The statistical determinants of adaptation rate in human reaching

Johannes Burge
UC Berkeley, School of Optometry, Berkeley, CA, USA

Marc O. Ernst
Max Planck Institute for Biological Cybernetics, Tübingen, Germany

Martin S. Banks
UC Berkeley, School of Optometry, Berkeley, CA, USA

Rapid reaching to a target is generally accurate but also contains random and systematic error. Random errors result from noise in visual measurement, motor planning, and reach execution. Systematic error results from systematic changes in the mapping between the visual estimate of target location and the motor command necessary to reach the target (e.g., new spectacles, muscular fatigue). Humans maintain accurate reaching by recalibrating the visuomotor system, but no widely accepted computational model of the process exists. Given certain boundary conditions, a statistically optimal solution is a Kalman filter. We compared human to Kalman filter behavior to determine how humans take into account the statistical properties of errors and the reliability with which those errors can be measured. For most conditions, human and Kalman filter behavior was similar: Increasing measurement uncertainty caused similar decreases in recalibration rate; directionally asymmetric uncertainty caused different rates in different directions; more variation in systematic error increased recalibration rate. However, behavior differed in one respect: Inserting random error by perturbing feedback position causes slower adaptation in Kalman filters but had no effect in humans. This difference may be due to how biological systems remain responsive to changes in environmental statistics. We discuss the implications of this work.

Keywords: visuomotor, reaching, adaptation, recalibration, Kalman filter, cue combination

Citation: Burge, J., Ernst, M. O., & Banks, M. S. (2008). The statistical determinants of adaptation rate in human reaching. *Journal of Vision, 8*(4):20, 1–19, http://journalofvision.org/8/4/20/, doi:10.1167/8.4.20.

Introduction

Consider a reaching movement too fast for visual feedback to adjust the movement in flight. Such open-loop reaching requires an appropriate mapping between the visual estimate of the target's location and the motor output required to arrive there. In everyday circumstances, the relationship between perceived location and appropriate motor command may change. Maintenance of an appropriate mapping therefore requires continual updating; the updating is called visuomotor calibration.

Visuomotor calibration was first studied by Helmholtz (von Helmholtz, 1867). He presented a visible target and subjects reached to it, a task they accomplished without difficulty. The visuomotor mapping was then altered with a prism that shifted the apparent direction of the target rightward. Initial reaches after the prism was introduced were in the visually specified direction, missing the target to the right. However, after only a few reaches, subjects corrected and reached to the true location.

Recalibration is integral to the proper functioning of perceptual and perceptual–motor systems. Its time course varies greatly from one system to another. In an attempt to understand how and why time course varies, we describe an optimal adaptor, show how its calibration rate is affected by various stimulus properties, and then compare its pattern of rate changes to human performance.

We assume that the goal of calibration is to maximize accuracy and hence that calibration involves error minimization. Consider a target whose location relative to the body is L. Reaching toward the target requires a visual estimate of its location and a movement that will bring the hand to that location. The accuracy of the reach can be assessed by comparing the visual estimates of the target location, $\hat{L}_V$, and the reach endpoint feedback, $\hat{F}_V$. If the subject observes an error, $\hat{E}$, it could be systematic due to miscalibration (as when first wearing a prism), it could be random (such as motor noise), or it could be both. (Proprioception can also provide feedback (Sober & Sabes, 2003, 2005; van Beers, Wolpert, & Haggard, 2002). However, the experimental setup and the task instructions were designed to make proprioception relatively unuseful. Subjects reached in a plane orthogonal to the plane in which the visual feedback appeared, so the locations of proprioceptive and visual feedback were very different. Subjects were also instructed to minimize the discrepancy between the visual feedback and visual target

doi: 10.1167/8.4.20 Received August 18, 2007; published April 23, 2008 ISSN 1534-7362

location. We therefore ignored proprioception in our analysis. But it is nevertheless possible that such input influenced behavior. The consequences of this possibility are discussed in the Methods section.

When error is due only to miscalibration (non-zero mean), the best strategy is straightforward: measure the error signal, $\hat{E} = \hat{F}_V - \hat{L}_V$, and adjust the visual location estimate and/or the motor command to make the error zero on the next trial:

$$\hat{X}_{t+1} = \hat{X}_t + \hat{E}_t, \tag{1}$$

where $\hat{E}_t$ is the observed reach error for the current reach, $\hat{X}_t$ is the visuomotor mapping estimate on the current reach, and $\hat{X}_{t+1}$ is the visuomotor mapping estimate that will be used for the next reach.

If the system is calibrated and the error is due to random noise (zero mean), a change of the mapping estimate would introduce rather than eliminate miscalibration by pulling the system moment by moment away from the correct calibration state thereby increasing response variance. Thus, the mapping estimate should not be changed; the best strategy for minimizing error is to average over all previous observations:

$$\hat{X}_{t+1} = \frac{\sum_{i=0}^{n} \hat{E}_{t-i}}{n} \quad \mathrm{n} \to \infty, \tag{2}$$

where n is the number of observations in the average. Of course, averaging over many observations would make the system slow to respond to abrupt changes in systematic error. In sum, the optimal strategies for dealing with systematic and random errors are quite different.

Everyday changes in sensorimotor mapping are probably best characterized by an accumulation of small changes rather than by large step changes (Baddeley, Ingram, & Miall, 2003; Kagerer, Contreras-Vidal, & Stelmach, 1997), so a given reach error is generally due to both systematic and random error (where random error includes error in reach execution). Thus, neither of the aforementioned strategies is optimal for the world we live in. This creates a fundamental problem: To determine the appropriate response to a given error, the nervous system should respond differently depending on how much of a reach error was due to systematic error and how much to random error. The system cannot determine those relative proportions from any single observation, so a calibration method is required that balances the competing needs to average out random error while remaining sensitive to changes in systematic error.

We can construct an algorithm that uses the available information optimally to solve this problem. By studying the properties of this optimal adaptor, we can determine how the system should react to different kinds of variation in the observed error. We can then compare the behavior of the optimal adaptor to human behavior. If the behaviors are similar, this is evidence that humans use the available information efficiently; if the behaviors are dissimilar, this is evidence that humans use the available information suboptimally or that the assumptions underlying the 'optimal' adaptor are inappropriate. Such comparisons between optimal and observed behavior have proven useful in many other domains of neuroscience (Crowell & Banks, 1996; Geisler, 1989; Najemnik & Geisler, 2005; Watson, 1987). The optimal adaptor specifies the adaptation rate that will minimize reaching error given the statistics of reaching errors and the precision with which they can be measured. Changes in those properties therefore lead directly to behavioral predictions for how the adaptation rate should change. We can present the same conditions to humans and thereby determine how well they take the available information into account when adapting.

In the three scenarios described above—error due to miscalibration, to random error, or to both—the properties of the environment determine the optimal adaptation rate. Nearly all previous models of motor (Bhushan & Shadmehr, 1999; Donchin, Francis, & Shadmehr, 2003; Shadmehr & Mussa-Ivaldi, 1994) and visuomotor recalibration (Cheng & Sabes, 2006; Diedrichsen, Hashambhoy, Rane, & Shadmehr, 2005) do not take environmental properties as input but rather attempt to identify the control parameters that allow a good fit to the data.

An optimal solution is provided by the Kalman filter, an algorithm widely used in engineering for the purpose of estimating dynamic systems (Baddeley et al., 2003; Ghahramani, Wolpert, & Jordan, 1997; Korenberg & Ghahramani, 2002; Maybeck, 1979). A Kalman filter is optimal if the system it is trying to estimate is linear, has stationary statistics, has Gaussian-distributed noise, and its errors are penalized by a least squares cost function; these assumptions have been made previously and are not unreasonable over short time periods. By combining noisy measurements over time with prior information, the Kalman filter maximizes the accuracy and precision of its estimates. The position of a satellite orbiting the earth, for example, can be estimated most accurately if position measurements, made uncertain by atmospheric distortions and sensor noise, are combined with the information that satellites travel in mostly elliptical paths. The optimal combination of measurement and prior information can be similarly applied to visuomotor behavior (Körding & Wolpert, 2004) and calibration (Ghahramani et al., 1997; Korenberg & Ghahramani, 2002). As the visuomotor mapping changes in time, the visuomotor system can maximize the accuracy of its mapping estimates by combining measurements of reaching error, made uncertain by visual blur, with knowledge based on prior experience of how stable the mapping generally is.

For a visuomotor mapping, X_t, changing in time, the Kalman filter continuously updates the estimate of the mapping by a weighted combination of the most recent observed error and the current mapping estimate based on previous estimates:

$$\hat{X}_{t+1} = \hat{X}_t + K(\hat{E}_t), \tag{3}$$

where K is the proportion of the observed error by which the mapping estimate will be adjusted (the Kalman gain). By using multiple observations, weighted by recency, to estimate the mapping error, the impact of random error is reduced. In this way, the system retains the ability to respond to changes in systematic error while not overreacting to random error.

Figure 1 shows the Kalman filter's response to step changes that are equivalent to introducing a prism and later removing it. As the filter adjusts the visuomotor mapping, the error between target and reach position decreases exponentially with time. For a given gain K:

$$\lambda = -\log(1-K), \tag{4}$$

where λ is the exponential time constant: Faster adaptation is associated with higher gain. There is no analytical solution for K (see Appendix A for the exact equation at steady state), but to close approximation,

$$K \cong \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{\hat{z}}^2}, \tag{5}$$

where $\sigma_{\hat{z}}$ is the uncertainty associated with measuring the location of the feedback (measurement uncertainty) and σ_x is the uncertainty of the visuomotor mapping that would produce unbiased performance (mapping uncertainty) (Korenberg & Ghahramani, 2002).

The values of the parameters $\sigma_{\hat{z}}$ and σ_x should be consistent with the distributions of measurements made by the observer before the step change was introduced. If an observed error is likely to have been caused by an erroneous measurement, $\sigma_{\hat{z}}$ should be large, and the visuomotor system should adapt slowly (Equations 4 and 5). If the visuomotor motor mapping is changing quickly, σ_x should be large, and the system should adapt quickly to minimize what is likely to be systematic error. Figure 1 shows the effects of those parameters on the adaptation rate of a Kalman filter. Higher values of $\sigma_{\hat{z}}$ yield slower adaptation and higher values of σ_x yield faster adaptation. The Kalman filter is provided the statistics, but humans must learn the values of $\sigma_{\hat{z}}$ and σ_x from the distributions of measurements (Baddeley et al., 2003). We presented many trials before the step change so that human subjects could learn the statistics and set their internal parameters appropriately before having to respond to the step change.

We compared the adaptation rates of humans and the Kalman filter for the same conditions. Specifically, we examined the effects of varying measurement uncertainty (by blurring the feedback isotropically and anisotropically) and of varying the statistical properties of the mapping. By comparing human and Kalman filter behavior, we ask whether humans take into account the properties of the feedback and the mapping in a manner consistent with the filter to minimize reaching error in changing environments.

