delayed state $\mathbf{x}_\Delta^{(k)}$. The extended "state update equation" was as follows:

$$\mathbf{x}_e^{(k+1)} = \begin{bmatrix} A & 0 & 0 & \dots & 0 & 0 \\ I_{8 \times 8} & 0 & 0 & \dots & 0 & 0 \\ 0 & I_{8 \times 8} & 0 & \dots & 0 & 0 \\ 0 & 0 & I_{8 \times 8} & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & I_{8 \times 8} & 0 \end{bmatrix} \mathbf{x}_e^{(k)} + \begin{bmatrix} B \\ O_{88 \times 2} \end{bmatrix} \mathbf{u}^{(k)} + \boldsymbol{\varepsilon}_e^{(k)}, \quad (3)$$

in which the covariance of the extended state noise is $Q = \text{diag}(0, 0, 0, 0, \sigma_w^2, \sigma_w^2, \sigma_{T_w}^2, \sigma_{T_w}^2, \sigma_{T_w}^2, \sigma_{T_w}^2, 0_{1 \times 88})$. As a result, our "controller" was able to observe the hand position and the target position only in the most delayed state $\mathbf{x}_\Delta^{(k)}$. This idea was reflected in the "measurement equation":

$$\mathbf{y}^{(k)} = H\mathbf{x}_e^{(k)} + \boldsymbol{\omega}^{(k)}$$

$$H = [O_{4 \times 88}, H_n]$$

$$H_n = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}. \quad (4)$$

The measurement noise $\boldsymbol{\omega}$ was zero mean with covariance $R = \text{diag}(\sigma_p^2, \sigma_p^2, \sigma_f^2, \sigma_f^2)$.

State estimation. The objective is to optimally (least variance) estimate the states of our body and the environment despite the noise and the delay in the sensory measurements. The standard Kalman filter provides the solution to our problem. Suppose that, in a trial, the discrete time steps are referred to with variable k . Let us define $\hat{\mathbf{x}}^{(k|k-1)}$ to be a prior state estimate at step k given the prior observations up to and including step $k - 1$. We define $\hat{\mathbf{x}}^{(k|k)}$ to be a posterior state estimate at step k given measurement $\mathbf{y}^{(k)}$. We define the estimation error of $\hat{\mathbf{x}}^{(k|k-1)}$ and $\hat{\mathbf{x}}^{(k|k)}$ as

$$\mathbf{e}^{(k|k-1)} = \mathbf{x}^{(k)} - \hat{\mathbf{x}}^{(k|k-1)}$$

$$\mathbf{e}^{(k|k)} = \mathbf{x}^{(k)} - \hat{\mathbf{x}}^{(k|k)}. \quad (5)$$

Then, the uncertainty of the state estimation is as follows:

$$P^{(k|k-1)} = E[\mathbf{e}^{(k|k-1)} \mathbf{e}^{(k|k-1)T}]$$

$$P^{(k|k)} = E[\mathbf{e}^{(k|k)} \mathbf{e}^{(k|k)T}]. \quad (6)$$

The objective of the Kalman filter, the optimal state estimator, is to update the state estimation with the following equation:

$$\hat{\mathbf{x}}^{(k|k)} = \hat{\mathbf{x}}^{(k|k-1)} + K^{(k)}(\mathbf{y}^{(k)} - H\hat{\mathbf{x}}^{(k|k-1)}), \quad (7)$$

where $K^{(k)}$ is the Kalman gain. The Kalman gain is a function of the uncertainty of the estimated state and the measurement noise such that

$$K^{(k)} = P^{(k|k)} H^T (HP^{(k|k-1)} H^T + \hat{R}^{(k)})^{-1}, \quad (8)$$

where R is a measurement noise. The Kalman gain updates the uncertainty of the estimation:

$$P^{(k|k)} = (I - K^{(k)} H) P^{(k|k-1)}. \quad (9)$$

For this computation, the estimator needs a model of the system dynamics (indicated by hat $\hat{\cdot}$), system noise, and the observation noise:

$$\begin{aligned} \hat{\mathbf{x}}^{(k+1|k)} &= \hat{A}_e \hat{\mathbf{x}}^{(k|k)} + \hat{B}_e \mathbf{u}^{(k)} \\ P^{(k+1|k)} &= \hat{A}_e P^{(k|k)} \hat{A}_e^T + \hat{Q}, \end{aligned} \quad (10)$$

where $\mathbf{u}^{(k)}$ is the efference copy of the motor command.

Optimal feedback control. The target jump task requires subjects to react to the shift of the target. The reaction to this change depends on the “gain” of the sensorimotor feedback loop. Optimal control theory predicts that this gain depends on the accuracy requirements of the task (Todorov and Jordan, 2002; Diedrichsen, 2007; Liu and Todorov, 2007; Chen-Harris et al., 2008; Izawa et al., 2008). We used the theory to formulate the feedback controller and then used it in conjunction with the state estimator to predict hand trajectories.

