

Flexible, Task-Dependent Use of Sensory Feedback to Control Hand Movements

David C. Knill,¹ Amulya Bondada,² and Manu Chhabra²

¹Center for Visual Sciences and ²Department of Brain and Cognitive Sciences, University of Rochester, Rochester, New York 14618

We tested whether changing accuracy demands for simple pointing movements leads humans to adjust the feedback control laws that map sensory signals from the moving hand to motor commands. Subjects made repeated pointing movements in a virtual environment to touch a button whose shape varied randomly from trial to trial—between squares, rectangles oriented perpendicular to the movement path, and rectangles oriented parallel to the movement path. Subjects performed the task on a horizontal table but saw the target configuration and a virtual rendering of their pointing finger through a mirror mounted between a monitor and the table. On one-third of trials, the position of the virtual finger was perturbed by ± 1 cm either in the movement direction or perpendicular to the movement direction when the finger passed behind an occluder. Subjects corrected quickly for the perturbations despite not consciously noticing them; however, they corrected almost twice as much for perturbations aligned with the narrow dimension of a target than for perturbations aligned with the long dimension. These changes in apparent feedback gain appeared in the kinematic trajectories soon after the time of the perturbations, indicating that they reflect differences in the feedback control law used throughout the duration of movements. The results indicate that the brain adjusts its feedback control law for individual movements “on demand” to fit task demands. Simulations of optimal control laws for a two-joint arm show that accuracy demands alone, coupled with signal-dependent noise, lead to qualitatively the same behavior.

Introduction

Research has shown that humans use online visual information about both a target object and the moving hand throughout goal-directed hand movements to control the movements (Goodale et al., 1986; Pélisson et al., 1986; Prablanc and Martin, 1992; Sarlegna et al., 2003; Saunders and Knill, 2003, 2004, 2005). In an optimal system, both the reliability of sensory and motor signals and the goals and constraints of the motor task (Todorov and Jordan, 2002; Li and Todorov, 2007) determine how different sensory signals influence online control of movements. The first factor acts via its influence on state estimation—noisy feedback signals should contribute less to online control than reliable ones because they have less influence on internal state estimates. Recent studies have shown that humans exhibit this reliability-based weighting of online visual signals both for visual feedback signals from the moving hand (Körding and Wolpert, 2004; Saunders and Knill, 2004, 2005) and for visual signals from a target that jumps at movement onset (Izawa and Shadmehr, 2008).

Task constraints act by shaping the control law that generates motor commands. Both subjective (e.g., energy conservation) and objective (e.g., minimizing endpoint variance) task constraints have been used to account for average movement trajectories of ballistic movements (Flash and Hogan, 1985; Uno et al.,

1989; Harris and Wolpert, 1998). Similar constraints also shape the sensory feedback component of an optimal feedback-driven controller (Todorov and Jordan, 2002; Todorov and Li, 2004). The optimal control law for any given task will tend to minimize variance in task-relevant dimensions at the expense of increasing variance in irrelevant dimensions—a form of “minimum-intervention principle.”

A number of recent studies provide evidence that the CNS adapts its online control law to changing task demands (Liu and Todorov, 2007; Diedrichsen and Dowling, 2009; Diedrichsen and Gush, 2009). When performing a pointing task, subjects corrected less for visible shifts in target position when they had to stop at the target than when they were allowed to hit it (Liu and Todorov, 2007). Optimal controllers derived for the different task conditions showed a similar behavior—a pattern that results from an inherent stability/accuracy trade-off imposed by motor noise. In a bimanual control task, subjects showed coupled corrections of the left and right hands when both hands controlled the movement of a single cursor to a target, but decoupled corrections when controlling independent cursors to two different targets (Diedrichsen and Dowling, 2009; Diedrichsen and Gush, 2009).

Previous studies blocked different task conditions; some explicitly measured the rate of adaptation to changing task demands. We asked whether the CNS shapes its feedback control law to the accuracy demands of individual pointing movements when those demands change from trial to trial (by changing the shapes of targets). Subjects' corrective responses to visual perturbations of their fingers changed from trial to trial to match the accuracy constraints of different shapes. Simulations of optimal

Received July 18, 2009; revised April 23, 2010; accepted July 20, 2010.

This work was supported by National Institutes of Health Grant R01 EY013319 (D.C.K.).

Correspondence should be addressed to David C. Knill, Center for Visual Sciences, 274 Meliora Hall, University of Rochester, Rochester, NY 14618. E-mail: knill@cvs.rochester.edu.

DOI:10.1523/JNEUROSCI.3522-09.2011

Copyright © 2011 the authors 0270-6474/11/311219-19\$15.00/0

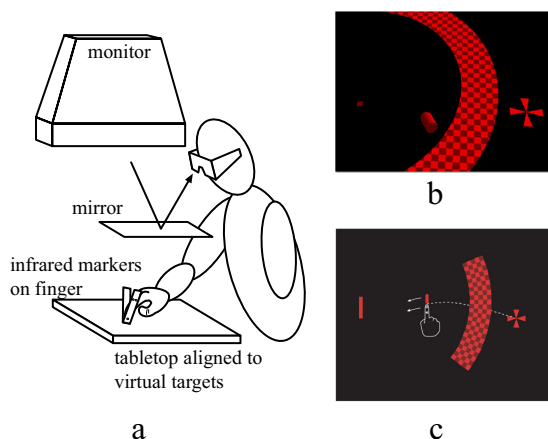


Figure 1. *a*, A side view of the experimental apparatus used in the experiments. *b*, Subjects' view of the tabletop environment (seen from the top) halfway through a trial. Subjects viewed the display stereoscopically through LCD stereo glasses with geometrically correct rendering of the three-dimensional scene. The start position and target rectangle appeared flat in the plane of the tabletop (approximately coextensive with the virtual image of the monitor). The semi-circular occluder appeared 10 cm above the tabletop, so that subjects' fingers disappeared as they passed behind the occluder. *c*, A schematic rendering of a perturbation trial, in which the position of the virtual finger was shifted up by 1 cm (in the plane of the tabletop) relative to the true finger position.

controllers for the different shapes show that the observed changes in feedback control result naturally from optimizing a simple performance constraint—maximizing the probability of hitting a target. We show that subjects' corrective responses are linear—the time course of corrections satisfies the superposition property—corrections to a visual perturbation that is a weighted sum of two other perturbations for visual perturbations are linear superpositions of corrections to those two perturbations.

Materials and Methods

Overview of experimental design

Since we were interested in studying the feedback control law used during “normal” movements, we used an experimental paradigm in which we perturbed the visual position of the moving hand without subjects' awareness of the perturbations. Subjects pointed to targets presented in a virtual environment that included a visual display of their fingers. We measured the time course and magnitude of subjects' corrective movements in response to small, undetected (consciously) perturbations of the virtual finger's position. We used small perturbations of the virtual finger when it disappeared behind an occluder, so that subjects were never aware of the changes in finger position induced by a perturbation. We randomly varied the accuracy demands of the pointing task from trial to trial by changing the shape of the target [e.g., in experiment 1, we used square, vertical rectangle, horizontal rectangle targets (Fig. 1)].