Experiment 1: Measurement uncertainty

We first investigated whether changing measurement uncertainty ($\sigma_{\hat{z}}$) changes adaptation rate in humans, as it does in a Kalman filter: specifically, whether increasing $\sigma_{\hat{z}}$ causes a slowing of adaptation. To do this, we used the set-up in Figure 2 to examine adaptation rate when the

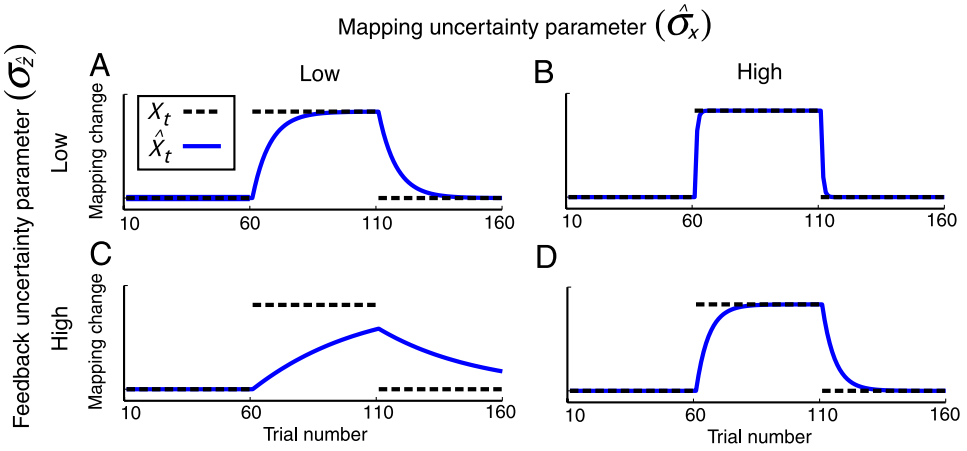

Figure 1. Kalman filter responses to step changes. The dashed black lines in each panel represent the mapping between the position of the reach endpoint and the position of the visual feedback. This relationship is the visuomotor mapping. As in our experiments, there are three phases: pre-step (trials 1–60), step (61–110), and post-step (111–160). A step change in the mapping occurs during the step phase; the initial mapping is restored after the step phase. The blue curves represent the visuomotor mapping estimates ($\hat{X}_t$) over time. The upper and lower rows show estimates when the measurement uncertainty ($\sigma_{\hat{z}}$) is small and large, respectively. An increase in $\sigma_{\hat{z}}$ causes a decrease in adaptation rate. The left and right columns show responses when the mapping uncertainty (σ_x) is small and large, respectively. An increase in σ_x causes an increase in adaptation rate; the effect is larger when $\sigma_{\hat{z}}$ is large.

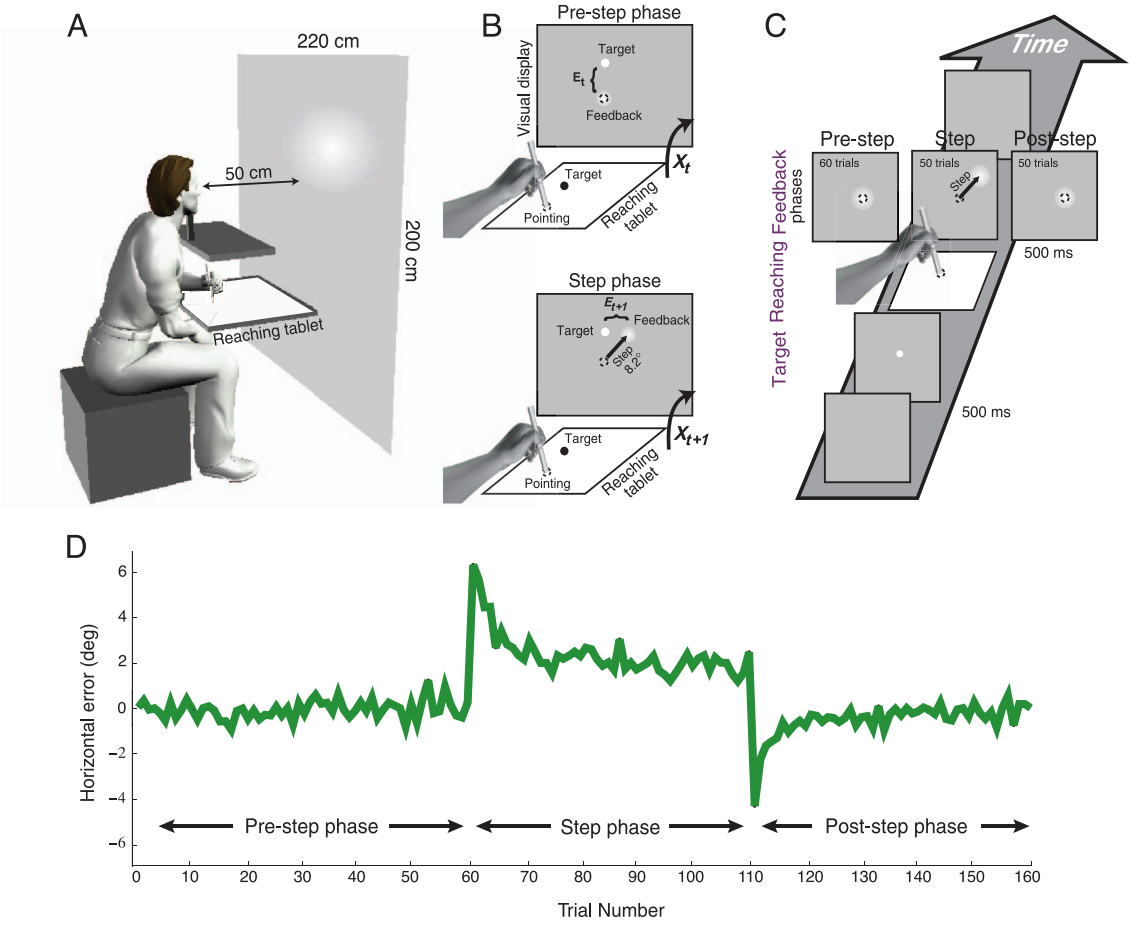

Figure 2. Experimental setup and procedure. (A) Experimental setup. Subjects sat in front of a visual display that they viewed binocularly from a distance of 50 cm. Their heads were restrained by a chin-and-forehead rest. Subjects held a stylus with their preferred hand on a horizontal graphics tablet. The hand was not visible. The *x*–*y* positions of the stylus were recorded on the tablet. (B) The visual target, actual reach endpoint, and feedback stimulus. Visual feedback was available only at the end of each reach. Reach endpoint was defined as first place the stylus touched the graphics tablet after the reach was initiated. The upper panel shows the target (bright circle), reach endpoint (dashed circle), and visual feedback (bright Gaussian blob superimposed on dashed circle) for trials in which the feedback and endpoint had the pre-step mapping X_t. The lower panel shows the mapping X_{t+1} after the mapping change of the step phase: the feedback was displaced by 8.2° up and to the right (5.8° horizontally and 5.8° vertically) relative to the pre-step mapping, as indicated by the offset between the dashed circle and Gaussian blob. A reach error E_t was defined as the difference between the target and the feedback locations. (C) Time course for a trial. The target was presented at a random position for 500 ms. Then, the subject made a rapid reach in response to the target's position. The position of the reach endpoint was recorded. Then a 500-ms feedback stimulus appeared whose position was determined by the current mapping X_t. (D) Average errors (n = 24) across all three phases (pre-step, step, and post-step) for a single condition. Post-step error reduction was generally faster than step-phase error reduction, but the same basic trends were followed. Because it is advantageous to work with slower rates, we focus our analysis on the step rather than on the post-step data (see Methods for details).

localizability of the visual feedback stimulus was varied. The feedback stimuli were blurry isotropic Gaussian blobs with different standard deviations (Figure 3A). Using blur to reduce measurement uncertainty has two advantages for analyzing the system experimentally. First, reaching adaptation often occurs quite quickly in humans, sometimes in only a few trials (von Helmholtz, 1867). Fast rates are difficult to quantify from noisy data, so it is advantageous to work with slower adaptation rates. Second, unlike other forms of measurement uncertainty (see Experiment 3), we can use psychophysical techniques to determine exactly how uncertain a given amount of blur makes the measurement of feedback position.

We first conducted a visual discrimination experiment to determine the relationship between the amount of blur (quantified as the standard deviation of the blob stimulus, σ_{blur}) and the ability to localize it visually. The just-noticeable difference (JND) increased monotonically with σ_{blur} (Figure 3B). Thus, blurring had the desired effect of increasing the uncertainty of the measurement of feedback position. Henceforth, our estimate of feedback uncertainty σ_z will equal the measured JND/$\sqrt{2}$.