The accuracy and motor costs of the tasks were specified by

$$J = \sum_{k=0}^{n+r-1} \mathbf{u}^{(k)T} L \mathbf{u}^{(k)} + \mathbf{x}_e^{(k+1)T} S \mathbf{x}_e^{(k+1)}, \quad (11)$$

in which the first term penalizes the motor commands and the second term penalizes the reach/tracking error. L and S weigh each cost term. Because the cost function is quadric, system dynamics are linear, and the noises follow a Gaussian distribution (LQG problem), the problem of the optimal feedback control comes down to the two separate problems: estimation of the gain for the deterministic system (LQ problem) and estimation of Kalman gain. The cost is minimized with the following feedback control law:

$$\mathbf{W}^{(k)} = S^{(k)} \quad (k = n + r)$$

$$\mathbf{W}^{(k-1)} = \hat{A}_e^T (I + \mathbf{W}^{(k)} \hat{B}_e L^{-1} \hat{B}_e^T) \mathbf{W}^{(k)} \hat{A}_e + S^{(k-1)} \quad (k < n + r) \quad (12)$$

$$\mathbf{u}^{(k)} = (L^{(k)} + \hat{B}_e^T \mathbf{W}^{(k+1)} \hat{B}_e)^{-1} \hat{B}_e^T \mathbf{W}^{(k+1)} \hat{A}_e \hat{\mathbf{x}}_e^{(k|k-1)}$$

For the simulation, our cost function was

$$J = w_p (p_x - T_x)^2 + w_p (p_y - t_y)^2 + w_r (u_x^2 + u_y^2), \quad (13)$$

where w_r was set to 0.1^8 throughout the reaching duration, and w_p was 5×0.1^5 between 0 and 500 ms and grew linearly thereafter with slope 5×0.1^4 .

Model parameters. We assumed that, when a target blob was presented, it produced measurement noise with a variance that was scaled by the variance of the blob. We supposed that the scaling factor of the subject’s uncertainty of the target was 3% of SD that we used to project the visual blob. We explored this scaling factor so that the endpoint variance of the simulation could be similar to the experimental data (compare Figs. 1C, 2D). In the simulation of the target jump task, the target position in the state vector was placed at (0, 20) cm and jumped at time step $k = 12$ to (1.75, 20) cm. A target update noise (a random walk model of the target position) represented the increase of the uncertainty of the target position attributable to the target jump, which was $\sigma_{Tu} = 0.0015$. We set the measurement noise of the hand position at $\sigma_p = 0.01$, and the SD of the initial uncertainty of the hand position was 0.0035. All noise factors were

scaled by the time step of the simulation. The complete Matlab code used for all simulations is available at www.shadmehrlab.org.

Results

Subjects moved a handheld LED in a darkened room and placed it as accurately as possible in the center of a blob-like visual target (Fig. 1A). The handheld LED was illuminated at all times. In our control experiment (control 1), we wanted to determine whether the luminance of the visual stimulus (σ in Eq. 1) was an effective method in controlling the subject’s uncertainty about the location of the center of the target. Indeed, we found that, as the luminance properties changed, both the latency of the reaction (i.e., movement start time) and the distribution of the movement endpoints was altered (Fig. 1C). That is, as the σ of the blob increased, the SD of reach endpoints (two-way ANOVA, $F_{(5,135)} = 4.04$; $p < 0.002$), as well as reaction times (two-way ANOVA, $F_{(5,135)} = 23.19$; $p < 0.0001$), tended to increase. However, despite the longer reaction times for blobs with larger σ , the movement speeds (peak acceleration, $F_{(5,135)} = 1.45$; $p > 0.21$) were not affected by the properties of the target. Furthermore, although target uncertainty affected endpoint variability, it did not introduce a bias: reach endpoints on average undershot the target center by only 0.19 mm (about our display accuracy), and changing blob σ did not introduce a change in the endpoint mean ($F_{(5,135)} = 1.38$; $p > 0.2$). It is noteworthy that, although endpoint variability was affected by the σ of the blob, it was not equal to it. Figure 1C suggests that the subject’s uncertainty about target center was approximately an order of magnitude smaller than the displayed σ . Therefore, a blob in a darkened room introduced uncertainty in the subject’s estimate of target position. As the luminance of the blob was changed through manipulation of the parameter σ (Eq. 1), movements exhibited longer reaction times and increased endpoint variance.

The computational problem of reacting to a change in the environment

In our main experiment, as the movement started, we occasionally moved the center of the visual stimulus by a small amount (± 1.75 cm), changed its luminance, or both. In this section, we consider how in theory one should react to these changes.

The objective of the task is to place the handheld LED at the center of the blob. We can represent the problem of estimating the center via the graphical model in Figure 2A. In this figure, the shaded circles are known quantities and the unshaded circles represent unknown quantities. The arrows indicate conditional probabilities. The state variable \mathbf{x} describes the position and velocity of the handheld LED and the position of the center of the blob. Because of delay and noise in our sensory measurements, these states are unavailable to us. When at time t we make a measurement $\mathbf{y}^{(t)}$, the sensory signal reflects a noisy measure of the state of the LED and the blob center at Δ time ago, $\mathbf{x}_\Delta^{(t)}$. The best way (least variance) that we can estimate the current state $\hat{\mathbf{x}}_0^{(t)}$, i.e., the position of the handheld LED and our target, is to take what we measured about the past and then integrate it forward in time with what we predict will be the future of these states to produce an estimate of what the state is at current time.

When the sensory system reports the properties of the target at time $t - \Delta$, this information should be combined with the prior predictions about the immediate future of that target, i.e., predicted states for time $t - \Delta$ to t , written as $\hat{\mathbf{x}}_\Delta^{(t)}, \hat{\mathbf{x}}_{\Delta-\delta}^{(t)}, \dots, \hat{\mathbf{x}}_0^{(t)}$ (Fig. 2B). As a result, when the sensory system reports that the target has jumped (Fig. 2C, gray line, top subplot), our estimate of the target position does not jump but rather converges gradually to

the sensory measurements. The key prediction is that, on each trial, the rate of this convergence will depend on the uncertainties that we have about our measurement and our predictions. The change in the uncertainties dictates how the motor output will change in response to the change in the stimulus.