The hypothesis that humans shape their feedback control laws anew for each movement based on the accuracy demands of the movement makes a clear qualitative prediction—that subjects will correct less for perturbations of the finger in the “short” dimension of a target button than in the “long” dimension. Differences in feedback corrections should appear on a trial-by-trial basis (conditioned on the shape of the target button presented in each trial) and should appear early in movements. To study the constraints that might drive such behavior, we derived different optimal control laws for each different target button shape using cost functions that include weighted terms for errors in task execution and for smoothness or energy consumption. In this way, we were able to determine whether optimizing task performance alone would predict the changes in feedback control laws that we observed or whether one also needs to resort to energy minimization as an explanation for the observed behavior.

Experiment 1 tested the basic predictions of the theory of stochastic optimal control using square and rectangular target shapes. We found that subjects did adjust their feedback control laws on a trial-by-trial basis to match the shape of the target—correcting less for motor errors in the direction of the long axis of a rectangular target than in the short axis. Experiment 2 tested the “linearity” of the control law by measuring subjects' corrective responses to perturbations in multiple directions and by replacing the square targets used in experiment 1 with targets shaped as crosses. A linear controller (or a system whose kinematic input–output mapping is at least approximately linear) would show corrections for errors in one direction that were a predictable linear sum of the corrections for errors in two other directions. Furthermore, this would limit the spatial pattern of corrections to a complex shape like a cross to appear suboptimal, since for such a shape, the error constraints lead to an optimal pattern of corrections in which corrections for errors in diagonal directions (relative to the arms of the cross) would be larger than corrections for errors in the orthogonal directions of the arms of the cross—a pattern that cannot arise from a linear system.

Subjects performed a simple pointing task in which they moved their right index finger from a starting position on the right-hand side of a tabletop to touch a target button on the left-hand side of the tabletop. As illustrated in Figure 1, subjects viewed a virtual scene in which starting positions and target buttons were displayed to be spatially coincident with the tabletop. A subject's index finger was displayed within the virtual environment as a virtual rendering of the finger. In most trials, the virtual finger was coincident with a subject's real finger, but on perturbation trials, the virtual finger was offset by 1 cm from the subject's real finger when it went behind an annular occluder that was rendered to appear 10 cm above the tabletop near the starting position. The occluder served to mask the visual transients associated with the perturbation. All of the subjects reported being unaware of the perturbations when questioned after the experiment, even when told what they would have looked like. By measuring the time course of subjects' corrective responses to the perturbations, we obtained data reflecting the control law that subjects used to adjust movements online based on visual feedback from the finger.

In experiment 1, three different target button shapes were used in the stimuli, interleaved randomly from trial to trial in each experimental block—squares, vertical rectangles (oriented perpendicular to the axis between the starting position and the target), and horizontal rectangles (oriented to be aligned with the axis between the starting position and the target). For simplicity, we refer to vertical as the direction perpendicular to the start–target axis and horizontal as the direction parallel to the axis. In reality, the orientation of the start–target axis varied randomly around the horizontal. On two-thirds of the trials in the experiment, the virtual finger remained coaligned with a subject's real finger. These trials served as the basis for measuring baseline aspects of motor performance (e.g., changes in mean endpoints and endpoint distributions as a function of target button shape). On one-third of the trials, the virtual finger was perturbed by 1 cm in either the vertical or horizontal directions (as defined above). Positive and negative perturbations were randomly interleaved within an experimental block. These trials served as the basis for measuring feedback-driven corrections. In experiment 2, the square buttons were replaced by crosses and perturbation trials included 1 cm shifts of the virtual finger position in a diagonal direction relative to the axis between the starting position and the target.

Subjects

Different sets of eight subjects participated in the two experiments. All were naive to the purposes of the experiment and were paid to participate. Subjects were undergraduates at the University of Rochester who provided informed consent in accordance with guidelines from the University of Rochester Research Subjects Review Board. Our apparatus required that subjects use their right hand, so only right-handed subjects were accepted.

Experiment 1

Apparatus and display. Visual displays were presented in stereo on a computer monitor viewed through a mirror (Fig. 1), using CrystalEyes shutter glasses to present different stereo views to the left and right eyes.

The mirror was half-silvered to allow subjects to view both rendered stimuli and real objects (e.g., a subject's finger) during calibration. The left and right eye views were accurate perspective renderings of the simulated environment. Displays had a resolution of 1024×768 pixels and a refresh rate of 120 Hz (60 Hz for each eye's view). The stimuli and feedback were all drawn in red to take advantage of the comparatively faster red phosphor of the monitor and prevent interocular cross talk. During the experiment, a matte black occluder was positioned behind the mirror to prevent view of a subject's hand. Visual feedback about hand position was instead provided by a virtual finger that moved in real time along with the subject's actual finger. The virtual finger was represented as a cylinder with a rounded, hemispherical tip. It had a 1 cm radius and was 5 cm long from base to tip.

A starting position was rendered as a cross on the right-hand side of the display (Fig. 1). Targets for subjects' movements were squares (1×1 cm), vertical rectangles (1×6 cm), or horizontal rectangles (6×1 cm), positioned 30 cm away from the starting position. Vertical rectangles were oriented perpendicular to the axis between starting and target positions. Horizontal rectangles were oriented parallel to this axis. The target shapes were rendered to appear aligned with an unseen tabletop (~ 55 cm from the eyes), which provided consistent haptic feedback when a subject's finger touched a target. The positions of the starting cross and the target rectangle were symmetric around the center of the virtual tabletop. On each trial, a start-target axis orientation was randomly chosen between 10° above and 10° below the horizontal. The start cross was displayed 15 cm to the right of center and the target 15 cm to the left of center along this axis.

In addition to the start position, target button, and the virtual fingertip, displays included a planar annular-shaped occluder, rendered to appear 10 cm above the tabletop. At this height, the virtual finger passed behind the occluder on normal movements. The center of the target served as the center for the arc defining the annular occluder. The occluder was 6 cm wide with the outer boundary positioned 6 cm from the start position and the inner boundary positioned 12 cm from the start position. The configuration insured that, regardless of the start and target positions, the finger always emerged from behind the occluder when it was 18 cm away from the target and at approximately the same time during the movement.