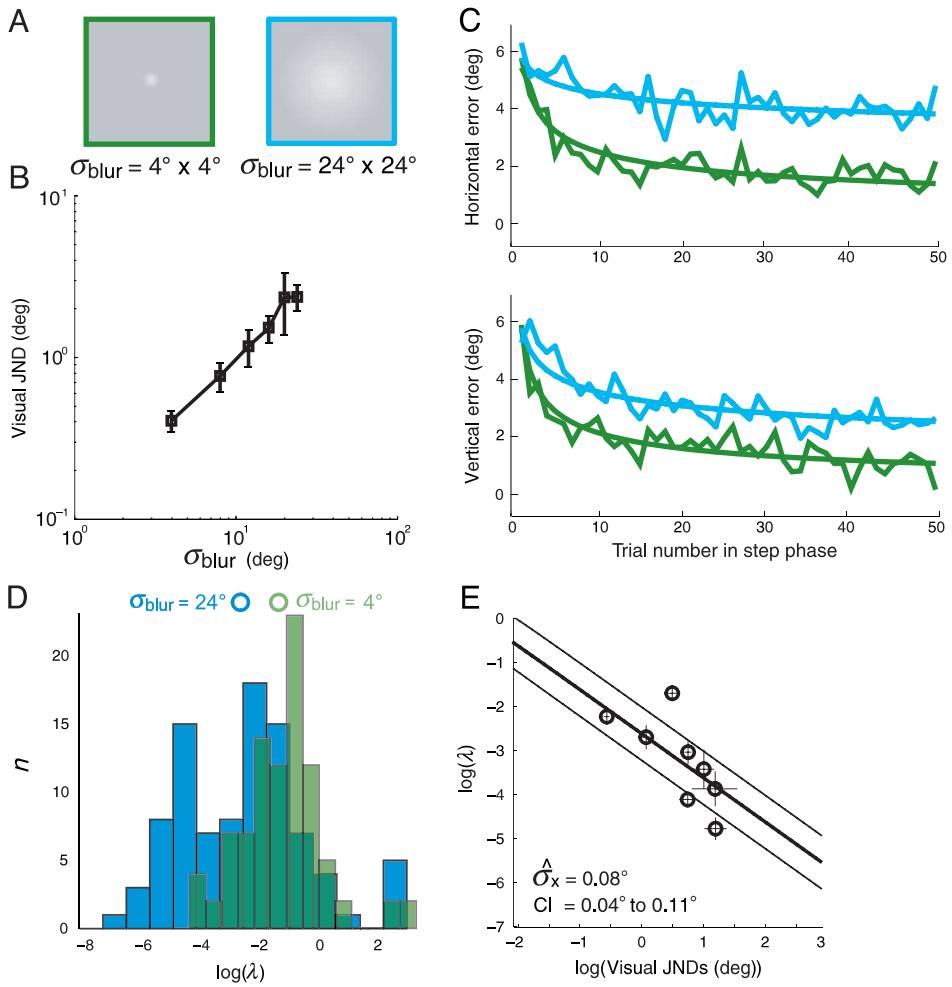

Figure 3. Stimuli and results from Experiment 1. (A) Isotropic visual feedback stimuli. They were isotropic Gaussian blobs with σ_{blur} = 4° × 4° and 24° × 24°. (B) Just-noticeable differences (JNDs) in the visual localization experiment as a function of σ_{blur}. Visual discrimination thresholds increased monotonically from 0.4° to 2.4° as σ_{blur} was increased from 4° to 24°. $y = 0.16*x - 0.085$, $R^2 = 0.9821$. (C) Average adaptation profiles for isotropic feedback stimulus. On the first trial, subjects had no information about the size and the direction of the step, so initial errors were roughly equal to the step. Subjects gradually adjusted their visuomotor mapping estimate, $\hat{X}_t$, so that the observed errors approached zero. Light blue represents the data when σ_{blur} = 24° × 24° and green represents the data when σ_{blur} = 4° × 4°. The upper panel plots horizontal error (in degrees) against trial number and the lower panel vertical error (also in degrees) against trial number across trials. One degree corresponds to ~2.5 mm on the tablet, so the shift of 8.2 deg corresponds to ~2 cm on the tablet. The line segments represent the data, and the smooth curves are the averages of the best-fitting exponentials (Equation 6). Exponentials and power laws, each with two free parameters, were fit to all individual subject data (exponential: $E_t = (b + C)e^{-\lambda(t-1)} + b$; power: $E_t = (b + C)t^{-\lambda} + b$). Exponentials provided a better fit 78% of the time ($p < 1.7*10^{-27}$; sign test, n = 384 (24 subjects × 8 conditions [4° × 4°, 8° × 8°, 12° × 12°, 16° × 16°, 20° × 20°, 24° × 24°, 24° × 4°, and 4° × 24°] × 2 dimensions [x, y])). (D) Histogram of all vertical and horizontal adaptation rates for subjects in which σ_{blur} was 4° (green) or 24° (light blue) in the respective direction. The histogram includes data from conditions with 4° × 4°, 4° × 24°, 24° × 4°, and 24° × 24° feedback. Rates were significantly slower in the large blur condition (t test, $p < .0001$; n = 96). (E) Average time constants as a function of average visual JND for each condition of the experiment. Exponential fits were performed along the diagonal (C = 8.2°). This includes the two conditions with anisotropic blobs plotted in Figure 4 and four additional conditions with isotropic blobs (8° × 8°, 12° × 12°, 16° × 16°, and 20° × 20°). Stimuli and data from these conditions are shown in Supplementary Figure 1. Error bars indicate one standard deviation. The thick black line (actually a curve) shows the predicted change of adaptation rate with feedback uncertainty for $\hat{\sigma}_x$ = 0.08°. The thin black lines are the 95% confidence intervals (R^2 = 0.61). As $\hat{\sigma}_x$ changes, the y-intercept of the Kalman filter prediction changes, but the slope is effectively unchanged.

We then conducted a reaching adaptation experiment consisting of three phases (Figure 2C). During the pre-step phase, subjects learned the initial mapping, X_t, between the reach endpoint and the visual feedback (Figure 2B) and gained experience with the environmental conditions so their internal parameters could be properly set. In the

step phase, we introduced a step change in the mapping: Visual feedback was shifted 8.2° up and to the right. The step persisted until the post-step phase when the original mapping was restored. Data from a full run of 160 trials are shown in Figure 2D. Adaptation was assessed by examining reaching errors during the step phase (see Methods section for why we did not focus on post-step data). For our experiments, reaching error E_t is the difference between the locations of the visual target and visual feedback on trial t. To quantify adaptation rate, we fit an exponential function separately for the x and y directions to each subject's reaching errors across trials:

$$E_t = (b + C)e^{-\lambda(t-1)} + b, \quad (6)$$

where λ is a time constant, b is a bias term, C is a constant equal to the shift, and t is the trial number. We defined adaptation rate as the time constant (λ) of the best-fitting exponential.

Figure 3C shows average reaching errors across trials for the small and the large feedback blobs. Reaching responses were averaged because trial-by-trial errors for individual subjects were quite variable. The smooth curves are the average of all subjects' best-fitting exponentials. Note that the curves look more like power laws than exponentials. There are two possibilities for this: (a) the individual profiles could be best characterized by power laws, a possibility that would fatally undermine the Kalman filter model, or (b) the individual profiles could be best characterized by exponentials since the average of multiple exponentials with different time constants approximates a power law (Heathcote, Brown, & Mewhort, 2000; Wickens, 1999). To determine the decay function most appropriate to the individual subject data, we fit all individual error profiles from Experiments 1 and 3 with both exponential and power functions with the same number of free parameters (see caption for details). We found that exponentials provided better fits 78% of the time, a highly significant result ($p < 1.7*10^{-27}$; sign test). This justifies our use of exponentials to analyze individual subject data and indirectly supports the Kalman filter model.

Although the Kalman filter predicts asymptotes at zero, many individual adaptation profiles appeared to asymptote at values greater than zero. The non-zero asymptote, b, presumably manifests a bias to either mislocalize the feedback or reach toward the center of the display screen and was not incorporated in our Kalman filter model. Because we are interested primarily in learning rate and since removing constant bias from the data has been shown not to affect estimates of rate (Cheng & Sabes, 2006), we focus on the time constants required to fit the data and not on the bias.

Figure 3D shows the distribution of the best-fitting exponential time constants across subjects for the small and large blur conditions. Figure 3E summarizes the adaptation data by plotting the average time constants as a function of feedback uncertainty ($\sigma_{\hat{z}}$). Like a Kalman filter, subjects adapted more slowly when the position of the feedback was less certain.

We can examine the relationship between human and Kalman behavior in another way. Once a Kalman filter has settled to steady state, K is fully determined by the visuomotor-mapping uncertainty, σ_x, and measurement uncertainty, $\sigma_{\hat{z}}$ (Appendix A). We set $\sigma_{\hat{z}}$ to the value of the visual JND/$\sqrt{2}$ (tantamount to making the reasonable assumption that the JND/$\sqrt{2}$ manifests the measurement uncertainty). We then assumed different values for σ_x and determined how the filter's adaptation rate changes with $\sigma_{\hat{z}}$. For each assumed value of σ_x, there is a line relating the visual JND to adaptation rate. All the lines have approximately the same slope. The line in Figure 3E shows the behavior of the Kalman filter that fits our data best (R^2 = 0.61). The change in adaptation rate observed in humans is similar to the pattern of rate changes in the optimal adaptor. The value of $\hat{\sigma}_x$ that provides the closest agreement is 0.08° (95% CI = 0.04°–0.11°). Clearly, humans adapted at rates consistent with the Kalman filter as the feedback was made less certain.

We also wondered whether adaptation rates can be spatially anisotropic (Cheng & Sabes, 2006). If measurement uncertainty is greater in one direction than in another, a Kalman filter would adapt more slowly in the direction of greater uncertainty. We investigated this by presenting feedback blobs that were elongated horizontally or vertically (Figure 4A). After the introduction of a step change in the mapping, reaching errors were indeed reduced more slowly in the direction of greater blur (Figure 4B). Adaptation rate in a given direction was influenced only by the measurement uncertainty in that direction; as predicted by the Kalman filter, anisotropic rates do indeed occur.

Adaptation rates that differ by direction yield interesting spatial patterns of reaching over time. Figure 4C shows the Kalman filter's pattern of horizontal and vertical reaching errors in the adaptation phase. With isotropic feedback uncertainty, the filter's horizontal and vertical rates are similar, and the spatial distribution of errors over time is linear. With anisotropic feedback, however, the filter's horizontal and vertical rates differ, so the distributions are curved (Figure 4C). Replotting the data in Figure 4B reveals that these effects occurred in humans as well (Figure 4D). With symmetrically blurred feedback, subjects reached progressively more downward and leftward along a straight path toward the target. With anisotropic blur, subjects' reaches progressed along curved paths. Thus, in both humans and in the Kalman filter, adaptation rate slows when uncertainty increases, and the slowing in a given direction is influenced only by the uncertainty in that direction.

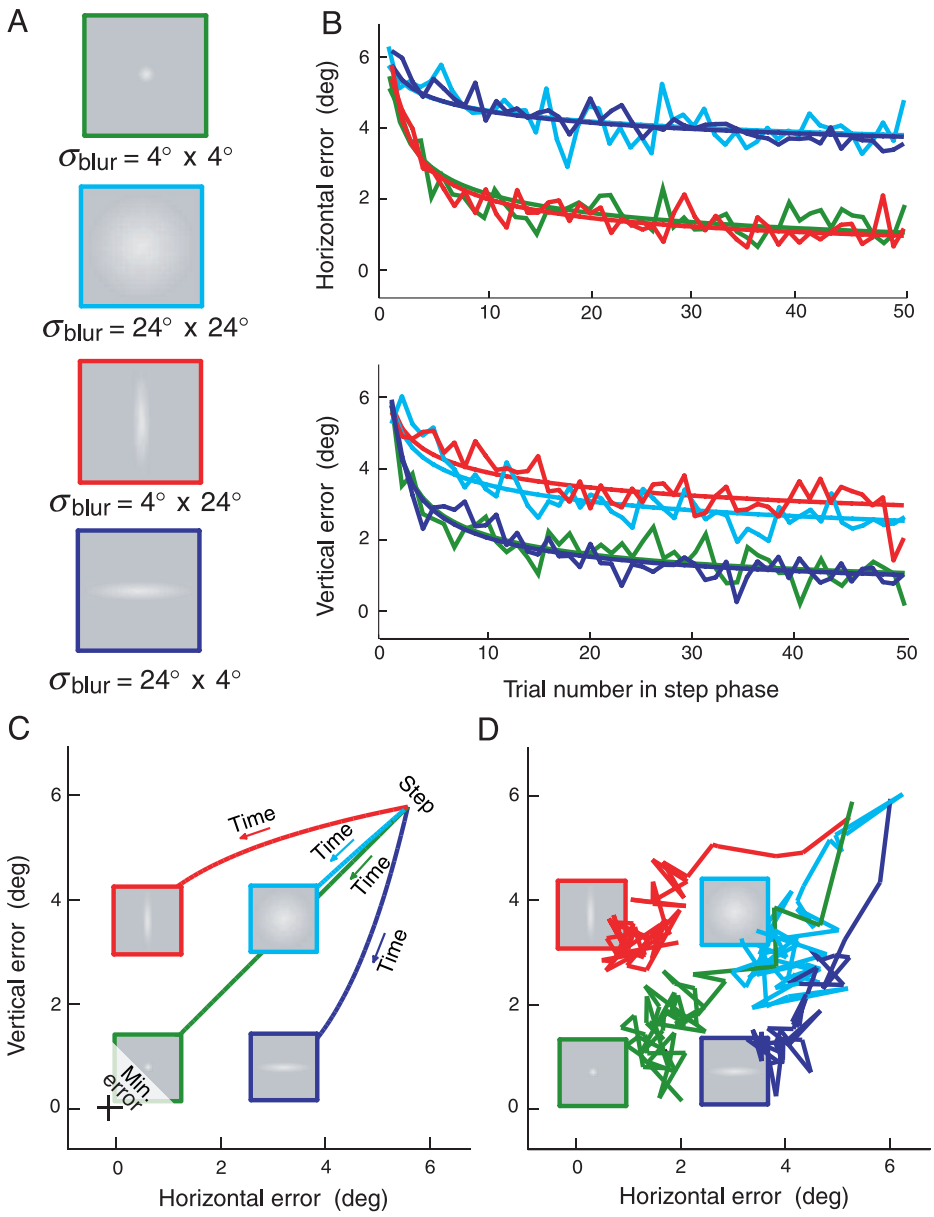

Figure 4. Stimuli and results from Experiment 1 comparing anisotropic and isotropic visual feedback conditions. (A) Gaussian blobs with σ_{blur} = 4° × 4°, 24° × 24°, 4° × 24°, and 24° ×4° used as visual feedback. (B) Average adaptation profiles for isotropic and anisotropic feedback stimuli. Red and dark blue represent the data when the feedback stimulus was anisotropic: σ_{blur} = 4° × 24° and 24° × 4°, respectively. As in Figure 3, green and light blue represent the data when the feedback was isotropic: σ_{blur} = 4° × 4° and 24° × 24°, respectively. The upper panel plots horizontal error across trials, and the lower panel plots vertical error across trials. The line segments represent the data, and the smooth curves are the average of the best-fitting exponentials. A repeated measures ANOVA on the best-fitting exponential rates with blur as a factor revealed that the low blur conditions were significantly faster than the high blur conditions ($F(1,23) = 54.4$, $p < .0001$). Multiple comparison tests showed that the conditions with low blur in each direction were significantly faster than the high-blur conditions, and that the low-blur conditions did not differ significantly from each other. (C) Spatial profiles for a Kalman filter in response to a step change in the mapping. Horizontal error is plotted as a function of vertical error. Initially, the reaching error corresponds to the step change (5.8° horizontally and vertically), so the error is large (upper right). As the filter adjusts its responses over time toward the goal of zero error (lower right), the horizontal and the vertical errors change. When the blur is isotropic (green and light blue), the horizontal and the vertical adaptation rates are the same, so errors progress along a diagonal toward zero error. When the blur is anisotropic (red and dark blue), horizontal and vertical adaptation rates differ, so errors progress along curves, the direction of the curve depending on the direction of least blur. The Kalman filter's parameter for measurement uncertainty $\sigma_{\hat{z}}$ was set equal to the human visual JNDs for the corresponding conditions. The filter's parameter for mapping uncertainty $\hat{\sigma}_x$ was set equal to 0.08°, the estimated value of baseline visuomotor mapping uncertainty (see Figure 3E). (D) Spatial profiles for the average human data plotted in the same format.