For example, suppose that the first target has luminance properties described by σ_S , i.e., small SD. As a result, we are fairly certain about this target position (Fig. 1C). Now suppose that the stimulus jumps to a new location and acquires a medium SD σ_M (we refer to this as the S–M condition) (Fig. 2C). Contrast this condition to a situation in which the first target has a large SD σ_L (we refer to this as the L–M condition). One may predict that, because the second target is the same in the S–M, M–M, and L–M trials, the response to the target jump will be similar. That is, during presentation of the second target, the subject will estimate its center and correct the hand trajectory toward it. However, theory predicts that the subject’s estimate of target center will gradually move from the first target to the second target (Fig. 2C, sensory estimate). The rate of this apparent motion will be faster in the L–M condition than in the S–M condition (Fig. 2C, sensory estimate, left column). As a result, the motor response will be larger, producing faster hand acceleration and an earlier peak in the L–M condition than S–M condition (Fig. 2C, left column). Similarly, if the first target has σ_M properties and changes its center and acquires a small SD (M–S condition), the peak hand acceleration should be larger and the timing of the peak should be earlier than the M–M or M–L conditions (Fig. 2C, right column).

Although the model predicts that the reaction to a given change in the stimulus will be affected by the luminance properties of both stimuli, it also predicts that the latency of the reaction will be constant (Fig. 2C, motor response or hand acceleration). This is somewhat puzzling because earlier we saw that the latency of reactions were affected by the luminance of the stimuli (Fig. 1C). Why should there be a difference in reaction times in one case but not the other?

In Figure 2E, we have plotted the uncertainty of the model about the position of the target as a function of time (i.e., Eq. 6). Before the first target is displayed, the uncertainty is very large. When a target with σ_S luminance is displayed, uncertainty of the target rapidly declines as a function of time. Importantly, the rate of decline in the uncertainty is faster for σ_S compared with σ_L . If we assume that a movement starts when uncertainty falls below an arbitrary threshold, then reaction times will be faster for σ_S than σ_L ,