An Optotrak 3020 system recorded the time-varying position of a subject's finger at 240 Hz. The data were used to dynamically update the position of the virtual fingertip. Subjects wore a finger splint on their right index finger, which had a triad of active infrared markers. The position of a subject's fingertip within the splint was computed from the position of the three markers attached to the splint. The optical measurements had very small latencies (< 2 ms), but the speed of the graphical rendering loop imposed a significant additional latency on visual feedback, ~ 25 ms. We measured this by rendering a small red spot in the middle of the monitor and alternating the color between dark and light every four times through the graphics loop. A light-sensitive diode measured the brightness of the dot. At the time that we refreshed the video buffer (at the end of the display loop, timed to be just before the drawing of a new frame), we sent a signal to the serial port that covaried with the brightness of the spot. We measured the phase shift between the signal from the light-sensitive diode and the serial port on an oscilloscope. This provided an estimate of the delay between software draws and the appearance of a stimulus on the midline of the screen. The shift was consistently 25 ms (one and one-half video frames), reflecting a one video frame delay in the monitor (the extra one-half frame results from the time to scan to the midpoint of the monitor). Finger position on the screen varied over time and between trials, causing some variance in the display delay. We estimated the mean and SD in the display delay from subjects' finger positions. The SD computed across all subjects in experiment 1 was 3.2 ms, with no subject having a SD > 3.7 ms.

When computing the rendered position of the virtual fingertip, we compensated for the delay by linearly extrapolating the "current" position and orientation of the finger from the latest marker positions, using the positions from previous frames to estimate velocity. Linear extrapolation contributed a small bias in the x position of the virtual finger because of acceleration of the hand, varying between approximately -0.3

cm at the point of peak acceleration to 0.3 cm at peak deceleration. Biases in the y direction caused by extrapolation were comparatively small, ranging between approximately ± 0.02 cm. Extrapolation also had the effect of amplifying the intrinsic noise in Optotrak measurements; however, even after amplification, this noise remained small (root mean square error, 0.016 cm after extrapolation).

Spatial calibration of the virtual environment required computing the coordinate transformation from the reference frame of the Optotrak to the reference frame of the computer monitor, and the location of a subject's eyes relative to the monitor. These parameters were measured at the start of each experimental session using an optical matching procedure. The backing of the half-silvered mirror was temporarily removed, so that subjects could see their hand and the monitor simultaneously, and subjects aligned an Optotrak marker to a sequence of visually cued locations. Cues were presented monocularly, and matches were performed in separate sequences for left and right eyes. Thirteen positions on the monitor were cued, and each position was matched twice in different depth planes. The combined responses for both eyes were used to determine a globally optimal combination of three-dimensional reference frame and eye position. After the calibration procedure, a rough test was performed in which subjects moved a marker viewed through the half-silvered mirror and checked that the position of a rendered dot was perceptually coincident with the marker. Calibration was deemed acceptable if deviations appeared < 1 – 2 mm. Otherwise, the calibration procedure was repeated.

The transformation from the Optotrak measurements of marker position on the finger to the fingertip was computed in two steps. First, a subject placed her finger in a metal disk resting on the tabletop with a slot cut into it to hold the finger. This assured that, during the initial step of calibration, a subject's finger was rigid and parallel to the tabletop. The virtual finger was displayed so as to appear at the center of the virtual screen, oriented parallel to the tabletop pointing "up" in screen coordinates and positioned at a height above the virtual screen (and hence the tabletop) matching the height of the finger slot. The subject positioned his or her finger so as to visually align it with the rendered finger and hit a mouse button when satisfied that the two overlapped. This allowed us to compute the transformation from the three-dimensional coordinate frame defined by the three infrared markers on the finger flag and the rendered finger. To calibrate the position of a subject's fingertip (the part of the finger that a subject actually makes contact with the table when touching a target), we had subjects touch 10 small square buttons displayed at random positions on the tabletop at what they perceived to be the centers of the squares. We then calculated the average vector (in a finger-centered coordinate frame) between the center of the finger coordinate frame (the center of the three markers) and the centers of the squares. This represented the position of the fingertip for purposes of data analysis.

Procedure. At the beginning of each trial, the display refreshed, and the start cross, the target, and the annular occluder appeared in the virtual workspace. The subject was instructed to touch the center of the cross with the tip of his or her index finger and to hold that position until hearing a beep. At 500 ms after detecting that a subject's finger was touching the start cross, the system emitted a beep signaling subjects to point to touch the target. If the fingertip was within 2.5 mm of the boundary of the visible target region when it touched the tabletop, the trial was recorded as a success and the target visibly exploded; otherwise, nothing happened (the 2.5 mm "slop" was used to accommodate a successful touch anywhere on the finger pad). Two seconds after the beginning of the trial (the time the subject touched the start cross), the display was cleared and a 300 ms intertrial interval ensued before the display for the next trial was presented. Pointing movements were constrained to be > 600 ms, and the entire duration of each trial, from the time at which the subject touched the start position to the time they completed the pointing movement, was constrained to be < 2 s. If a subject moved before the beep or completed the movement outside the time constraints, an appropriate message was displayed and the trial condition was randomly placed into the stack of remaining trials and a new trial was begun. Subjects found movements within the allotted time window quite natural.

Before the main experiment blocks, subjects were allowed 20–40 practice trials to familiarize themselves with the task and the timing constraints. Subjects were instructed that touching the target buttons anywhere along their entire extent would lead to a success (blowing up the target). Importantly, subjects were instructed during the practice trials to “test” their calibration by touching the target buttons in different locations to make sure they exploded no matter where they touched. This was used to reinforce the constraint that success derived from touching a button anywhere along its entire extent.

Subjects participated in four experimental sessions on separate days. Each began with calibration of the virtual environment, followed by practice trials to familiarize themselves with the task, and then six blocks of experimental trials separated by a brief break. Subjects performed 72 trials in each experimental block for a total of 432 trials per session. The experiment contained 15 different conditions, corresponding to three different target button shapes—square, vertical rectangles, and horizontal rectangles—and five different finger perturbation conditions—1 cm perturbation forward along the start-target axis, 1 cm back along the axis, 1 cm up perpendicular to the axis, 1 cm down perpendicular to the axis, and no perturbation. The no-perturbation condition served as a baseline condition for analyzing movement kinematics. Forty-eight of the 72 trials (16 per target shape) were no-perturbation conditions and 24 were perturbation conditions (2 per target shape/perturbation direction/perturbation sign).