Experiment 2: Mapping uncertainty

Increasing the uncertainty of the visuomotor mapping, σ_x, produces faster adaptation in the Kalman filter (Equations 4 and 5). As we said earlier, this makes sense. If the mapping changes frequently, the system should adjust its mapping estimate quickly in response to an error because the previous mapping estimate is less likely to be correct. If the mapping has been stable, adjustments should be slow because it is likely that the mapping from previous trials is correct and that the error was due to measurement error.

To examine the effect of mapping uncertainty, we imposed a random walk on the relationship between the reach endpoint and the visual feedback position in addition to the step presented in Experiment 1. A random walk is a good mathematical description of many natural processes such as the motion of a molecule in a gas: its direction changes frequently and unpredictably. The walk value was:

$$W_{t+1} = W_t + N(0, \sigma_{\text{walk}}), \tag{7}$$

where $N(0, \sigma_{\text{walk}})$ is a Gaussian random variable with mean 0 and standard deviation σ_{walk}. In the pre- and post-step phases, $(x', y') = (x + W_x, y + W_y)$, where (x', y') were the coordinates of the reach endpoint on the graphics tablet, $(x + W_x, y + W_y)$ were the screen coordinates of the feedback, and W_x and W_y were the horizontal and vertical

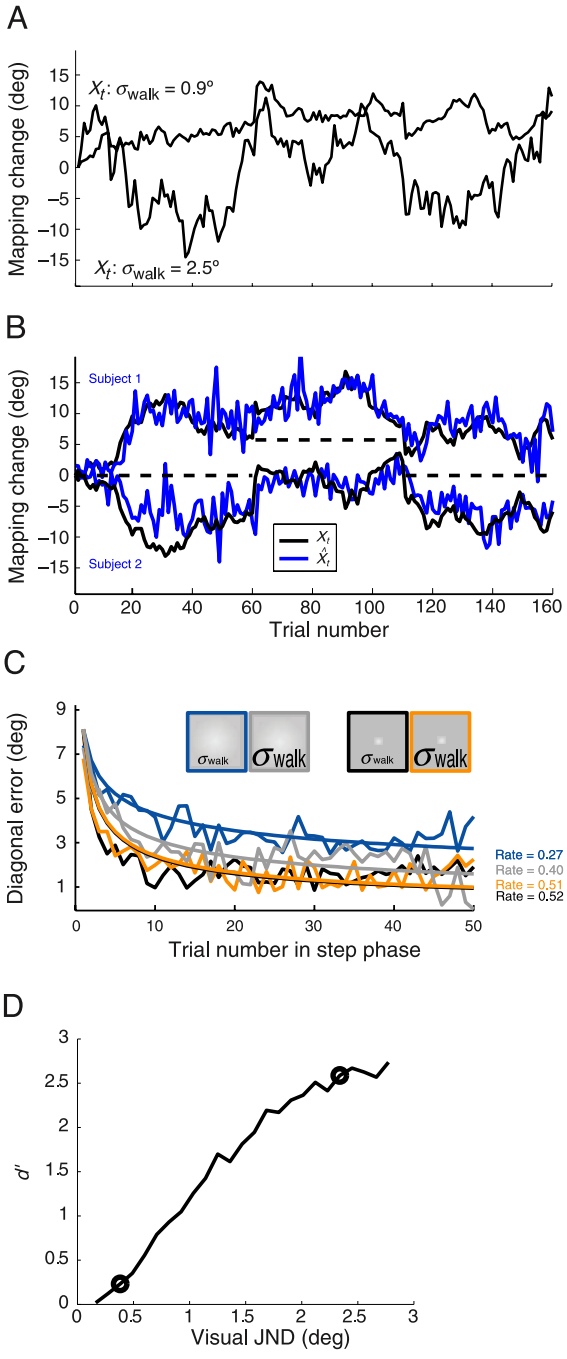

Figure 5. Stimuli and results from Experiment 2. (A) Example random walks (plus step change) with σ_{walk} = 0.9° and σ_{walk} = 2.5°. (B) Reach endpoints multiplied by −1 from two representative subjects (blue traces) in response to two random walks with σ_{walk} = 0.9° applied to the mapping between reach endpoint and visual feedback positions (solid black traces). Subject mapping estimates clearly followed the walks. The random walks that were added to the step changes were mirror images of one another. The average of the walks was identical to the trial-by-trial stimulus values in Experiment 1 (dashed black trace). Variation due to the walks was nulled by averaging across all walks in Experiment 2. (C) Average adaptation profiles for the random walks during the step phase. Errors are plotted against trial number. Because walk statistics were isotropic within a condition, error was calculated in the direction of the constant shift (rather than in x and y separately). Dark blue: σ_{walk} = 0.9° × 0.9°; σ_{blur} = 24° × 24°. Gray: σ_{walk} = 2.5° × 2.5°, σ_{blur} = 24° × 24°. Black: σ_{walk} = 0.9° × 0.9°, σ_{blur} = 4° × 4°. Orange: σ_{walk} = 2.5° × 2.5°, σ_{blur} = 4° × 4°. Walk standard deviation was chosen to be equivalent to the visual JNDs for σ_{blur} = 4° × 4° and 24° × 24°. The curves are the best-fitting power laws to the data averaged across subjects. Exponentials could not be fit to the individual data because of the drifts associated with the random walks. The rates displayed on the right are the exponents of the best-fitting power laws. An increase in walk variability yields an increase in adaptation rate with large blur. (D) d' for discriminating adaptation rates as a function of visual JND. Circles indicate the two blur conditions in our experiment. Assuming that the system estimate of mapping uncertainty, $\hat{\sigma}_x$, is roughly equal to σ_{walk}, rate predictions were made for the two levels of σ_{walk} imposed in Experiment 2 at multiple blur levels (Equations 4 and 5; Appendix A). Circles mark the blurs tested in Experiment 2. 1000 exponential adaptation profiles were generated with the predicted λ. After corrupting each profile with motor noise (estimated from the pre-step phase of Experiment 1 (σ_{motor} = ~2°), we fit each individually with an exponential. d' was calculated from the resultant distributions of best-fit λs. d' was 0.2 and 2.6 in the low- and high-blur conditions, respectively. Thus, high rates are difficult to measure reliably because subjects adapt quickly over the first few trials. We therefore do not expect a measurable effect of σ_{walk} when blur is low.

components of W_t. After the pre-step trials, step-phase trials were presented with a constant shift between the reach endpoint and the feedback: $(x', y') = (x + W_x + 5.8°, y + W_y + 5.8°)$. We manipulated mapping variability by changing σ_{walk}. Figure 5A shows examples of how the mapping between reach endpoint and feedback changed for two random walks with different values of σ_{walk}. Figure 5B shows trial-by-trial mapping estimates from two representative subjects in response to two random walks with the same variability. Clearly, subjects adjusted their mapping estimates to follow the walk.

We expect that faster adaptation rates will be associated with walks of greater variability. However, our ability to measure this effect should be affected by measurement uncertainty. As in Experiment 1, we expect that conditions with large blur will on average yield slower adaptation than those with small blur. When adaptation is slow, it is relatively easy to quantify the rate because error decreases over many trials. When adaptation is fast, however, it is quite difficult to quantify the rate because the error is minimized in only a few trials. Assuming that measurement uncertainty is equal to the visual JND/$\sqrt{2}$ (Experiment 1) and that the system estimate of mapping uncertainty $\hat{\sigma}_x$ is equal to σ_{walk}, the initial error would be reduced by ~97% in two trials when the blur is small. Thus, we expect to find a large effect of walk variability with large blur and little if any effect when blur is small.

Figure 5C shows reach errors averaged across subjects. As expected, with large blur, adaptation was faster with larger σ_{walk}, and with small blur, there was no discernible effect. There was also a clear effect of the localizability of the feedback: Adaptation was slower with large blur. Because of the previously mentioned difficulty of measuring changes in adaptation rate at high rates, we could not determine whether human rates slowed with increasing σ_{walk} in the small-blur conditions of this experiment. But for the conditions in which we can observe rate changes reliably, the pattern of human and filter rate changes were similar. This supports the idea that humans adjust adaptation rate with mapping variability changes in a manner consistent with a Kalman filter.