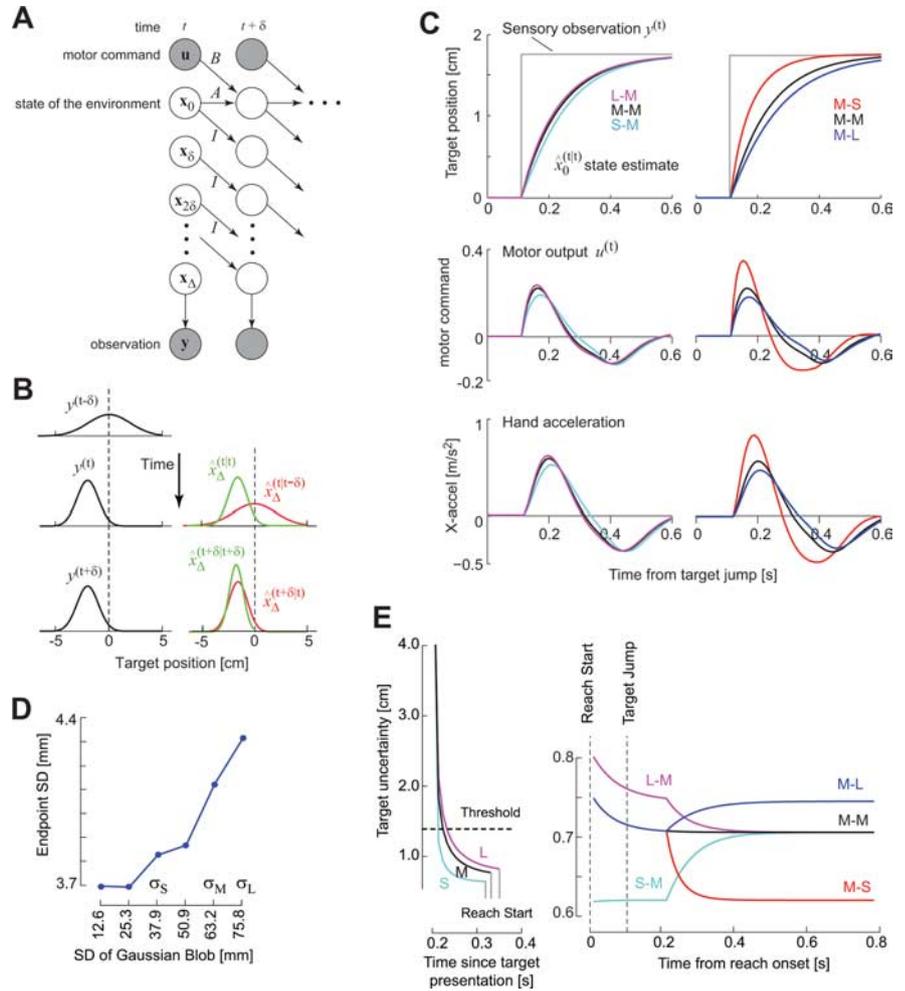


Figure 2. The computational problem of estimating the state of the environment despite delay and noise in the sensory measurements. **A**, A “graphical model” representation of the delay and noise problems. Each circle depicts a random variable. The shaded circles are variables that are observed, and the unshaded circles are unobserved variables that must be estimated. The arrows describe conditional probabilities, or simply causality. The vector \mathbf{x} represents information about state of the hand (position and velocity) and the target (its center). Observation $\mathbf{y}^{(t)}$ is a measure of the delayed state $\mathbf{x}_\Delta^{(t)}$. Given this observation, we need to estimate the current state $\mathbf{x}_0^{(t)}$ and then produce motor commands \mathbf{u} such that the hand arrives in the center of the target. **B**, A simplified example of the estimation problem. The observation suddenly changes at time t . $\mathbf{x}_\Delta^{(t-\delta)}$ is the prior estimate before observation $\mathbf{y}^{(t)}$ and $\hat{\mathbf{x}}_\Delta^{(t)}$ is the posterior estimate after the observation. With repeated observations, the mean of the estimate converges onto the observation. **C**, Simulations of a point mass system controlled by a time-delayed controller. Top row, The target jumps at $t = 0$. This change is observed after a time delay $\Delta = 100$ ms. After the system observes the change in the target position, the estimated target position gradually converges to the observed target position. However, the speed of the convergence is affected by both the uncertainty of the prior (position of the target before it jumped) and the uncertainty of the observation (position of the target after it jumped). If the first target had a small uncertainty with respect to the second uncertainty, then the convergence rate is slow (e.g., S–M condition or M–L condition). If the first target had a large uncertainty with respect to the second uncertainty, then the convergence rate is fast (L–M and M–S conditions). Middle row, The motor command in response to the target jump. The peak response to the target jump is highest for L–M and M–S conditions. The duration of the response (time at which the curve crosses 0) is smallest for the L–M and M–S conditions. Bottom row, Motion of the simulated hand in response to the target jump. The peak response is higher and earlier in the L–M and M–S conditions. **D**, The endpoint SD of the simulated trajectories (500 trials in each condition). **E**, The time evolution of target uncertainty. Left, The uncertainty during the reaction time. Before the first target is displayed, uncertainty is very large. When the first target is displayed, the rate of reduction in the uncertainty is faster for the S target than for M or L targets. If we suppose that the system starts to initiate a reach a constant time after the uncertainty crosses an arbitrary threshold, then reaction times are shorter for S targets than for M or L. Right, The uncertainty during the reach when the target changes. By the time the reach starts, target uncertainty is near asymptotic levels. The uncertainty increases or decreases when the target changes. In the L–M condition, the uncertainty is higher than in the M–M or S–M conditions after the target jump, resulting in a greater reliance in the delayed sensory observations and therefore a faster correction.

just as we saw in our data. Now as the movement is proceeding, the target changes. The uncertainty will increase or decrease, which in turn affects the rate of change in the motor output but not the latency of the reaction because the change in the uncertainties can never