Data analysis. We analyzed three features of subjects’ movements—the endpoints of movements on baseline trials, the endpoints of movements on perturbation trials, and the time course of online corrections. To measure changes in subjects’ planned movements for different target shapes, we calculated the average and covariance (in table coordinates) of endpoint fingertip positions for no-perturbation trials in each of the different target shape conditions. To obtain a summary measure of the contribution of visual feedback to the pointing movements, we computed the average corrections for positive and negative perturbations for perturbations along the start-target axis and perturbations perpendicular to the start-target axis. This was given by the following:

$$\Delta \mathbf{x}_{\text{vert}} = \frac{1}{2} (\mathbf{x}_{\text{vert}}^- - \mathbf{x}_{\text{vert}}^+) \quad (1)$$

and

$$\Delta \mathbf{x}_{\text{horiz}} = \frac{1}{2} (\mathbf{x}_{\text{horiz}}^- - \mathbf{x}_{\text{horiz}}^+), \quad (2)$$

where $\mathbf{x}_{\text{vert}}^-$ and $\mathbf{x}_{\text{vert}}^+$ represent the average endpoints for negative and positive perturbations perpendicular to the start-target axis and $\mathbf{x}_{\text{horiz}}^-$ and $\mathbf{x}_{\text{horiz}}^+$ represent the average endpoints for negative and positive perturbations along the start-target axis. The vectors are represented in a coordinate frame aligned with the start-target axis, so a positive vertical perturbation is represented by the vector $[0, 1]^T$ and a negative vertical perturbation by the vector $[0, -1]^T$. Corresponding vectors for horizontal perturbations are $[1, 0]^T$ and $[-1, 0]^T$. The order of the terms in the difference equation derives from the fact that the proper correction for a perturbation of the visual finger is in the opposite direction of the perturbation.

To derive a sensitive measure of the temporal evolution of subjects’ corrective responses, we fit the following linear model to the movement kinematics at each time point (Saunders and Knill, 2002) as follows:

$$x(t) = w_1(t)x(t-1) + w_2(t)x(t-2) + \dots + w_n(t)x(t-n) - w_{\text{pert}}(t)\Delta_{\text{pert}}, \quad (3)$$

$$y(t) = v_1(t)y(t-1) + v_2(t)y(t-2) + \dots + v_n(t)y(t-n) - \gamma_{\text{pert}}(t)\Delta_{\text{pert}}, \quad (4)$$

where $x(t)$ is the position of the finger along the start-target axis at time t , $y(t)$ is the position of the finger perpendicular to the start-target axis at time t , and Δ_{pert} is the perturbation in the virtual finger (± 1 cm on perturbation trials and 0 on no-perturbation trials). The weights $w_i(t)$

and $v_i(t)$ capture the strong correlations between fingertip positions at successive points in time, as reflected in the smoothness of finger trajectories. The weights $w_{\text{pert}}(t)$ and $v_{\text{pert}}(t)$ represent the residual trajectory data (after accounting for autocorrelations) that can be attributed to corrections for the perturbations. We refer to these as perturbation influence functions; they show the temporal evolution of subjects’ corrections to perturbations. Since $w_{\text{pert}}(t)$ and $v_{\text{pert}}(t)$ are subtracted from the finger coordinates positions predicted by the previous history of positions, they should be positive during epochs in which subjects generate corrections for the perturbations. We set $n = 6$ for the analysis but found that the results were insensitive to the exact value for $n > 6$.

To group each subject’s data for analysis, their trajectories were lined up at the time at which the finger first reappeared from behind the occluder. We recorded the rendered position of the fingertip at the time of the call to the software draw—this was the fingertip extrapolated in time by 25 ms to account for the delay in the draw cycle rather than the true fingertip position at that time. Since the fingertip appears on the monitor at this position ~ 25 ms later, we used the time in software at which the rendered fingertip position passed the edge of the occluder plus 25 ms as the time the fingertip reappeared to subjects. We refer to this as the perturbation time, because it is the first time at which the visual effects of a perturbation appeared to an observer. The time variable t in the regression Equations 3 and 4 was set to 0 at the perturbation time.

We smoothed the influence functions using a causal exponential filter with a time constant of four frames (33.33 ms). The time constant was chosen using a random subsampling cross-validation technique to find the smoothing kernel that best fit the data. For each subject and condition, we repeatedly and randomly split trials into two equal-sized subsamples (training and validation). We fit perturbation influence functions to the training subsamples, smoothed them by an exponential filter with a specified time constant, and measured the mean-squared error between the smoothed influence functions and the training subsamples. The time constant that gave the smallest mean-squared error for predicting the validation subsample (averaged across random splits of the data) was chosen as the best-fitting filter constant. This varied between subjects and conditions. A time constant of four frames was the smallest best-fitting parameter across all subjects and conditions.

Experiment 2

Experiment 2 was equivalent to experiment 1 in every way with two exceptions. First, the square buttons were replaced by crosses. The crosses were 6 cm tall by 6 cm wide, with the arms of the crosses 1 cm wide (equivalent to overlaying the horizontal and vertical rectangles used in experiment 1). Second, perturbation trials included ± 1 cm shifts in the position of the virtual finger in three rather than two directions—vertical, horizontal, and diagonal (45° from the axis between the start position and the target). This resulted in a total of 84 trials per block—48 baseline, unperturbed trials and 36 perturbed trials (2 per target shape/perturbation direction/perturbation sign) and 504 trials per session.

The model

Rather than modeling a full, three-joint arm moving in three dimensions, the optimization of which would be very difficult if not intractable, we simulated a simplified model of a two-joint arm performing a planar pointing task that was in all other respects the same as the task performed by subjects. The model captures the principal nonlinearities of the arm as they affect the pointing task, in particular, the mapping of torques in joint space to movements in Euclidean space (in which the accuracy demands are specified). Figure 2 shows the arm model. The arm is modeled as a second-order nonlinear system (Hollerbach and Flash, 1982) as follows:

$$\tau = M(\theta)\ddot{\theta} + C(\theta, \dot{\theta}) + B\dot{\theta}, \quad (5)$$

where τ is a two-dimensional vector of joint torques (shoulder and elbow), θ , $\dot{\theta}$, and $\ddot{\theta}$ are two-dimensional vectors of joint angles, angular velocities, and angular accelerations, respectively, $M(\theta)$ is an inertial matrix, $C(\theta, \dot{\theta})$ is a vector of coriolis forces, and B is a joint friction matrix. Using these equations, the forward dynamics are given by the following:

$$\ddot{\theta} = M(\theta)^{-1}(\tau - C(\theta, \dot{\theta}) - B\dot{\theta}); \quad (6)$$




Figure 2. A schematic view of the two-joint arm model used for the simulations.