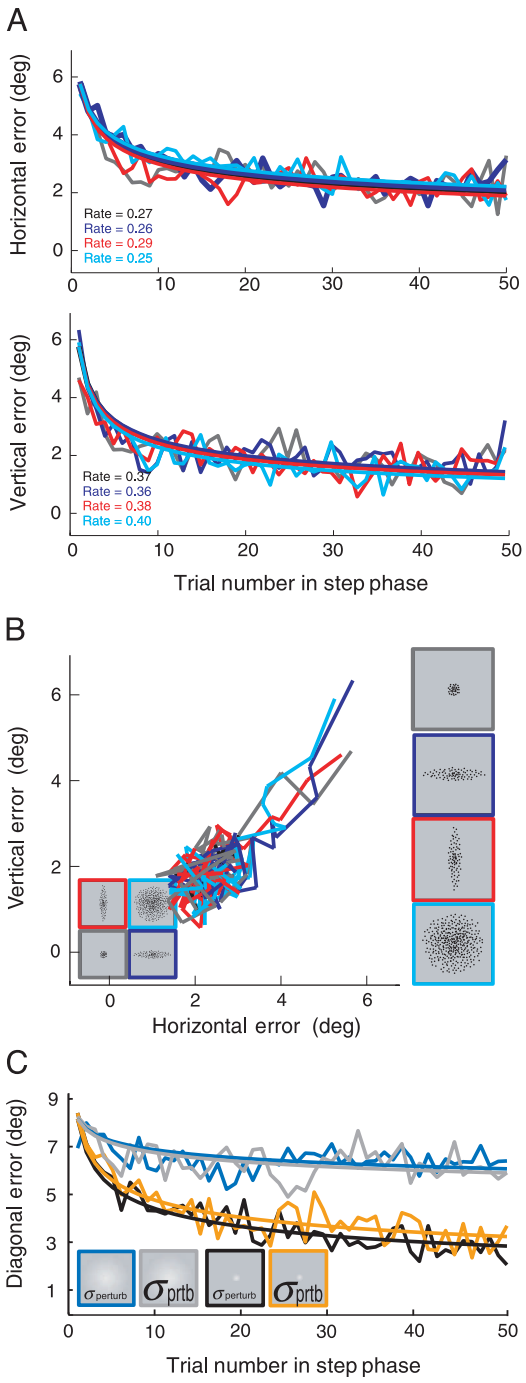

Figure 6. Results from Experiment 3. (A) Results from Experiment 3a. The insets in panel b represent $\sigma_{perturb}$ for the various conditions; σ_{blur} was fixed at 4° × 4° for all conditions. Black: $\sigma_{perturb}$ = 0.9° × 0.9°. Light blue: $\sigma_{perturb}$ = 2.5° × 2.5°. Dark blue and red: $\sigma_{perturb}$ = 2.5° × 0.9° and 0.9° × 2.5°, respectively. Note that these standard deviations are equivalent to the visual JNDs for the corresponding conditions in Experiment 1. The upper data plot shows horizontal error over trials and the lower plot shows vertical error over trials. The line segments represent the data, and the smooth curves are the averages of the best-fitting exponentials. A repeated measures ANOVA with perturbation as a factor showed no effect of $\sigma_{perturb}$ on adaptation rate ($F(1,23)$ = 0.25, $p < 0.62$). (B) Average spatial profiles of reach errors over trials. Vertical error is plotted as a function of horizontal error as in Figure 4D. The profiles were linear for all forms of random perturbation. Note the different pattern of results in Experiment 1. (C) Stimuli and results from Experiment 3b. The insets represent $\sigma_{perturb}$ for the various conditions; σ_{blur} was varied. Light blue: $\sigma_{perturb}$ = 0.9° × 0.9°; σ_{blur} = 24° × 24°. Dark blue: $\sigma_{perturb}$ = 2.5° × 2.5°; σ_{blur} = 24° × 24°. Black: $\sigma_{perturb}$ = 0.9° × 0.9°; σ_{blur} = 4° × 4°. Red: $\sigma_{perturb}$ = 2.5° × 2.5°; σ_{blur} = 4° × 4°. An ANOVA showed there was again no effect of $\sigma_{perturb}$ ($F(1,23)$ = 0.19, $p < 0.67$). Blur continued to produce a main effect ($F(1,23)$ = 14.52, $p < 0.0009$).

Experiment 3: Random trial-by-trial variability

In the first experiment, we showed that blurring the stimulus (increasing σ_{blur}) made the feedback difficult to localize and resulted in slower adaptation, an effect probably caused by an increase in the system estimate of measurement uncertainty ($\hat{\sigma}_{\hat{z}}$). In the second experiment, we showed that adding a random walk to the visuomotor mapping (increasing σ_{walk}) resulted in faster adaptation; an effect probably caused by an increase in the system estimate of mapping variability ($\hat{\sigma}_x$). There is a third source of variability in rapid reaching ($\sigma_{perturb}$): random trial-by-trial variation in the reach endpoints due, for example, to noise in reach execution (motor noise). We next asked how the system responds to such variations.

Random perturbation, like blur, reduces the certainty with which the system can use a given error to estimate the difference between its current mapping estimate $\hat{X}_t$ and the mapping necessary to produce unbiased performance X_t (Baddeley et al., 2003). Random perturbation should therefore affect the measurement uncertainty parameter $\sigma_{\hat{z}}$. The adaptation rate of an optimal adaptor, assuming stationary statistics (i.e., constant measurement and mapping uncertainties), should therefore slow down. This is similar to the effect of blurring the feedback (Experiment 1), but unlike blur, the amount of random perturbation cannot be estimated from a single observation, so $\sigma_{perturb}$ may have more or less effect depending on how the system estimates measurement uncertainty.

Random perturbation could conceivably be treated as mapping uncertainty, thereby affecting the system's estimate of σ_x. This would be non-optimal, however, because with random perturbation the mean of the mapping never changes, so the variance of the change in the mapping producing unbiased performance is unaffected by $\sigma_{perturb}$ (Appendix B). In fact, the variance of the change in the mapping yielding unbiased performance is 0. Thus, one expects that random perturbation should not cause an increase the system's estimate of mapping uncertainty.

To investigate how the system responds to perturbations that are uncorrelated from trial to trial, we conducted two experiments. In the pre- and post-step phases, a reach endpoint at (x', y') generated feedback at $(x + P_x, y + P_y)$, where $P_t = N(0, \sigma_{perturb})$ and N is Gaussian with mean 0 and standard deviation, $\sigma_{perturb}$. As before, the pre-step phase allowed subjects to learn the properties of the mapping and feedback. During the step phase, reach position (x', y') generated feedback at $(x + P_x + 5.8°, y + P_y + 5.8°)$. In Experiment 3a, σ_{blur} was fixed and small, and the values of $\sigma_{perturb}$ were essentially equivalent to the visual JNDs for isotropic and anisotropic blobs in Experiment 1. Interestingly, $\sigma_{perturb}$ had no measurable effect on adaptation rate (Figure 6A). The spatial error profiles across trials were therefore roughly linear whether $\sigma_{perturb}$ was anisotropic or not (Figure 6B). Experiment 3b was similar to Experiment 2 except that the mapping was changed via random perturbation instead of a random walk. The blur of the feedback was also varied. There was again no measurable effect of $\sigma_{perturb}$. Adaptation rates were determined only by feedback localizability: They were fast with small blur and slow with large blur (Figure 6C). The same pattern of results was observed in Experiment 1 when $\sigma_{perturb}$ was 0°.

The adaptation rate of the Kalman filter slows when random perturbation is added to the feedback. We found in contrast to this that increasing random perturbation has no measurable effect on adaptation rate in humans. The lack of an effect of adding random perturbation has been observed previously (Diedrichsen et al., 2005; Donchin et al., 2003; Smith & Shadmehr, 2004; but see Baddeley et al., 2003, and Appendix C). In this case, human behavior is not consistent with that of an optimal adaptor at least for the assumptions we used in constructing the algorithm. We will examine those assumptions in the Discussion section.

Discussion

We developed an optimal adaptor and studied how it responds to changes in the statistics of the input. By doing so, we could determine the degree to which humans use the available sensory information to minimize reaching errors over time and thereby maintain visuomotor calibration. In our first experiment, we found that adaptation rate is reduced in humans and the optimal adaptor in the same direction-specific manner when feedback uncertainty is increased by spatially blurring the stimulus. In the second experiment, we found that adaptation rate increases in both humans and optimal adaptor when the variability of the visuomotor mapping is increased by a random walk. In the third experiment, we observed different behaviors in humans and optimal adaptor: Injecting random perturbation in the feedback causes adaptation to slow down in the optimal adaptor but had no effect in humans. In the remainder of the discussion, we consider the implications of our results.

Measurement uncertainty

The fact that feedback localizability affects adaptation rate and does so in a manner consistent with the Kalman filter has important implications for visuomotor and intersensory calibration. Any stimulus property that affects the reliability of the system's measurements of feedback position should also affect adaptation rate. When the feedback is visual, such properties include blur (Kayargadde & Marten, 1996), luminance (Waugh & Levi, 1993a), contrast (Waugh & Levi, 1993b), duration (Waugh & Levi, 1993b), retinal eccentricity (Levi & Klein, 1985), and whether the

feedback is in the frontal plane or in depth (Gepshtein & Banks, 2003; van Beers et al., 2002). In inter-sensory calibration, measurement uncertainty should affect adaptation rate in the fashion we observed with open-loop reaching. Consider, for example, recalibrating auditory–visual location maps after conflicts are introduced via monaural earplugs or prism lenses (Hofman, Van Riswick, & Van Opstal, 1998; Knudsen & Knudsen, 1985; King et al., 2001; King, Parsons, & Moore, 2000; Zwiers, Van Opstal, & Paige, 2003). The nervous system should adjust the mapping at a rate determined in part by the uncertainties of the visual and the auditory signals: Adaptation should occur more rapidly when the positions of visual and auditory stimuli are well specified and more slowly when they are poorly specified (Ghahramani et al., 1997; Korenberg & Ghahramani, 2002). It would be interesting to see if this prediction is borne out.

Mapping uncertainty

We found that more variable random walks yielded faster adaptation (Figure 5C). It is sensible to speed up adaptation when current errors are predictive of future ones, as they are with positively correlated time series such as random walks. By similar reasoning, negative correlations should cause adaptation to slow down. Smith and Shadmehr (2004) observed these behaviors: Adaptation rate increased in response to positively correlated error signals (i.e., random walks) and decreased in response to negatively correlated signals.

These findings have implications for calibration between sensory signals. Two sensory estimators that are often miscalibrated should adapt more quickly than two that are tightly coupled. Consider, for example, two signals from two senses as opposed to two signals from within one sense. There are presumably more causes of miscalibration in the former case (e.g., vision and touch) than in the latter (e.g., two visual depth cues) (Hillis, Ernst, Banks, & Landy, 2002). By this reasoning, adaptation should be more sluggish within than between senses. This prediction is supported by experimental data. When two visual depth cues are re-calibrated relative to one another, adaptation is slow (Adams, Banks, & van Ee, 2001); when visual and proprioceptive location cues are recalibrated, adaptation is fast (Harris, 1963; Hay & Pick, 1966). The principle that frequent mapping changes should increase adaptation rate may generalize to other learning domains: The more frequently the system needs to learn, the faster it should learn.

Another reason for slow adaptation may be the complexity of the mapping changes, a factor not included in our instantiation of the Kalman filter. Numerous studies have shown that the visuomotor system adapts slowly or incompletely to non-uniform scalings (Bedford, 1989, 1993a, 1993b). For example, after receiving progressive lenses for correcting presbyopia, people often experience mild vertigo for 1 to 2 weeks (Brandt & Daroff, 1980). Progressive lenses introduce different magnifications for different elevations in the visual field. This change between vision and other sensorimotor systems is more complex than that associated with standard lenses. Although the relevant mapping change is inter-sensory, adaptation is slow probably because the change introduced by the progressive correction is outside the realm of experiences the system normally encounters.

Random perturbation

In an environment with stationary statistics, the adaptation rate of a Kalman filter decreases with increases in random perturbation. In contrast, we observed no change in rate as the amount of random perturbation was varied (Figure 6). As we said earlier, this result is consistent with earlier reports (Diedrichsen et al., 2005; Donchin et al., 2003; Smith & Shadmehr, 2004). There are at least three plausible explanations for non-optimality in the response to random perturbation. (1) The system estimate of feedback uncertainty ($\hat{\sigma}_{\hat{z}}$) may be derived only from information that is available in a single observation (e.g., stimulus blur) and not from information that must be obtained over many observations (random perturbation). However, in the random walk experiment that adaptation rate is adjusted based on experience with previous trials. Specifically, we found evidence that $\hat{\sigma}_x$ is adjusted based on experience with previous trials. It therefore seems unlikely that previous trials are not considered for the estimation of $\sigma_{\hat{z}}$. (2) Random perturbation may increase the system's estimates of both $\sigma_{\hat{z}}$ and σ_x such that their opposing effects cancel and no measurable change in adaptation rate occurs. This also seems unlikely because there is no obvious reason both estimates should be affected in the same way by random perturbation. (3) The speed at which changes in environmental statistics can be learned may differ for random perturbations and random walks. The rate for learning changes in the statistics should depend on how reliably the changes can be determined. Changes in uncorrelated noise (σ_{perturb}) that must be estimated from past measurements may be more difficult (i.e., slower) to learn than changes in correlated noise (σ_{walk}). If this is the case, the pre-step phase of $\hat{\sigma}_{\hat{z}}$ whereas the pre-step phase of Experiment 2 would have had enough trials to update $\hat{\sigma}_x$. Further research is needed to determine how the visuomotor system learns the statistics of the environment.