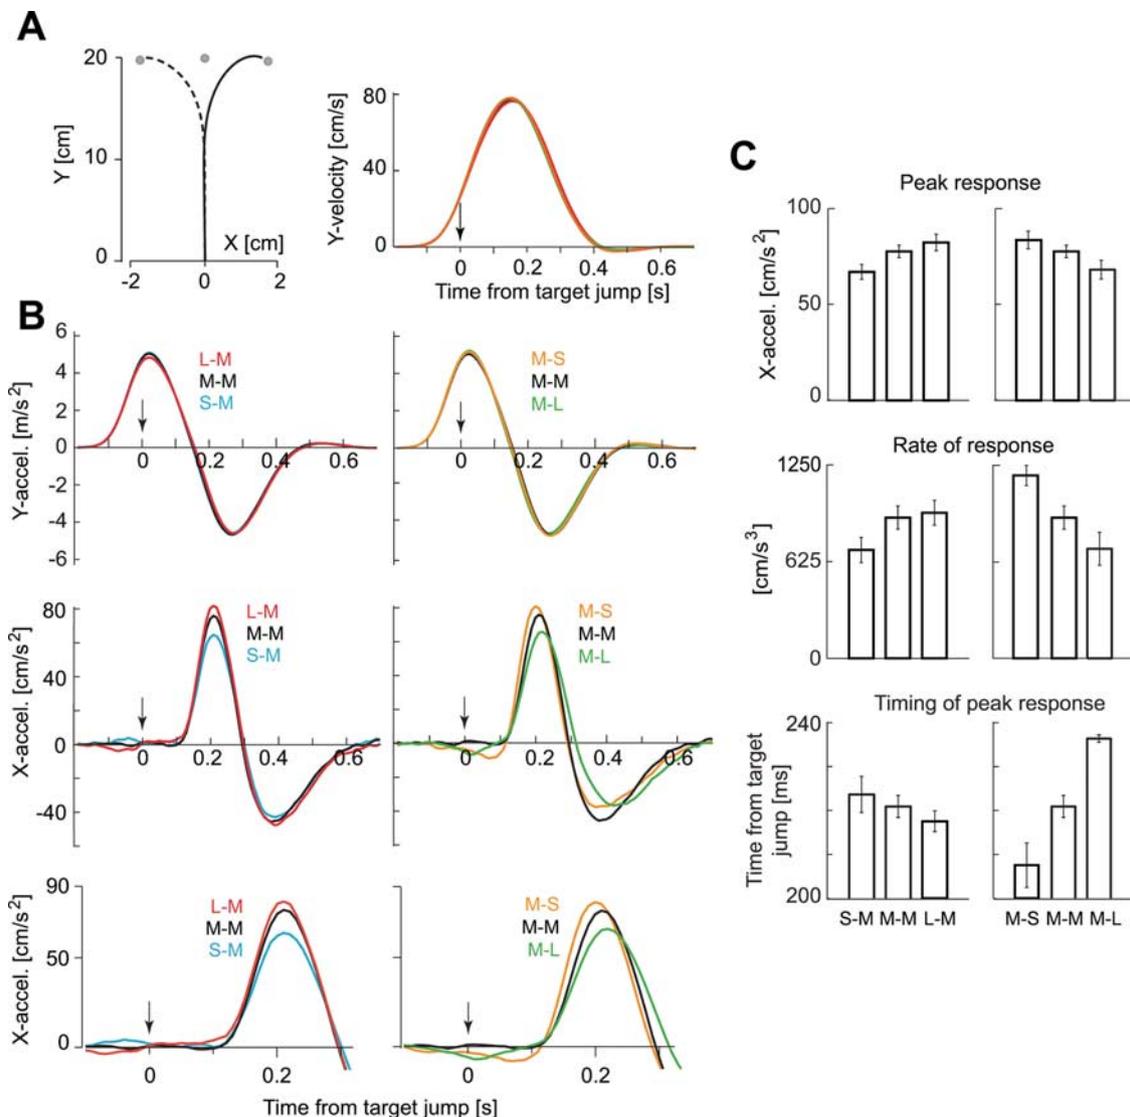


Figure 3. Main experiment. **A**, Left, The mean hand path of the response induced by target jump. The mean was calculated with data across all subjects. Right, Temporal patterns of y velocity. Data was aligned by the onset of target jump. The mean data for all conditions are shown but are almost precisely overlapping. **B**, Top, There were no detectable changes in the y accelerations as a function of target conditions. Middle, However, there were clear changes in the motor response along the x -axis (x acceleration) with changing uncertainty of the target. When the first uncertainty was small with respect to the second target, the motor response was large. When the uncertainty of the first target was large with respect to the second target, the motor response was small. Bottom, An enlarged scale of the motor response (x acceleration). The peak and zero-crossing point were delayed when the peak response increased. **C**, There were main effects associated with target uncertainties in peak response. Top, Rate of initial response (middle) and time of peak response (bottom).

increase beyond the asymptotic value produced by σ_L [presumably, the movement would come to a halt if uncertainty about target location increases above our arbitrary threshold]. Because the asymptotic uncertainty is different for each target luminance, the model also reproduces the observation that varying target luminance affects the variance of the movement endpoints (Fig. 2D).

Reacting to a change in the environment

In the main experiment, during each reach there was a one-third probability of a change in the center of the target and, given this event, a two-thirds probability of SD change. We rotated the reach data so that the first target was along the y -axis (Fig. 3A). As before, we found that, as the uncertainty of the first target decreased, reaction times decreased ($F_{(2,20)} = 32.41$; $p < 0.0001$). That is, movements started earlier when subjects were more certain about the center of the target. However, once the movement begun, the trajectories were indistinguishable, as illustrated by the y velocities and

accelerations in Figure 3, A and B (y peak velocity, $F_{(2,20)} = 1.58$, $p > 0.23$; y peak acceleration, $F_{(2,20)} = 0.77$, $p > 0.47$).

Examination of hand acceleration profiles suggested that the target jump produced a change only in the x -axis of the trajectory of the hand (Fig. 3B). However, the response to a given second target was larger if the first target of that trial had a larger σ . That is, when second target had σ_M characteristics (Fig. 3C, left column), peak acceleration was largest for the L–M condition (main effect of σ , $F_{(2,20)} = 10.8$, $p < 0.001$; L–M vs S–M, t test, $p < 0.001$). Conversely, for a given first target, response to a second target was larger if its σ was smaller than the first target. That is, when the first target had σ_M characteristics (Fig. 3C, right column), peak acceleration was largest for the M–S condition (main effect of σ , $F_{(2,20)} = 9.3$, $p < 0.002$; M–S vs M–L, t test, $p < 0.01$).

We next computed the rate of change in hand acceleration (the derivative at 150 ms post-target jump). When the second target had σ_M characteristics, the rate of change in acceleration

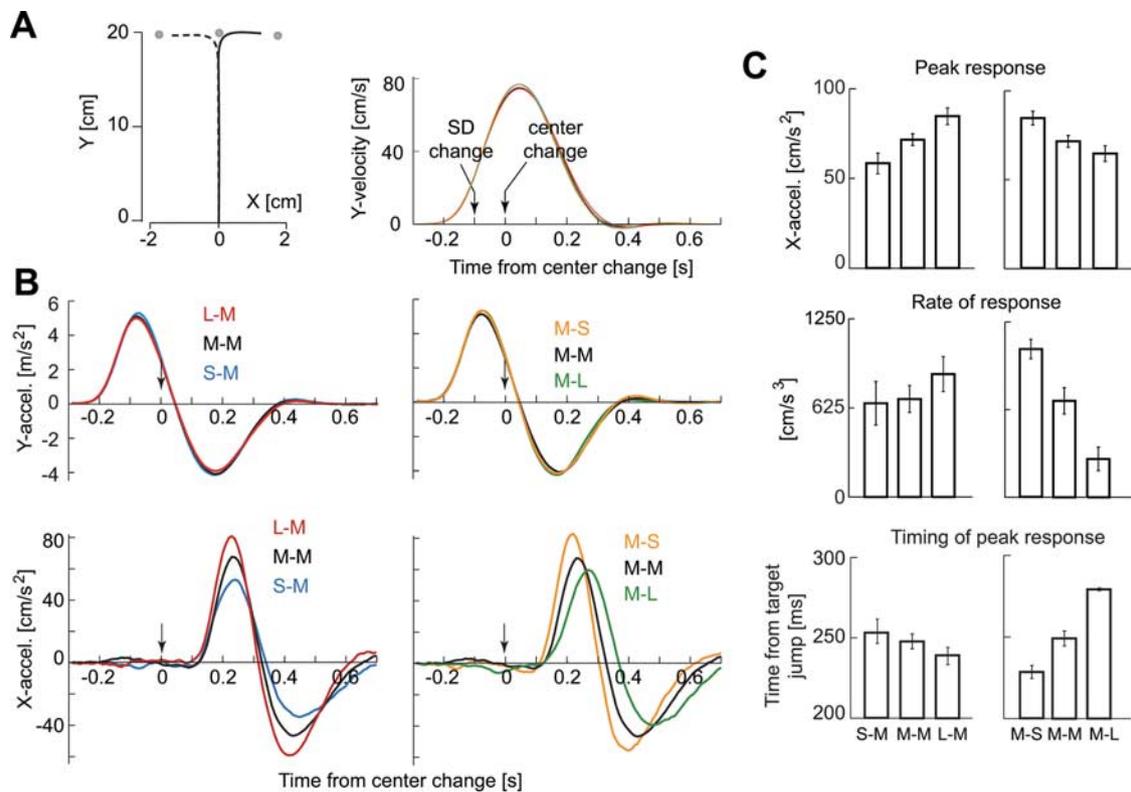


Figure 4. Control experiment 2. **A**, Left, Mean hand paths of the response induced by target jump. The mean was calculated with data across all subjects. Right, Temporal patterns of y velocity. **B**, Top, The response along the y -axis was indistinguishable across various conditions. Bottom, However, changing the target uncertainty produced clear variations in the response along the x -axis. When the uncertainty associated with the first target was small with respect to the second target, the motor response to the second target was large. When the uncertainty associated with the first target was large with respect to the second target, the response was small. The peak and zero-crossing points were delayed when the peak response increased. **C**, There were main effects associated with target uncertainties in peak response (top), rate of initial response (middle), and time of peak response (bottom).