M , C , and B are given by the following:

$$M(\theta) = \begin{pmatrix} a_1 + 2a_2 \cos\theta_2 & a_3 + a_2 \cos\theta_2 \\ a_3 + a_2 \cos\theta_2 & a_3 \end{pmatrix} \quad (7)$$

$$C(\theta, \dot{\theta}) = \begin{pmatrix} -\dot{\theta}_2(2\theta_1 + \dot{\theta}_2) \\ \dot{\theta}_1^2 \end{pmatrix} a_2 \sin\theta_2 \quad (8)$$

$$B(\theta) = \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \quad (9)$$

$$a_1 = I_1 + I_2 + m_2 l_1^2 \quad (10)$$

$$a_2 = m_2 l_1 s_2 \quad (11)$$

$$a_3 = I_2,$$

where $b_{11} = b_{22} = 0.05 \text{ kgm}^2 \text{ s}^{-1}$, $b_{12} = b_{21} = 0.025 \text{ kgm}^2 \text{ s}^{-1}$, m_1 and m_2 are the masses of the two links in the arm model (the upper arm and the forearm), set to 1.4 and 1 kg, respectively, l_i is the length of link i (30, 33 cm), s_i is the distance from the center of mass for link i to the joint center for the link (11, 16 cm), and I_i is the moment of inertia for link i (0.025, 0.045 kgm^2). These are the same model parameters used by Li and Todorov (2004).

We discretized the forward dynamics given in Equation 6 to obtain the following:

$$\dot{\theta}_{t+1} = \dot{\theta}_t + \Delta M(\theta_t)^{-1} (\tau_t - C(\theta_t, \dot{\theta}_t) - B\dot{\theta}_t), \quad (12)$$

where Δ is the time step for the discretization (10 ms in our simulations).

The position, \mathbf{x}_t , and velocity, \mathbf{v}_t , of the finger are given by the following:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \Delta \mathbf{v}_t \quad (13)$$

$$\mathbf{v}_t = g(\theta_t, \dot{\theta}_t), \quad (14)$$

where the mapping from joint space to velocities in Euclidean coordinates is given by the following:

$$\mathbf{v}_t = g(\theta_t, \dot{\theta}_t) = \begin{pmatrix} -l_1 \dot{\theta}_1 \sin\theta_1 - l_2 (\dot{\theta}_1 + \dot{\theta}_2) \sin(\theta_1 + \theta_2) \\ l_1 \dot{\theta}_1 \cos\theta_1 + l_2 (\dot{\theta}_1 + \dot{\theta}_2) \cos(\theta_1 + \theta_2) \end{pmatrix}. \quad (15)$$

Joint angles and velocities can be expressed as functions of position and velocity in Euclidean space, $\theta_t = h(\mathbf{x}_t)$, $\dot{\theta}_t = q(\mathbf{x}_t, \mathbf{v}_t)$. Inserting

these into Equations 12–14 gives a nonlinear state update equation as follows:

$$\begin{bmatrix} \mathbf{x}_{t+1} \\ \mathbf{v}_{t+1} \end{bmatrix} = \begin{bmatrix} I_{2 \times 2} & \Delta \\ 0 & 0_{2 \times 2} \end{bmatrix} \begin{bmatrix} \mathbf{x}_t \\ \mathbf{v}_t \end{bmatrix} + \begin{bmatrix} 0 \\ f(\mathbf{x}_t, \mathbf{v}_t, \tau_t) \end{bmatrix}, \quad (16)$$

where τ_t is a two-dimensional vector of torques.

The control signal that generates torques at the joints was modeled as the sum of planned and feedback components,

$$\mathbf{u}_t = \mathbf{I}_t + L_t [\hat{\mathbf{x}}_t, \hat{\mathbf{v}}_t]^T, \quad (17)$$

where $[\hat{\mathbf{x}}_t, \hat{\mathbf{v}}_t]^T$ is the estimate of the system of the position and velocity of the finger, and L_t is a feedback gain matrix. To simulate motor delays and low-pass filtering of the motor signal, we augmented the state vector to incorporate a one-step delay (10 ms) and to implement a second-order recursive filter using a set of coupled update equations as follows:

$$\mathbf{u}_{t+1}^{(1)} = \mathbf{u}_t + \phi_t, \quad (18)$$

$$\mathbf{u}_{t+1}^{(2)} = \mathbf{u}_t^{(1)}, \quad (19)$$

$$\tau_{t+1}^{(1)} = \left(1 - \frac{\Delta}{k_{\text{motor}}}\right) \tau_t^{(1)} + \frac{\Delta}{k_{\text{motor}}} \mathbf{u}_t^{(2)}, \quad (20)$$

$$\tau_{t+1} = \left(1 - \frac{\Delta}{k_{\text{motor}}}\right) \tau_t + \frac{\Delta}{k_{\text{motor}}} \tau_t^{(1)}, \quad (21)$$

where ϕ_t is a signal-dependent Gaussian noise source, k_{motor} is the time constant of the low-pass filter (set to 40 ms), and τ_t is the torque applied to the joints.

We similarly modeled the visual sensory signal as a low-pass-filtered copy of the state of the hand, delayed by two time steps (20 ms) and corrupted by additive, state-dependent Gaussian noise. The update equations for the augmented portion of the state matrix were given by the following:

$$\begin{bmatrix} \mathbf{x}_{t+1}^{(1)} \\ \mathbf{v}_{t+1}^{(1)} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_t \\ \mathbf{v}_t \end{bmatrix} + \omega_t, \quad (22)$$

$$\begin{bmatrix} \mathbf{x}_{t+1}^{(2)} \\ \mathbf{v}_{t+1}^{(2)} \end{bmatrix} = \begin{bmatrix} \mathbf{x}_t \\ \mathbf{v}_t \end{bmatrix}, \quad (23)$$

$$\begin{bmatrix} \mathbf{x}_{t+1}^{(3)} \\ \mathbf{v}_{t+1}^{(3)} \end{bmatrix} = \left(1 - \frac{\Delta}{k_{\text{sense}}}\right) \begin{bmatrix} \mathbf{x}_t^{(3)} \\ \mathbf{v}_t^{(3)} \end{bmatrix} + \frac{\Delta}{k_{\text{sense}}} \begin{bmatrix} \mathbf{x}_t^{(2)} \\ \mathbf{v}_t^{(2)} \end{bmatrix}, \quad (24)$$

$$\begin{bmatrix} \mathbf{x}_{t+1}^{(s)} \\ \mathbf{v}_{t+1}^{(s)} \end{bmatrix} = \left(1 - \frac{\Delta}{k_{\text{sense}}}\right) \begin{bmatrix} \mathbf{x}_t^{(s)} \\ \mathbf{v}_t^{(s)} \end{bmatrix} + \frac{\Delta}{k_{\text{sense}}} \begin{bmatrix} \mathbf{x}_t^{(3)} \\ \mathbf{v}_t^{(3)} \end{bmatrix}, \quad (25)$$

where ω_t is a state-dependent Gaussian noise source, and k_{sense} is the time constant of the low-pass filter (set to 40 ms). Because of the discretization of the dynamics, the model had an effective sensorimotor delay of eight time steps (80 ms). The motor and sensory delays were chosen based on simulations of the model to match the delayed temporal response of human subjects (for details, see Results, Modeling).