Our experimental results do not allow us to evaluate these explanations empirically, but the third one makes sense. When estimating measurement and mapping uncertainties from the environment, the nervous system is not given estimates of the correlated and the uncorrelated mapping variation directly; rather it must estimate them from the stream of incoming observations. The estimation could be done precisely with many observations if the

environmental statistics were stationary, as we assumed in the construction of the optimal adaptor. However, the statistics in the everyday environment might well not be stationary (i.e., measurement and mapping uncertainty change over time). Using many observations would make the system slow to respond to changes in these statistics. To retain responsiveness, the nervous system should estimate parameters over few observations, the exact number depending on how quickly the statistics are likely to change and how reliably the statistics can be determined. With few observations, uncorrelated mapping variation cannot be fully distinguished from correlated variation. The system might be biased to interpret uncorrelated as correlated variation because the statistics of random and systematic error probably change at different rates (Körding, Tenenbaum, & Shadmehr, 2007). Another possibility is that it may simply take longer to reliably learn changes in uncorrelated than correlated variation.

No matter which explanation proves correct, the observation that changes in random perturbation do not affect rate over short time periods, while changes in feedback blur do, is an important step toward identifying the behaviorally relevant variables. It will be important in future research to determine the form of an optimal adaptor for an environment with changing statistics.

Cue combination and cue calibration

The work presented here is relevant to recent work on how noisy sensory signals are combined to determine the most likely state of an environmental property. Bayes' Law prescribes the statistically optimal method for combining noisy signals; it can be re-written as a weighted sum of sensory estimates by making some reasonable assumptions (Cochran, 1937; Landy & Kojima, 2001). For k estimators providing sensory estimators of environmental property S:

$$\hat{S}_c = \sum_{i=1}^{k} w_i \hat{S}_i; \text{ where } w_i = \frac{\sigma_i^{-2}}{\sum_{j=1}^{k} \sigma_j^{-2}}, \quad (8)$$

where $\hat{S}_c$ is the combined estimate and $\hat{S}_i$ is the estimate from the ith estimator ($\hat{S}_i = f_i(S)$). The resultant, $\hat{S}_c$, is more precise (less variable) and more accurate (closer to correct) than the individual estimates alone (Ernst & Banks, 2002; Ghahramani et al., 1997; Oruç, Maloney, & Landy, 2003) if the individual estimators are calibrated. Substantial experimental evidence supports the idea that signals from different sensory modalities (Alais & Burr, 2004; Bresciani, Dammeier, & Ernst, 2006; Ernst & Banks, 2002; Gepshtein, Burge, Ernst, & Banks, 2005; Roach, Heron, & McGraw, 2006) and from within a modality (Helbig & Ernst, 2007; Hillis et al., 2002; Landy & Kojima, 2001; Saunders & Knill, 2003) are combined in this fashion.

However, the cue-combination model in Equation 8 assumes unbiased (calibrated) estimators without providing a mechanism for achieving their calibration. A measurement of an environmental property by a miscalibrated estimator will not agree on average with measurements from calibrated estimators. Combining such measurements would decrease rather than increase accuracy, thereby undermining one of the potential benefits of cue combination. To avoid this, the nervous system must incorporate a method for calibrating the estimators relative to one another. As we said, the Kalman filter is the optimal algorithm for such re-calibration (Ghahramani et al., 1997; Korenberg & Ghahramani, 2002). Consider sensory estimates $\hat{S}_1$ and $\hat{S}_2$ produced by estimators f_1 and f_2 and a current mapping estimate between the estimators of $\hat{X}_t$. The error with each new observation of $\hat{S}_1$ and $\hat{S}_2$ is $\hat{E}_t = \hat{S}_{1t} - \hat{S}_{2t}$; the filter sets the mapping $\hat{X}_{t+1}$ to minimize the error over time (Equation 3). The measurement uncertainty is $\sqrt{\sigma_1^2 + \sigma_2^2}$. The mapping uncertainty, σ_x, represents how variable over time measurements from two sensory estimators are relative to each other. A long-standing stable relationship between estimates should yield a low value for σ_x; a historically unstable relationship should yield a high value.

Mapping uncertainty is not expressed in the prevailing cue-combination model (Equation 8). Rather, the model assumes that the mapping is fixed and that the system estimates it without error or uncertainty. Under these conditions, we expect sensory signals to be fused into one estimate. However, the mapping may change and the mapping estimate can be uncertain. Interestingly, combined cues are not always fully fused; in some cases, access to the individual signals is retained (Hillis et al., 2002). To account for the observation that signals are fused to varying degrees, a coupling prior has been proposed (Ernst, 2005); it represents the probability distribution of naturally occurring mappings between two estimator signals and determines the degree to which the signals should be fused.

We propose that the coupling prior is equivalent to the mapping uncertainty described above. Thus, we predict that signals that are fully fused, or nearly so, should be associated with estimators that adapt very slowly to conflict, and that signals that are kept independent should be associated with estimators that adapt quickly.

Conclusion

We compared the rates at which visuomotor recalibration occurs in humans and an optimal adaptor. The comparison allowed us to identify the stimulus properties that determine rate and to provide a theory-driven account of how and why interaction with the environment adjusts the rate of recalibration in reaching. We showed that adaptation rates slow down and speed up in a fashion predicted by measurement reliability and environmental statistics. The response to random perturbation is non-optimal for an

environment with stationary statistics but might be optimal for changing statistics. The specific mechanisms by which the nervous system measures reliability and statistics should be a focus of future investigation.

Methods

Apparatus and stimuli

Visual stimuli were presented on a vertical 220 × 200-cm back-projection screen illuminated by a digital projector (JVC D-ILA DLA-C15 with XGA resolution). At the 50-cm viewing distance, the screen subtended 131° × 127°. Average luminance was 3.7 cd/m^2. Subjects' heads were restrained by a chin-and-forehead rest. A graphics tablet (AIPTEK™ Hyper Pen 1200 USB) was placed horizontally below the chin rest and in front of the subject's torso where subjects could not see it. The tablet's active area was 30.4 × 22.8 cm (Figure 2A).

The visual stimuli were targets and feedback. The target stimuli were small high-contrast circles (diameter = 1°). The feedback stimuli were low-contrast (~10%) Gaussian blobs with various widths (σ_{blur}) and aspect ratios. Target and feedback stimuli were presented for 500 ms.

Procedures

Visual feedback localization experiment

For the analysis of visuomotor behavior, we needed to know how well the various visual feedback stimuli could be localized. To determine this, we conducted a visual discrimination experiment. The stimuli were Gaussian blobs with different widths. Some were isotropic with standard deviations of 4° × 4°, 8° × 8°, 12° × 12°, 16° × 16°, 20° × 20°, and 24° × 24°; others were anisotropic with standard deviations of 24° × 4° and 4° × 24°. We used a two-interval, forced-choice procedure to measure localization thresholds. In the first interval, a standard blob appeared in a random location. In the second interval, a comparison blob appeared left or right of the first. The subject indicated with a key press whether the comparison was left or right of the standard. We also ran the experiment with the comparison above or below the standard in which the subject indicated whether the comparison was above or below.

The separation between the standard and the comparison blobs was varied with the method of constant stimuli (±0.2°, ±0.4°, ±0.6°, ±0.9° for the smallest blobs; ±1.2°, ±2.3°, ±3.5°, ±4.6° for the largest blobs). All blobs within an experimental block had the same dimensions. Each block had 160 trials. Each subject completed two blocks for each condition. Block order was randomized for each subject. In all, they completed 2560 trials (40 presentations of 8 comparison values for 8 blob sizes). We fit the resulting data with a cumulative Gaussian using a maximum-likelihood criterion (Wichmann & Hill, 2001a, 2001b). The standard deviation of the best-fitting Gaussian was the just-noticeable difference or JND. We divided the JND by $\sqrt{2}$ to determine the standard deviation of the underlying estimate (Wickens, 2002). JNDs were averaged across subjects to obtain the localization JND for each stimulus condition.

Reaching adaptation experiments

Subjects reached with a stylus held by the preferred hand. The starting position for each movement was a small knob close to the near end of the tablet. Before each trial, the subject moved the stylus to the starting position. When the subject pressed a button on the stylus, the trial began. The visual target appeared in a random screen position within an invisible 7.7° square centered straight ahead. Subjects were instructed to rapidly move the stylus to the position on the tablet that corresponded to the target position on the projection screen. Average movement duration was ~600 ms and average movement distance was ~15 cm. The visual feedback stimulus then appeared immediately on the projection screen at a position associated (via the visuomotor mapping) with the reach endpoint. Movement endpoint was determined by the stylus' first contact with the tablet after leaving the start position. Subjects were instructed to reach to the position that would make the center of the feedback stimulus appear in the same location as the target. Reach error E_t was defined as the difference between visual target and feedback locations. To perform the task, subjects had to learn the visuomotor mapping, X_t, from tablet coordinates to screen coordinates (Figure 2B). None had difficulty doing so.

Before each experimental session, subjects completed a training session of 200 trials to become familiar with the task and to learn the mapping between screen and tablet coordinates. During training, the visual feedback stimulus was a small circle that could be accurately localized.

Each experimental condition consisted of 160 trials. In the 60 pre-step phase trials, visual feedback appeared at the corresponding position of the reach endpoint: tablet position (x', y') generated feedback at screen position $(x + X_{t_x}, y + X_{t_y})$. In the step phase (50 trials), the feedback was displaced up and to the right by 8.2° from the corresponding position of the reach endpoint: tablet position (x', y') produced feedback at screen position $(x + X_{t_x} + 5.8°, y + X_{t_y} + 5.8°)$. In the post-step phase (50 trials), the step change was removed: tablet position (x', y') again generated feedback at screen position $(x + X_{t_x}, y + X_{t_y})$ (Figure 2C).

Each experiment contained several blocks that differed in the visual feedback provided and in the statistical properties of the mapping. To prevent practice effects, the blocks were presented in random orders that differed across subjects. Twenty practice trials, just like those in

the training session, were presented at the beginning of each block. An experimental session lasted about 45 min.

The experiment was designed to minimize the possible influence of proprioceptive feedback. First, subjects reached in a plane orthogonal to the plane in which the visual feedback appeared so the proprioceptive location was never similar to the visual feedback location. Thus, to use the proprioceptive signal, it would have had to be converted from tablet coordinates to screen coordinates, and that conversion would presumably make the signal less reliable. Second, subjects were instructed to minimize the discrepancy between the visual feedback and the visual target location, so the proprioceptive signal would not have been useful for satisfying this instruction. Although we think it is unlikely, it is nevertheless possible that proprioceptive feedback influenced behavior. If this were the case, it would have the effect of reducing the estimated error. For example, consider the first trial after the step change: the visual system would indicate a large error and the proprioceptive system would indicate zero error. If the magnitude of estimated errors were reduced, slower adaptation would result. This, in turn, would lead to an underestimation of the system estimate of mapping uncertainty, $\hat{\sigma}_x$ (see Figure 3E).

Post-step data followed the same trends as the step data, but we did not analyze it because post-step data are difficult to deal with for two reasons. (1) Post-step adaptation is generally faster (consistent with previous findings; Devane, 1968) and discriminating between fast rates is more difficult than discriminating between slow rates (see Figure 5D caption). (2) The error profiles in the post-step phase have an undetermined starting point because the error on the first trial of the post-step phase depends on the error of the final trial of the step phase. This necessitates a three-parameter exponential fit to individual subject error profiles instead of the two-parameter fit that was used throughout the paper.

Observers

Twenty-four subjects between the ages of 18 and 40 years participated in each reaching experiment. Ten subjects between the ages of 18 and 40 years participated in the visual localization experiment. All subjects were unaware of the experimental hypotheses and were paid for their participation. All had normal or corrected-to-normal vision.

Appendix A

This discussion follows Maybeck (1979), but with the terms and equations tailored to the issues raised in the manuscript.