was largest in the L–M condition (main effect of σ , $F_{(2,20)} = 5.5$, $p = 0.01$; L–M vs S–M, t test, $p < 0.001$). When the first target had σ_M characteristics, the rate of change was largest for the M–S condition (main effect of σ , $F_{(2,20)} = 17.2$, $p < 0.001$; M–S vs M–L, t test, $p < 0.001$).

We also computed the timing of the peak acceleration along the x -axis for a given second target. This measure was earliest if the σ of that target happened to be smaller than the σ of the target that preceded it. That is, when the second target had σ_M characteristics, the timing of the peak response was earliest in the L–M condition and latest in the S–M condition (main effect of σ , $F_{(2,20)} = 5.1$, $p = 0.016$; L–M vs S–M, t test, $p = 0.011$). When the first target had σ_M characteristics, the timing of the peak response was earliest in the M–S condition (main effect of σ , $F_{(2,20)} = 21.2$, $p < 0.0001$; M–S vs M–L, t test, $p < 0.001$).

All of these changes are consistent with the predictions of the theory. Furthermore, as the theory had predicted, there were no main effects of combinations of target uncertainty in the latency of the reaction to the jumped target ($F_{(4,40)} = 0.81$; $p = 0.526$; mean is 128 ms).

In summary, when during a reaching movement the target moved by a small distance and changed its luminance to affect the uncertainty with which the subject could estimate its center, the response (hand acceleration) to the second target became stronger, its rate of change became greater, and its peak shifted earlier as the uncertainty associated with the first target increased. For a given first target, the response to a second target became stronger, its rate of change became greater, and the peak acceleration occurred earlier as the uncertainty of the second target decreased.

Confounding factors

A concern is that, in the S–M, M–M, and L–M conditions, the response to the second target might have been different because of the response to the first target. Recall that the S target provided a more certain center location, and this might encourage a faster movement that produced a faster response in the S–M versus L–M condition. This prediction, however, is opposite of what we found (the response to the second target was fastest in the L–M condition). Regardless, we checked whether properties of the first target affected hand accelerations toward it and found no evidence to support this idea: in the control 1 experiment noted above, peak accelerations were unaffected by the changes in σ ($F_{(5,135)} = 1.6$; $p > 0.15$). This was also confirmed in the main experiment ($F_{(3,20)} = 2.22$; $p > 0.13$), as demonstrated by the accelerations plotted in Figure 3B.

Another concern is that the change of the luminance at the timing of the target jump may alter the subject's perception about the intensity of the target jump significantly. In other words, if we replot the dependent variables in Figure 3C as a function of change in target luminance (but not luminance), the trends are unchanged. To investigate this, we performed a second control experiment (control 2) in which, after a target was displayed, in a fraction of trials its SD changed for 100 ms, and then its center moved. In other words, the luminance of the target was identical before and after the change in its center. Figure 4 displays the results of this control experiment. We found that we were able to reproduce all the results of the main experiment. When the second target had σ_M characteristics (Fig. 4C, left column), peak acceleration, the rate of change in acceleration,

and timing of peak response were all significantly affected by the SD of the first target (paired t tests, S–M vs L–M; peak acceleration, $p < 0.001$; rate of acceleration, $p < 0.05$; timing, $p < 0.05$). Similarly, when the first target had σ_M characteristics (Fig. 4C, right column), peak acceleration, the rate of change in acceleration, and timing of peak response were all significantly affected by the SD of the second target (paired t tests, M–S vs M–L; peak acceleration, $p < 0.001$; rate of acceleration, $p < 0.001$; timing, $p < 0.001$).

Discussion

Our sensory system presents us with a delayed, noisy view of the world. To deal with delay, our brain might continuously predict the near future of sensory states. To deal with noise, our brain might assign measures of confidence to both our predictions and our observations so that, by integrating the two streams of information, we continuously have the best estimate of the state of the world. According to this theory, even during a single movement our motor response to a change in a sensory stimulus should carry the signature of the relative confidences that the brain has assigned to the two streams of information. Our goal here was to predict what this signature might look like and then see whether in healthy subjects the motor output depended on a continuous integration of the two hypothetical streams.

We considered a simple reaching task in a darkened room in which the goal was to place a handheld LED (always on) at the center of a blob. The blob was a collection of pixels with intensities described by a Gaussian with SD σ . Theory predicted that, as the sensory system continuously reported the state of the stimulus, uncertainty about its center would decrease (as a function of time). The rate of this decrease would be more rapid for a blob that had a smaller σ than for a blob with a larger σ . Suppose that the brain starts a movement when uncertainty about the location of the target falls below an arbitrary threshold. In that case, the reaction times should be faster, i.e., movements should start earlier toward a target for which one has greater confidence than one for which one is uncertain. Indeed, we found that, by increasing σ of the visual stimulus, subjects increased their reaction time and their endpoint variance. Therefore, it appeared that, by manipulating the visual stimulus, we increased the uncertainty in the subject's estimate of target center.