The augmented system is represented by the state vector, $\mathbf{X}_t = [\mathbf{x}_t, \mathbf{v}_t, \mathbf{u}_t^{(1)}, \mathbf{u}_t^{(2)}, \tau_t^{(1)}, \tau_t, \mathbf{x}_t^{(1)}, \mathbf{v}_t^{(1)}, \mathbf{x}_t^{(2)}, \mathbf{v}_t^{(2)}, \mathbf{x}_t^{(3)}, \mathbf{v}_t^{(3)}, \mathbf{x}_t^{(s)}, \mathbf{v}_t^{(s)}]^T$. The nonlinear state update equation is given by the following:

$$\mathbf{X}_{t+1} = \mathbf{A} \mathbf{X}_t + [0_2, f(\mathbf{x}_t, \mathbf{v}_t, \tau_t), 0_{24}]^T + \mathbf{I}_t + L_t (\hat{\mathbf{X}}_t - \bar{\mathbf{X}}) + \Phi_t, \quad (26)$$

where $f(\mathbf{x}_t, \mathbf{v}_t, \tau_t)$ is the nonlinear function mapping the current position and velocity of the finger and the joint torques to the velocity at the next time step. The state noise is given by a vector containing both motor and sensory noise terms, $\Phi_t = [0_4, \phi_t, 0_6, \omega_t, 0_{12}]^T$ ($\mathbf{0}_n$ is an n -dimensional vector of zeroes). Because of the augmentation of the state vector, L_t is a 28-dimensional vector that is zero everywhere except in the fifth and sixth elements, and the feedback gain matrix L_t is a 28×28 matrix that is

zero everywhere except for the 2×4 submatrix $L_{5:6,1:4}$. The matrix A is given by the following:

$$A = \begin{bmatrix} D & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & M & \mathbf{0} \\ S_1 & \mathbf{0} & S_2 \end{bmatrix}, \quad (27)$$

where D is the matrix form of Equation 16 and is given by the following:

$$D = \begin{bmatrix} I_{2 \times 2} & \Delta I_{2 \times 2} \\ 0_{2 \times 2} & 0_{2 \times 2} \end{bmatrix}, \quad (28)$$

where $I_{n \times n}$ is the $n \times n$ identity matrix and $0_{n \times n}$ is an $n \times n$ matrix of zeros. M captures the motor delay and filtering and is given by the following:

$$M = \begin{bmatrix} 0_{2 \times 2} & 0_{2 \times 2} & 0_{2 \times 2} & 0_{2 \times 2} \\ I_{2 \times 2} & 0_{2 \times 2} & 0_{2 \times 2} & 0_{2 \times 2} \\ 0_{2 \times 2} & \frac{\Delta}{k_{\text{motor}}} I_{2 \times 2} & \left(1 - \frac{\Delta}{k_{\text{motor}}}\right) I_{2 \times 2} & 0_{2 \times 2} \\ 0_{2 \times 2} & 0_{2 \times 2} & \frac{\Delta}{k_{\text{motor}}} I_{2 \times 2} & \left(1 - \frac{\Delta}{k_{\text{motor}}}\right) I_{2 \times 2} \end{bmatrix}, \quad (29)$$

S_1 and S_2 implement the fixed sensory delay and the sensory filter, respectively, and are given by the following:

$$S_1 = \begin{bmatrix} I_{4 \times 4} \\ 0_{4 \times 4} \\ 0_{4 \times 4} \\ 0_{4 \times 4} \end{bmatrix},$$

$$S_2 = \begin{bmatrix} 0_{4 \times 4} & 0_{4 \times 4} & 0_{4 \times 4} & 0_{4 \times 4} \\ I_{4 \times 4} & 0_{4 \times 4} & 0_{4 \times 4} & 0_{4 \times 4} \\ 0_{4 \times 4} & \frac{\Delta}{k_{\text{sense}}} I_{4 \times 4} & \left(1 - \frac{\Delta}{k_{\text{sense}}}\right) I_{4 \times 4} & 0_{4 \times 4} \\ 0_{4 \times 4} & 0_{4 \times 4} & \frac{\Delta}{k_{\text{sense}}} I_{4 \times 4} & \left(1 - \frac{\Delta}{k_{\text{sense}}}\right) I_{4 \times 4} \end{bmatrix}. \quad (30)$$

The observation model simply peels off the last four elements of the augmented state vector (the low-pass-filtered, noisy visual estimates of position and velocity),

$$\mathbf{y}_t = H\mathbf{X}_t, \quad (31)$$

where H is a 4×28 matrix as follows:

$$H = [0_{4 \times 24} \quad I_{4 \times 4}]. \quad (32)$$

No noise is added to \mathbf{y}_t because the sensory noise has been incorporated into the augmented state update model.

For a given control law, one can compute the mean trajectory using the nonlinear update equation (Eq. 26) and linearize the system dynamics around the mean trajectory to derive a Kalman filter to estimate the state of the hand. The linearized dynamics are given by the following:

$$\hat{\mathbf{X}}_{t+1} = \bar{\mathbf{X}}_{t+1} + J_t(\hat{\mathbf{X}}_t - \bar{\mathbf{X}}_t) + \mathbf{1}_t + L_t(\hat{\mathbf{X}}_t - \bar{\mathbf{X}}_t) + \Phi_t \quad (33)$$

$$\mathbf{y}_t = H\mathbf{X}_t, \quad (34)$$

where the state transition matrix J_t is given by the following:

$$J_t = \begin{bmatrix} D_1 & D_2 & \mathbf{0} \\ \mathbf{0} & M & \mathbf{0} \\ S_1 & \mathbf{0} & S_2 \end{bmatrix}. \quad (35)$$

D_1 and D_2 represent the linearized dynamics and are given by the following:

$$D_1 = \begin{bmatrix} I_{2 \times 2} & \Delta I_{2 \times 2} \\ \Delta \frac{\partial f}{\partial \mathbf{x}_t} \Big|_{\mathbf{x}_t = \bar{\mathbf{x}}_t} & \Delta \frac{\partial f}{\partial \mathbf{v}_t} \Big|_{\mathbf{v}_t = \bar{\mathbf{v}}_t} \end{bmatrix}, D_2 = \begin{bmatrix} 0_{2 \times 6} & 0_{2 \times 2} \\ 0_{2 \times 6} & \Delta \frac{\partial f}{\partial \boldsymbol{\tau}_t} \Big|_{\boldsymbol{\tau}_t = \bar{\boldsymbol{\tau}}_t} \end{bmatrix}, \quad (36)$$

where $\frac{\partial f}{\partial \mathbf{x}_t}$ is the Jacobian of $f()$ with respect to the position of the finger, $\frac{\partial f}{\partial \mathbf{v}_t}$ is the Jacobian of $f()$ with respect to the velocity of the finger, and $\frac{\partial f}{\partial \boldsymbol{\tau}_t}$ is the Jacobian of $f()$ with respect to the torque applied to the joints, all evaluated at the mean values for the relevant parameters at time t .