Consider two estimates $\hat{X}_1$ and $\hat{X}_2$ of the same environmental property. Equation 8 in the main text is the maximum-likelihood estimate for a linear system with Gaussian-distributed noise and conditional independence between the estimates. Rewriting that equation:

$$\hat{X}_c = w_1\hat{X}_1 + w_2\hat{X}_2; \quad \text{where} \quad \begin{aligned} w_1 &= \frac{\sigma^2_{\hat{X}_2}}{\sigma^2_{\hat{X}_1} + \sigma^2_{\hat{X}_2}} \\ w_2 &= \frac{\sigma^2_{\hat{X}_1}}{\sigma^2_{\hat{X}_1} + \sigma^2_{\hat{X}_2}}. \end{aligned} \tag{A1}$$

The variance is

$$\sigma^2_{\hat{X}_c} = \frac{\sigma^2_{\hat{X}_1}\sigma^2_{\hat{X}_2}}{\sigma^2_{\hat{X}_1} + \sigma^2_{\hat{X}_2}}. \tag{A2}$$

Now consider a visuomotor mapping X drifting in time and an estimate of that mapping $\hat{X}$. The best estimate at the current time just before a measurement is made is $\hat{X}_t^-$, and the measurement of the mapping at the current time-step is $\hat{Z}_t$. Again assuming Gaussian noise and conditional independence, the best estimate of the mapping (immediately after the current measurement) is

$$\hat{X}_t^+ = w_{\hat{X}_t^-}\hat{X}_t^- + w_{\hat{Z}_t}\hat{Z}_t; \quad \text{where} \quad \begin{aligned} w_{\hat{X}_t^-} &= \frac{\sigma^2_{\hat{Z}_t}}{\sigma^2_{\hat{X}_t^-} + \sigma^2_{\hat{Z}_t}} \\ w_{\hat{Z}_t} &= \frac{\sigma^2_{\hat{X}_t^-}}{\sigma^2_{\hat{X}_t^-} + \sigma^2_{\hat{Z}_t}}. \end{aligned} \tag{A3}$$

Adding and subtracting $w_{\hat{Z}_t}\hat{X}_t^-$ yields

$$\hat{X}_t^+ = w_{\hat{X}_t^-}\hat{X}_t^- + w_{\hat{Z}_t}\hat{Z}_t - w_{\hat{Z}_t}\hat{X}_t^- + w_{\hat{Z}_t}\hat{X}_t^-. \tag{A4}$$

Regrouping and using the fact that $w_{\hat{Z}_t} + w_{\hat{X}_t^-} = 1$ yields

$$\hat{X}_t^+ = \hat{X}_t^- + w_{\hat{Z}_t}(\hat{Z}_t - \hat{X}_t^-). \tag{A5}$$

Setting $K = w_{\hat{Z}_t}$ and defining $\hat{E}_t = \hat{Z}_t - \hat{X}_t^-$:

$$\hat{X}_t^+ = \hat{X}_t^- + K(\hat{E}_t), \tag{A6}$$

which is the standard form of the measurement-update equation of the Kalman filter. The variance of that best estimate is

$$\sigma^2_{\hat{X}_t^+} = \frac{\sigma^2_{\hat{X}_t^-}\sigma^2_{\hat{Z}_t}}{\sigma^2_{\hat{X}_t^-} + \sigma^2_{\hat{Z}_t}}, \tag{A7}$$

and the Kalman gain is

$$K_t = \frac{\sigma^2_{\hat{X}_t^-}}{\sigma^2_{\hat{X}_t^-} + \sigma^2_{\hat{Z}_t}}. \quad \text{(A8)}$$

These are the three measurement-update equations of the Kalman filter.

Now consider incorporating the mapping change dynamics into the mapping estimation algorithm. For simplicity, assume that the best model of the mapping changes has the form $\frac{\partial X}{\partial t} = u + x$, where u is the average change in mapping and x is the uncertainty associated with the average change in mapping. Assuming unit time steps, the best estimate of the mapping and the variance of that estimate just before the next measurement are

$$\hat{X}^-_{t+1} = \hat{X}^+_t + u, \quad \text{(A9)}$$

$$\sigma^2_{\hat{X}^-_{t+1}} = \sigma^2_{\hat{X}^+_t} + \sigma^2_x. \quad \text{(A10)}$$

These are the so-called state-update equations of the Kalman filter. Once the system has reached steady state, $\sigma^2_{\hat{X}_{t+1}^-} = \sigma^2_{\hat{X}_t^-}$ so by Equation A10:

$$\sigma^2_{\hat{X}_t^-} = \sigma^2_{\hat{X}^+_t} + \sigma^2_x. \quad \text{(A11)}$$

Substituting into Equation A8 yields

$$K_t = \frac{\sigma^2_{\hat{X}^+_t} + \sigma^2_x}{\sigma^2_{\hat{X}^+_t} + \sigma^2_x + \sigma^2_{\hat{Z}_t}}, \quad \text{(A12)}$$

which is the exact form of the Kalman gain at steady state. We can predict K from σ^2_x and $\sigma^2_{\hat{Z}_t}$ by substituting Equation A11 into Equation A7 and rearranging:

$$\left(\sigma^2_{\hat{X}^+_t}\right)^2 + \sigma^2_x\left(\sigma^2_{\hat{X}^+_t}\right) - \sigma^2_x\sigma^2_{\hat{Z}_t} = 0. \quad \text{(A13)}$$

To find $\sigma^2_{\hat{X}^+_t}$, we use the quadratic formula. One positive value is yielded:

$$\sigma^2_{\hat{X}^+_t} = \frac{-\sigma^2_x + \sqrt{(\sigma^2_x)^2 + 4\sigma^2_x\sigma^2_{\hat{Z}_t}}}{2}. \quad \text{(A14)}$$

Substituting Equation A14 into Equation A12 yields an analytic solution for K_t, based only on the measurement and mapping uncertainties, once the system has reached steady state.

Appendix B

From Appendix A, the general model of how the visuomotor mapping drifts in time is $\frac{\partial X}{\partial t} = u + x$, where u is the expected change in mapping with time, and x is the uncertainty associated with those mapping changes. The expected value of u is 0 whether the mapping change is a random walk or a random perturbation. The difference between random walks and random perturbations appears in x. Changes in mapping do not accumulate over time with perturbations—mapping is not a function of time—so the variance of x is 0 and $\sigma_x = 0$. With random walks, mapping changes do accumulate over time, so $\sigma_x = \sigma_{\text{walk}}\sqrt{t_2 - t_1}$. In our experiments, the interval between observations was always one trial: $t_2 - t_1 = 1$; thus, the standard deviation of the mapping estimate $\hat{X}_t$ should increase by σ_{walk} from one trial to the next when the mapping is changed via a random walk.

Appendix C

Here we examine the results of Baddeley et al. (2003): In particular, their conclusion that varying the amount of random perturbation changed the rate of adaptation. Subjects in their study performed a task like ours. The mapping between reach endpoint and visual feedback was varied by random walks, by random perturbations, and by their sum. The researchers assessed adaptation rate indirectly by calculating efficiency:

$$E = \frac{\text{MSE}_{\text{filter}}}{\text{MSE}_{\text{subject}}}, \quad \text{(C1)}$$

where $\text{MSE}_{\text{filter}}$ and $\text{MSE}_{\text{subject}}$ are the mean squared reaching errors of a Kalman filter and human subject, respectively. The filter's parameters, $\hat{\sigma}_x$ and $\sigma_{\hat{z}}$, were set equal to σ_{walk} and σ_{perturb}, respectively. Therefore, as σ_{perturb} was increased, filter adaptation rate decreased (Equations 4 and 5). Their results are plotted in Figure C1A. When σ_{walk} was large, efficiency was reasonably high and did not vary with σ_{perturb}. To explain the constant efficiency at large σ_{walk}, the authors reasoned that the Kalman filter and humans must have responded similarly to increases in σ_{perturb} by increasing the $\sigma_{\hat{z}}$ estimate. For small σ_{walk}, the increasing efficiency with σ_{perturb} was attributed to motor noise, which was present in humans but not in the model. From this, they concluded that the subjects' adaptation rate slowed with increases in perturbation.

One could, however, obtain the same efficiencies even if human adaptation rate did not slow with perturbation. To show this, we presented mixtures of a random walk and random perturbation to a Kalman filter and a simulated human observer. The filter and human simulations were

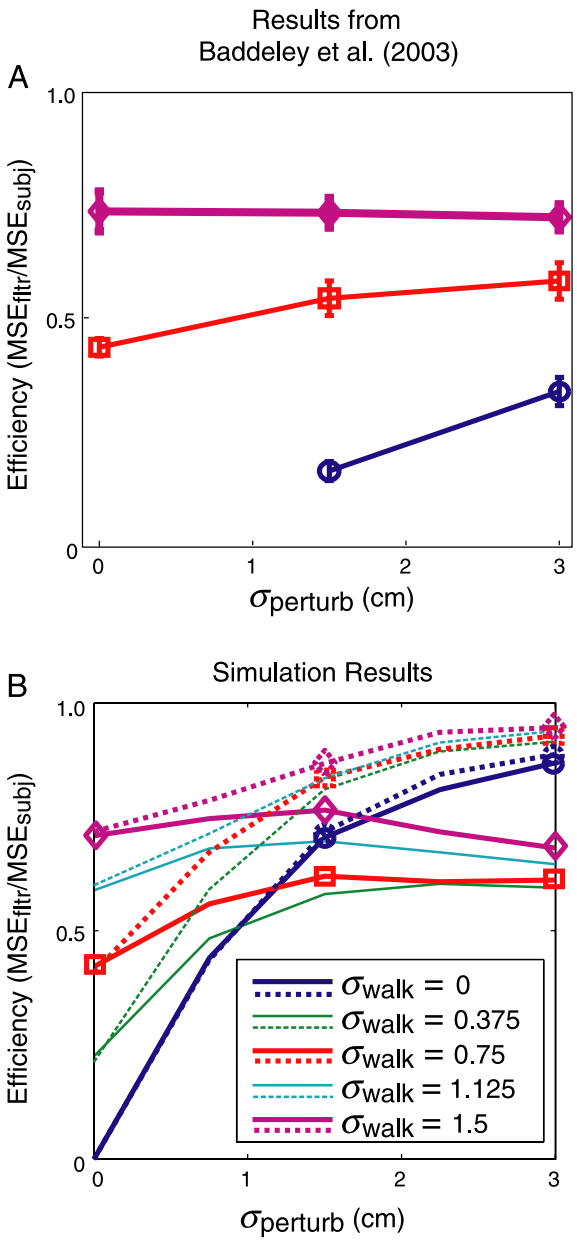

Figure C1. Efficiency ($MSE_{filter}/MSE_{subject}$) plotted as a function of $\sigma_{perturb}$ for different values of σ_{walk}. (A) Results from Baddeley et al. (2003). When σ_{walk} was high, efficiency was relatively high and constant as $\sigma_{perturb}$ varied. When σ_{walk} was small, efficiency increased with increasing $\sigma_{perturb}$. The authors argued the latter effect was caused by motor error, which was present in the human data but not in the Kalman filter. (B) Our simulation results. Solid curves: The Kalman filter's parameter $\sigma_{\hat{z}}$ was equal to $\sigma_{perturb}$, and the simulated subjects' $\sigma_{\hat{z}}$ was constant. Dashed lines: for comparison, efficiencies from a simulation in which $\sigma_{\hat{z}}$ for both the filter and the simulated subjects was unaffected. Thin lines indicate conditions that Baddeley et al. did not directly test (σ_{walk} = 0.375 and 1.125 cm). Note the consequence when the feedback uncertainty parameter in the simulated humans is not affected by perturbation but they are modeled as if it is. For a given perturbation variance, MSE_{filter} will be lower than it otherwise would be, so efficiency will be lower than it should be. The effect is disproportionately larger at large perturbations and therefore could be responsible for the constant efficiency levels at high σ_{walk} in Baddeley et al. In both simulations, the simulated subject had additive Gaussian motor noise with a standard deviation of 1 cm. In cases in which $\sigma_{\hat{z}}$ was unaffected by $\sigma_{perturb}$, $\sigma_{\hat{z}}$ was set to 0.133 cm, a reasonable value for the visual JND for localizing feedback of the size used in Baddeley et al.

the same with two important differences: (1) the filter's estimate of $\sigma_{\hat{z}}$ was set equal to $\sigma_{perturb}$, while the human estimate of $\sigma_{\hat{z}}$ was unaffected by $\sigma_{perturb}$, and (2) the human had motor noise while the filter had none. With those assumptions, we conducted a simulation to calculate efficiencies. The results are shown in Figure C1B. When σ_{walk} was large, efficiency was constant for a wide range of $\sigma_{perturb}$. When σ_{walk} was small, efficiency increased

with increasing $\sigma_{perturb}$. These results are similar to those of Baddeley and colleagues. It is therefore possible that perturbation had no effect on adaptation rates in their experiments, as we observed. We conclude that their observation of constant efficiency under some conditions does not necessarily support the conclusion that the nervous system's estimate of feedback uncertainty is affected by random perturbation.