We next considered how the subjects reacted to a change in the stimulus as the movement was unfolding. The idea was to start the movement with a stimulus that had a given uncertainty, presumably affecting the uncertainty with which the brain predicted the near future of that stimulus. Next, we changed the stimulus in mid-movement, producing a difference in the delayed sensory measurement (bottom-up stream) with respect to its predicted future (top-down stream). Theory predicted that the response to the change in the target would depend on the relative uncertainties associated with the two streams. That is, despite the fact that the sensory stream reported a jump in the target position, the subject's estimate of target position would only gradually converge to the sensory observation with a rate that was a direct measure of the relative confidences that the brain assigned to the two streams.

Indeed, we found that the response (hand acceleration) to a given second target was larger, had a faster rate of change, and an earlier peak if the first target had a larger σ . For a given first target, the response to the second target was larger, had a faster rate of change, and an earlier peak if the second target had a smaller σ . Therefore, people not only delayed their movement initiation when they were uncertain about the target of the movement, but, during the movement, they made more forceful corrections when a change in the target decreased their uncertainty.

These results appeared unrelated to low-level visual properties

such as absolute luminosity of the second target. In a control experiment in which we decoupled the time at which σ changed (affecting luminosity) from the time at which the center changed, the main results remained unchanged.

When one is reaching to a target and the target changes, the simplest view is that the brain uses the information about the new target to correct the movement. This would predict that the response to a given second target should depend only on its properties. Our results in Figure 3A reject this view. An alternative view is that, when the target jumps, the new estimated target position is a weighted combination of the first and second targets. Such a single-step integration model was used previously to explain corrections in response to brief sensory feedback (Körding and Wolpert 2004). Although this approach may be reasonable for conditions in which sensory feedback is very brief, its application here would predict that reach endpoints should be biased by the first target, something that we never observed (supplemental Fig. S1B, available at www.jneurosci.org as supplemental material). Finally, another view may be that our problem is akin to a decision-making process in which the evidence for the new target location is compared with the evidence for the old. In this process, one reacts to the new target when the ratio of its evidence to the old exceeds some threshold (Carpenter and Williams, 1995; Kim and Shadlen, 1999; Reddi and Carpenter, 2000), predicting that the latency of the corrective movement should depend on the relative uncertainties associated with the internal representations of each target (supplemental Fig. S1A, available at www.jneurosci.org as supplemental material). However, we did not see a change in timing of the start of the corrective movements. Rather, although the correction started at a constant time after the target changed, the rate of the change in the motor output depended on the properties of both the first and the second targets. If we imagine that the brain predicts the future properties of the target and combines it with the delayed measurements, the continuous integration of these two streams predicts the effects observed in our data.

In one respect, however, the model failed to predict an aspect of the data. The simulations predicted that we should see a greater range of responses for M–L and M–S conditions versus S–M and L–M conditions (Fig. 2C, compare right and left columns). In fact, the range of responses was similar (Fig. 3B, right and left columns). The reason for this is that, in our simple integration model (Kalman filter), the target jump produced a fairly rapid convergence of uncertainty to the noise inherent in the observation. That is, in our model, the “prior belief” rapidly converged to the sensory feedback (Fig. 2E). The data, however, suggest that this convergence took longer than expected. This failure of the model is not attributable to assumptions regarding parameter values but rather the mechanism with which one updates one's uncertainty. The reason for this discrepancy in the speed of the uncertainty update is unknown to us.

Regardless, our simulations predicted and our experiment confirmed that the convergence of the internal estimates to the delayed sensory measures must take much longer than 100 ms (Fig. 2C). The idea that the integration period for a step change in the sensory information is longer than 100 ms is directly supported by the neurophysiological experiments in which monkeys attempt to guess the direction of motion of a visual scene. Huk and Shadlen (2005) found that a 100 ms step change in the ongoing visual scene produced changes in neural discharge that after a 220 ms delay, gradually increased over a 200 ms period before declining back to baseline after an additional 300 ms. Why does it take so long for new information to be integrated with old?

In any time-delayed system in which measurements at time t inform us about state of the world at time $t - \Delta$, if we could predict the states from time $t - \Delta$ to t , then the delayed observation can be

integrated with predictions and then projected forward to produce an updated estimate of state at time t . This is advantageous because it allows us to change our beliefs about the current state despite the fact that we only have a measurement about the past. The disadvantage, however, is that, in our response to a step change in the environment, we are unable to change our estimates in a step. Rather, the rate of this gradual change depends on integration of our predictions about the near future with the stream of delayed sensory observations. Our results here demonstrate that the rate of correction in response to a change in the stimulus is a proxy for the process in which the brain integrates real-time sensory information about the past with prior predictions about its future.

Previous works have demonstrated that, as the uncertainty of the measured sensory information increases, the brain increasingly relies on a prior. For example, Körding and Wolpert (2004) showed that, during a reach, subject's reaction to a given sensory feedback depended on the long-term history of the subject's experience in that task. In their example, it was assumed that the prior predictions are combined with current observations in a single integration step. Although this is a reasonable approximation, it is difficult to directly apply it to more realistic scenarios in which there is a continuous stream of delayed sensory feedback, requiring continuous processing of that information. For example, when the target jumps, people do not reach to a point somewhere in between the first and the second targets. Rather, the timing and rate of change of their motor reaction depends on the properties of the two targets, suggesting that there are continuous changes in the prior and continuous integration of it with the sensory information.