The optimal estimate of the finger state is given by the following:

$$\hat{\mathbf{X}}_{t+1} = \bar{\mathbf{X}}_{t+1} + J_t(\hat{\mathbf{X}}_t - \bar{\mathbf{X}}_t) + \mathbf{1}_t + L_t(\hat{\mathbf{X}}_t - \bar{\mathbf{X}}_t) + K_t(\mathbf{y}_t - H\hat{\mathbf{X}}_t), \quad (37)$$

where K_t is the Kalman gain. This is given by the following:

$$K_t = J_t \Sigma_t^e H^T (H \Sigma_t^e H^T)^{-1}, \quad (38)$$

where Σ_t^e is the error covariance of the estimator ($\Sigma_t^e = E[(\hat{\mathbf{X}}_t - \mathbf{X}_t)(\hat{\mathbf{X}}_t - \mathbf{X}_t)^T]$). This is given by the following:

$$\Sigma_{t+1}^e = (J_t - K_t H) \Sigma_t^e (J_t - K_t H)^T + \Sigma_t^\Phi, \quad (39)$$

where Σ_t^Φ represents the covariance of the system noise, which is a composite of the motor noise and the sensory noise,

$$\Sigma_t^\Phi = \begin{bmatrix} 0_{4 \times 4} & 0_{4 \times 2} & 0_{4 \times 6} & 0_{4 \times 4} & 0_{4 \times 12} \\ 0_{2 \times 4} & \Sigma_t^\phi & 0_{2 \times 6} & 0_{2 \times 6} & 0_{2 \times 12} \\ 0_{6 \times 4} & 0_{6 \times 2} & 0_{6 \times 6} & 0_{6 \times 4} & 0_{6 \times 12} \\ 0_{4 \times 4} & 0_{4 \times 2} & 0_{4 \times 6} & \Sigma_t^\phi & 0_{4 \times 12} \\ 0_{12 \times 12} & 0_{12 \times 2} & 0_{12 \times 6} & 0_{12 \times 4} & 0_{12 \times 12} \end{bmatrix}, \quad (40)$$

The motor noise was assumed to be proportional to the control signal, so the covariance is given by the following:

$$\Sigma_t^\phi = \sum_{i=1}^2 C_i (I_i I_i^T + L_i \bar{\mathbf{X}}_i \bar{\mathbf{X}}_i^T L_i^T + L_i \Sigma_{i,t}^x L_i^T) C_i^T, \quad (41)$$

where

$$C_1 = \begin{bmatrix} 0_{4 \times 4} & 0_{4 \times 1} & 0_{4 \times 1} & 0_{4 \times 22} \\ 0_{1 \times 4} & c & 0 & 0_{1 \times 22} \\ 0_{1 \times 4} & 0 & 0 & 0_{1 \times 22} \\ 0_{22 \times 4} & 0_{22 \times 1} & 0_{22 \times 1} & 0_{22 \times 22} \end{bmatrix},$$

$$C_2 = \begin{bmatrix} 0_{4 \times 4} & 0_{4 \times 1} & 0_{4 \times 1} & 0_{4 \times 22} \\ 0_{1 \times 4} & 0 & 0 & 0_{1 \times 22} \\ 0_{1 \times 4} & 0 & c & 0_{1 \times 22} \\ 0_{22 \times 4} & 0_{22 \times 1} & 0_{22 \times 1} & 0_{22 \times 22} \end{bmatrix}. \quad (42)$$

We set $c = 0.06$ based on a coarse search to find a value that gave endpoint errors for the optimal controllers with approximately the same scatter as those of subjects.

We used published psychophysical data on spatial localization and motion acuity to parameterize the sensory noise model. Positional acuity is inversely proportional to eccentricity; thus, it is well modeled by a noise source with SD proportional to the radial position of the hand in retinal coordinates (Levi et al., 1988; Burbeck and Yap, 1990; Whitaker and Latham, 1997). Similarly, motion acuity, in both speed and direction, varies with target speed. Motion discrimination thresholds are well fit by a model in which the velocity components in both the direction of motion and the perpendicular direction are corrupted by a mixture of proportional noise, whose SD is proportional to the speed of the motion, and a constant noise component (Orban et al., 1985; De Bruyn and Orban, 1988).

For a given control law, we approximated the effects of state-dependent noise by using the average fingertip position and motion at each time to calculate the sensory noise covariance. We used results from two-point interval discrimination studies to set the parameters for visual

noise on position estimates (Burbeck, 1987; Burbeck and Yap, 1990; Whitaker and Latham, 1997). The data from these studies are consistent with a Weber fraction of 0.05 on position estimates beyond several degrees away from the fovea. This value is invariant to a large number of properties of the target (Burbeck, 1987; Toet and Koenderink, 1988). The resulting SDs (in meters) are given by the following:

$$\sigma_t^x = \sigma_t^y = (0.0005 + 0.05 \sqrt{x_t^2 + y_t^2}). \quad (43)$$

The constant additive term models a minimum SD in position estimates of 3' arc in the center of the fovea, which, for the viewing distance used in the experiment equates to ~ 0.5 mm (we used a small angle approximation in treating position in tabletop coordinates as proportional to visual angle). We multiplied the SD in Equation 43 by a factor of $0.25/\Delta t$ so that it gives the correct SD for an optimal estimator viewing a stimulus for 250 ms, the display time used in the experiments measuring localization acuity (this renders the noise model invariant to the particular time discretization used).

Results from speed and direction discrimination studies show a somewhat more complicated behavior than position perception. Up to speeds of 64°/s (close to the peak velocity measured in our experiments), Weber fractions for speed decrease to a minimum of 0.08 for viewing durations of 500 ms (Mateeff et al., 2000). These results are consistent across a number of studies and types of stimuli (Orban et al., 1985; De Bruyn and Orban, 1988). Subjects' threshold curves are well fit by a mixed constant and proportional noise model in which the SD of visual estimates of speed is given by the following:

$$\sigma_t^{v^{\parallel}} = 2^\circ/s + 0.08v_t^{\parallel}. \quad (44)$$

Using a small angle approximation to convert this to units of distance along the tabletop (assuming an average viewing distance of 52 cm) gives the following:

$$\sigma_t^{v^{\parallel}} = 1.8 \text{ cm/s} + 0.08v_t^{\parallel}. \quad (45)$$