Acknowledgments

We wish to thank Alexandra Reichenbach and Thomas Wiecki for running subjects and Joern Diedrichsen, Rich Ivry, and Simon Watt for helpful comments on the manuscript. This research was supported by AOF's William C. Ezell Fellowship (JB), "ImmerSence" IST-2006-027141 (MOE), Human Frontiers Science Program Grant RGP0003/2006-C (MOE), and NIH Grant EY12851 (JB, MSB).

Commercial relationships: none.
Corresponding author: Johannes Burge.
Email: jburge@berkeley.edu.
Address: 509 Minor Hall, School of Optometry, UC Berkeley, Berkeley, CA 94720-2020, USA.

References

Adams, W. J., Banks, M. S., & van Ee, R. (2001). Adaptation to three-dimensional distortions in human vision. *Nature Neuroscience, 4,* 1063–1064. [PubMed] [Article]

Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology, 14,* 257–262. [PubMed] [Article]

Baddeley, R. J., Ingram, H. A., & Miall, R. C. (2003). System identification applied to a visuomotor task: Near-optimal human performance in a noisy changing task. *Journal of Neuroscience, 23,* 3066–3075. [PubMed] [Article]

Bedford, F. L. (1989). Constraints on learning new mappings between perceptual dimensions. *Journal of Experimental Psychology, 15,* 232–248.

Bedford, F. L. (1993a). Perceptual learning. In D. Medin (Ed.), *The psychology of learning and motivation* (pp. 1–60). New York: Academic Press.

Bedford, F. L. (1993b). Perceptual and cognitive spatial learning. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 517–530. [PubMed]

Bhushan, N., & Shadmehr, R. (1999). Computational nature of human adaptive control during learning of reaching movements in force fields. *Biological Cybernetics, 81,* 39–60. [PubMed]

Brandt, T., & Daroff, R. B. (1980). The multisensory physiological and pathological vertigo syndromes. *Annals of Neurology, 7,* 195–203. [PubMed]

Bresciani, J. P., Dammeier, F., & Ernst, M. O. (2006). Vision and touch are automatically integrated for the perception of sequences of events. *Journal of Vision, 6*(5):2, 554–564, http://journalofvision.org/6/5/2/, doi:10.1167/6.5.2. [PubMed] [Article]

Cheng, S., & Sabes, P. N. (2006). Modeling sensorimotor learning with dynamical systems. *Neural Computation, 18,* 760–793. [PubMed]

Cochran, W. G. (1937). Problems arising in the analysis of a series of similar experiments. *Royal Statistical Society, 4,* 102–118.

Crowell, J. A., & Banks, M. S. (1996). Ideal observer for heading judgments. *Vision Research, 36,* 471–490. [PubMed]

Devane, J. R. (1968). Proaction in the recovery from practice under visual displacement. *Perceptual and Motor Skills, 27,* 411–416.

Diedrichsen, J., Hashambhoy, Y., Rane, T., & Shadmehr, R. (2005). Neural correlates of reach errors. *Journal of Neuroscience, 25,* 9919–9931. [PubMed] [Article]

Donchin, O., Francis, J. T., & 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. *Journal of Neuroscience, 23,* 9032–9045. [PubMed]

Ernst, M. O. (2005). A Bayesian view on multimodal cue integration. In G. Knoblich, M. Grosjean, I. Thornton, & M. Shiffrar (Eds.), *Perception of the human body from the inside out*. New York: Oxford University Press.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415,* 429–433. [PubMed]

Geisler, W. S. (1989). Sequential ideal-observer analysis of visual discriminations. *Psychological Review, 96,* 267–314. [PubMed]

Gepshtein, S., & Banks, M. S. (2003). Viewing geometry determines how vision and haptics combine in size perception. *Current Biology, 13,* 483–488. [PubMed] [Article]

Gepshtein, S., Burge, J., Ernst, M. O., & Banks, M. S. (2005). The combination of vision and touch depends on spatial proximity. *Journal of Vision, 5*(11):7, 1013–1023, http://journalofvision.org/5/11/7/, doi:10.1167/5.11.7. [PubMed] [Article]

Ghahramani, Z., Wolpert, D. M., & Jordan, M. I. (1997). Computational models of sensorimotor integration. In P. G. Morasso & V. Sanguineti (Eds.), *Self-organization, computational maps and motor control*. Amsterdam: Elsevier Press.

Harris, C. S. (1963). Adaptation to displaced vision: Visual, motor, or proprioceptive change? *Science, 140,* 812–813. [PubMed]

Hay, J. C., & Pick, H. L., Jr. (1966). Visual and proprioceptive adaptation to optical displacement of the visual stimulus. *Journal of Experimental Psychology, 71,* 150–158. [PubMed]

Heathcote, A., Brown, S., & Mewhort, D. J. (2000). The power law repealed: The case for an exponential law of practice. *Psychonomic Bulletin & Review, 7,* 185–207. [PubMed]

Helbig, H. B., & Ernst, M. O. (2007). Optimal integration of shape information from vision and touch. *Experimental Brain Research, 179,* 595–606. [PubMed] [Article]

Hillis, J. M., Ernst, M. O., Banks, M. S., & Landy, M. S. (2002). Combining sensory information: Mandatory fusion within, but not between senses. *Science, 298,* 1627–1630. [PubMed]

Hofman, P. M., Van Riswick, J. G., & Van Opstal, A. J. (1998). Relearning sound localization with new ears. *Nature Neuroscience, 1,* 417–421. [PubMed] [Article]

Kagerer, F. A., Contreras-Vidal, J. L., & Stelmach, G. E. (1997). Adaptation to gradual as compared with sudden visuo-motor distortions. *Experimental Brain Research, 115,* 557–561. [PubMed] [Article]

Kayargadde, V., & Martens, J. B. (1996). Perceptual characterization of images degraded by blur and noise: Experiments. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 13,* 1166–1177. [PubMed]

King, A. J., Kacelnik, O., Mrsic-Flogel, T. D., Schnupp, J. W., Parsons, C. H., & Moore, D. R. (2001). How plastic is spatial hearing? *Audiology and Neuro-Otology, 6,* 182–186. [PubMed]

King, A. J., Parsons, C. H., & Moore, D. R. (2000). Plasticity in the neural coding of auditory space in the mammalian brain. *Proceedings of the National Academy of Sciences of the United States of America, 97,* 11821–11828. [PubMed] [Article]

Knudsen, E. I., & Knudsen, P. F. (1985). Vision guides the adjustment of auditory localization in young barn owls. *Science, 230,* 545–548. [PubMed]

Kording, K. P., Tenenbaum, J. B., & Shadmehr, R. (2007). The dynamics of memory as a consequence of optimal adaptation to a changing body. *Nature Neuroscience, 10,* 779–786. [PubMed]

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

Korenberg, A. T., & Ghahramani, Z. (2002). A Bayesian view of motor adaptation. *Current Psychology of Cognition, 21,* 537–564.

Landy, M. S., & Kojima, H. (2001). Ideal cue combination for localizing texture-defined edges. *Journal of the Optical Society of America A, Optics, Image Science, and Vision, 18,* 2307–2320. [PubMed]

Levi, D. M., & Klein, S. A. (1985). Vernier acuity, crowding and amblyopia. *Vision Research, 25,* 979–991. [PubMed]

Maybeck, P. (1979). *Stochastic models, estimation, and control*. New York: Academic Press.

Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature, 434,* 387–391. [PubMed] [Article]

Oruç, I., Maloney, L. T., & Landy, M. S. (2003). Weighted linear cue combination with possibly correlated error. *Vision Research, 43,* 2451–2468. [PubMed]

Roach, N. W., Heron, J., & McGraw, P. V. (2006). Resolving multisensory conflict: A strategy for balancing the costs and benefits of audio-visual integration. *Proceedings of the Royal Society B: Biological Sciences, 273,* 2159–2168. [PubMed] [Article]

Saunders, J. A., & Knill, D. C. (2003). Humans use continuous feedback from the hand to control fast reaching movements. *Experimental Brain Research, 152,* 341–352. [PubMed] [Article]

Shadmehr, R., & Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *Journal of Neuroscience, 14,* 3208–3224. [PubMed] [Article]

Smith, M., & Shadmehr, R. (2004). Modulation of the rate of error-dependent learning by the statistical properties of the task. In *Society for neuroscience conference, advances in computational motor control.*

Sober, S. J., & Sabes, P. N. (2003). Multisensory integration during motor planning. *Journal of Neuroscience, 23,* 6982–6992. [PubMed] [Article]

Sober, S. J., & Sabes, P. N. (2005). Flexible strategies for sensory integration during motor planning. *Nature Neuroscience, 8,* 490–497. [PubMed]

van Beers, R. J., Wolpert, D. M., & Haggard, P. (2002). When feeling is more important than seeing in sensorimotor adaptation. *Current Biology, 12,* 834–837. [PubMed] [Article]

von Helmholtz, H. (1867). *Handbuch der physiologischen optik*. Leipzig: Leopold Voss.

Watson, A. B. (1987). The ideal observer concept as a modelling tool. In *Frontiers in visual science: Proceedings of the 1985 symposium* (pp. 32–37).

Waugh, S. J., & Levi, D. M. (1993a). Visibility, luminance, and vernier acuity. *Vision Research, 33,* 527–538. [PubMed]

Waugh, S. J., & Levi, D. M. (1993b). Visibility, timing, and vernier acuity. *Vision Research, 33,* 505–526. [PubMed]

Wichmann, F. A., & Hill, N. J. (2001a). The psychometric function: I. Fitting, sampling and goodness of fit. *Perception & Psychophysics, 63,* 1293–1313. [PubMed] [Article]

Wichmann, F. A., & Hill, N. J. (2001b). The psychometric function: II. Bootstrap-based confidence intervals and sampling. *Perception & Psychophysics, 63,* 1314–1329. [PubMed] [Article]

Wickens, T. D. (1999). Measuring the time course of retention. In C. Izawa (Ed.), *On memory: Evolution, progress, and reflections on the 30th anniversary of the Atkinson-Shiffrin model*. Mahwah, NJ: Erlbaum.

Wickens, T. D. (2002). *Elementary signal detection theory*. New York: Oxford University Press.

Zwiers, M. P., Van Opstal, A. J., & Paige, G. D. (2003). Plasticity in human sound localization induced by compressed spatial vision. *Nature Neuroscience, 6,* 175–181. [PubMed]