The reaction time that we observed (~ 128 ms) is in the range of movements termed "autopilot" (Prablanc and Martin, 1992; Brenner and Smeets, 1997; Day and Lyon, 2000; Saunders and Knill, 2004). Damage to the posterior parietal cortex extinguishes the autopilot adjustment (Desmurget et al., 1999; Pisella et al., 2000; Grea et al., 2002). A change in the reach target engages the posterior parietal cortex (Diedrichsen et al., 2005), and a single pulse of transcranial magnetic stimulation to the anterior interparietal area slows the rate of adjustments in the grip aperture when the target of the reach changes orientation (Tunik et al., 2005). Indeed, we (Shadmehr and Krakauer, 2008) and others (Deneve et al., 2007) have speculated that one of the major functions of the posterior parietal cortex is to integrate the predicted and observed sensory information, changing the integration ratio as a function of the uncertainties of the two streams. Uncertainty of visual information is reflected in the discharge of cells (Gold and Shadlen, 2001; Huk and Shadlen, 2005), and the rate of increase in the discharge reflects a process of accumulating information about the stimulus. Once the discharge reaches a threshold, monkeys start their movement. In our task, we found that people started their movements later when they were less certain of the target's center, which is consistent with the idea that the process of accumulating information had to reach a threshold before initiation of action. A strong prediction of our results is that, in experiments that include a sudden change in the stream of sensory information (Huk and Shadlen, 2005), the change in activity of cells in lateral interparietal area in response to a given change in the stimulus will depend on both the uncertainties of the new stimulus and the uncertainties of the predictions regarding the prior stimulus.

References

- Ariff G, Donchin O, Nanayakkara T, Shadmehr R (2002) A real-time state predictor in motor control: study of saccadic eye movements during unseen reaching movements. *J Neurosci* 22:7721–7729.
- Baldassi S, Burr DC (2004) "Pop-out" of targets modulated in luminance or colour: the effect of intrinsic and extrinsic uncertainty. *Vision Res* 44:1227–1233.
- Brenner E, Smeets JB (1997) Fast responses of the human hand to changes in target position. *J Mot Behav* 29:297–310.
- Carpenter RH, Williams ML (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62.
- Chen-Harris H, Joiner WM, Ethier V, Zee DS, Shadmehr R (2008) Adaptive control of saccades via internal feedback. *J Neurosci* 28:2804–2813.
- Day BL, Lyon IN (2000) Voluntary modification of automatic arm movements evoked by motion of a visual target. *Exp Brain Res* 130:159–168.
- Denève S, Duhamel JR, Pouget A (2007) Optimal sensorimotor integration in recurrent cortical networks: a neural implementation of Kalman filters. *J Neurosci* 27:5744–5756.
- Desmurget M, Epstein CM, Turner RS, Prablanc C, Alexander GE, Grafton ST (1999) Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nat Neurosci* 2:563–567.
- Diedrichsen J (2007) Optimal task-dependent changes of bimanual feedback control and adaptation. *Curr Biol* 17:1675–1679.
- Diedrichsen J, Hashambhoy Y, Rane T, Shadmehr R (2005) Neural correlates of reach errors. *J Neurosci* 25:9919–9931.
- Gold JI, Shadlen MN (2001) Neural computations that underlie decisions about sensory stimuli. *Trends Cogn Sci* 5:10–16.
- Grea H, Pisella L, Rossetti Y, Desmurget M, Tilikete C, Grafton S, Prablanc C, Vighetto A (2002) A lesion of the posterior parietal cortex disrupts on-line adjustments during aiming movements. *Neuropsychologia* 40:2471–2480.
- Heron J, Whitaker D, McGraw PV (2004) Sensory uncertainty governs the extent of audio-visual interaction. *Vision Res* 44:2875–2884.
- Huk AC, Shadlen MN (2005) Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J Neurosci* 25:10420–10436.
- Izawa J, Rane T, Donchin O, Shadmehr R (2008) Motor adaptation as a process of reoptimization. *J Neurosci* 28:2883–2891.
- Kalman RE (1960) A new approach to linear filtering and prediction problems. *Trans ASME J Basic Engineering* 82:35–45.
- Kim JN, Shadlen MN (1999) Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2:176–185.
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247.
- Liu D, Todorov E (2007) Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27:9354–9368.
- Miall RC, Weir DJ, Wolpert DM, Stein JF (1993) Is the cerebellum a Smith predictor? *J Mot Behav* 25:203–216.
- Pisella L, Grea H, Tilikete C, Vighetto A, Desmurget M, Rode G, Boisson D, Rossetti Y (2000) An "automatic pilot" for the hand in human posterior parietal cortex: toward reinterpreting optic ataxia. *Nat Neurosci* 3:729–736.
- Prablanc C, Martin O (1992) Automatic control during hand reaching at undetected two-dimensional target displacements. *J Neurophysiol* 67:455–469.
- Reddi BA, Carpenter RH (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–830.
- Saunders JA, Knill DC (2004) Visual feedback control of hand movements. *J Neurosci* 24:3223–3234.
- Schofield AJ, Georgeson MA (1999) Sensitivity to modulations of luminance and contrast in visual white noise: separate mechanisms with similar behaviour. *Vision Res* 39:2697–2716.
- Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. *Exp Brain Res* 185:359–381.
- Solomon JA (2002) Noise reveals visual mechanisms of detection and discrimination. *J Vis* 2:105–120.
- Solomon JA, Lavie N, Morgan MJ (1997) Contrast discrimination function: spatial cuing effects. *J Opt Soc Am A Opt Image Sci Vis* 14:2443–2448.
- Tassinari H, Hudson TE, Landy MS (2006) Combining priors and noisy visual cues in a rapid pointing task. *J Neurosci* 26:10154–10163.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.
- Tunik E, Frey SH, Grafton ST (2005) Virtual lesions of the anterior intraparietal area disrupt goal-dependent on-line adjustments of grasp. *Nat Neurosci* 8:505–511.
- Vaziri S, Diedrichsen J, Shadmehr R (2006) Why does the brain predict sensory consequences of oculomotor commands? Optimal integration of the predicted and the actual sensory feedback. *J Neurosci* 26:4188–4197.
- Wolpert DM, Miall RC (1996) Forward models for physiological motor control. *Neural Netw* 9:1265–1279.