Direction discrimination thresholds behave in a qualitatively similar fashion to speed discrimination thresholds, but when converted into units of speed in a direction perpendicular to the path of motion, thresholds are lower by more than a factor of 8. For the SD of velocity estimates perpendicular to the direction of motion, therefore, we have the following:

$$\sigma_t^{v^{\perp}} = 0.25^\circ/s + 0.01v_t^{\perp}. \quad (46)$$

Converting this to tabletop coordinates and adjusting for perspective foreshortening, we have for the SD the following:

$$\sigma_t^{v^{\perp}} = 0.35 \text{ cm/s} + 0.014v_t^{\perp}. \quad (47)$$

We scaled the parameters by a constant factor of $0.50/\Delta t$ so that an optimal estimator viewing a constant velocity stimulus for 500 ms would give velocity estimates with SDs listed above. Using Equations 45 and 47, we can compute the covariance of velocity estimates in Euclidean coordinates as follows:

$$\Sigma_t^v = \begin{bmatrix} \cos\theta_t & -\sin\theta_t \\ \sin\theta_t & \cos\theta_t \end{bmatrix} \begin{bmatrix} (\sigma_t^{v^{\parallel}})^2 & 0 \\ 0 & (\sigma_t^{v^{\perp}})^2 \end{bmatrix} \begin{bmatrix} \cos\theta_t & \sin\theta_t \\ -\sin\theta_t & \cos\theta_t \end{bmatrix}, \quad (48)$$

where θ_t is the angle between the direction of finger motion at time t ($\theta_t = \tan^{-1} \frac{v_t^y}{v_t^x}$) and the x -axis. These expressions determine the sensory noise covariance $\Sigma_t^{\tilde{v}}$ for a given mean trajectory.

We used an iterative, conjugate gradient descent procedure to find the control law that minimized the performance cost as follows:

$$\Upsilon(\mathbf{l}_\bullet, \mathbf{L}_\bullet) = \pi_{\text{miss}} + \lambda_{\text{stop}} E[|\mathbf{v}_T|^2]. \quad (49)$$

π_{miss} is the probability that the fingertip will miss the target button (land outside the bounds of the button at the end of the movement – time T). $E[|\mathbf{v}_T|^2]$ is the expected squared velocity of the fingertip at the end of the movement. It captures a constraint to stop at the end of the

movement. Strictly speaking, subjects did not have to decrease their velocity just before hitting the target, since they pressed their fingers to the table; however, subjects' trajectories showed that their velocity in x and y dropped to near 0 just before making contact with the table. In the simulations reported here, we set the coefficient for the stopping velocity cost to 0.5.

We can rewrite Equation 49 in terms of the means and covariances of the finger position and velocity, \mathbf{x}_t and \mathbf{v}_t (the first four elements of the augmented state vector), as follows:

$$\Upsilon(\mathbf{l}_\bullet, \mathbf{L}_\bullet) = \lambda_{\text{miss}} \iint_{\text{target}} N(\bar{\mathbf{x}}_t, \Sigma_t^{\tilde{x}}) dx dy + \lambda_{\text{stop}} (|\bar{\mathbf{v}}_T|^2 + \Sigma_T^v), \quad (50)$$

where the means and covariances are implicitly functions of \mathbf{l}_\bullet and \mathbf{L}_\bullet . $N(\bar{\mathbf{x}}_t, \Sigma_t^{\tilde{x}})$ is the multivariate Gaussian with mean $\bar{\mathbf{x}}_t$ and covariance $\Sigma_t^{\tilde{x}}$. The first integral is simply the probability of the finger endpoint falling within the target region. This is the only term in the objective function that depends on target shape. The optimal control law is given by the following:

$$[\mathbf{l}_\bullet^{\text{opt}}, \mathbf{L}_\bullet^{\text{opt}}] = \arg \min[\Upsilon(\mathbf{l}_\bullet, \mathbf{L}_\bullet)]. \quad (51)$$

To calculate the expected cost of a given control law, $\{\mathbf{l}_\bullet, \mathbf{L}_\bullet\}$, one needs to calculate $\bar{\mathbf{x}}_T$, $\bar{\mathbf{v}}_T$, $\Sigma_T^{\tilde{x}}$, and Σ_T^v . To compute $\bar{\mathbf{x}}_T$, $\bar{\mathbf{v}}_T$, we ran the nonlinear state update equation in Equation 26 forward in time with the noise term, Φ_t , set to 0. In the noiseless case, the error term, $\tilde{\mathbf{X}}_t - \bar{\mathbf{X}}_t$, is 0, and the forward model becomes the following:

$$\bar{\mathbf{X}}_{t+1} = A\bar{\mathbf{X}}_t + [0_2, f(\bar{\mathbf{x}}_t, \bar{\mathbf{v}}_t, \bar{\tau}_t), 0_{24}]^T + \mathbf{I}_t, \quad (52)$$

with $\bar{\mathbf{x}}_0 = [0, 50]^T$, $\bar{\mathbf{v}}_0 = [0, 0]^T$, $\bar{\tau}_0 = [0, 0]^T$ (in centimeters). To simplify the computation of $f(\bullet)$, we converted finger position and velocity into joint space at each time step and used Equation 12 to update the joint angular velocities before converting back to finger velocity. The resulting mean trajectories were used to calculate the sensory noise covariance $\Sigma_t^{\tilde{v}}$ at each time step.

The full state covariance matrix, $\Sigma_t^{\tilde{\mathbf{x}}}$, was computed using the linearized form of the system update. The update equation is given by the following:

$$\begin{aligned} \Sigma_{t+1}^{\tilde{\mathbf{x}}} &= E[(\mathbf{X}_{t+1} - \bar{\mathbf{X}}_{t+1})(\mathbf{X}_{t+1} - \bar{\mathbf{X}}_{t+1})^T] \\ &= J_t \Sigma_t^{\tilde{\mathbf{x}}} J_t^T + L_t \Sigma_t^{\tilde{\mathbf{x}}} L_t^T - J_t (\Sigma_t^{\tilde{\mathbf{x}}} + \Sigma_t^{\tilde{\mathbf{x}}}) L_t^T - L_t (\Sigma_t^{\tilde{\mathbf{x}}} + \Sigma_t^{\tilde{\mathbf{x}}}) J_t^T + \Sigma_t^{\Phi}, \end{aligned} \quad (53)$$

where the covariance matrices for the state estimate and the estimation error are given by the following:

$$\begin{aligned} \Sigma_{t+1}^{\tilde{\mathbf{x}}} &= (J_t - L_t) \Sigma_t^{\tilde{\mathbf{x}}} (J_t - L_t)^T + K_t \Sigma_t^{\tilde{\mathbf{x}}} K_t^T + (J_t - L_t) \Sigma_t^{\tilde{\mathbf{x}}} K_t^T \\ &\quad + K_t \Sigma_t^{\tilde{\mathbf{x}}} (J_t - L_t)^T \end{aligned} \quad (54)$$

$$\Sigma_{t+1}^{\tilde{\mathbf{x}}e} = \Sigma_{t+1}^{\tilde{\mathbf{x}}} = (J_t - L_t) \Sigma_t^{\tilde{\mathbf{x}}} (J_t - K_t H)^T + K_t H \Sigma_t^{\tilde{\mathbf{x}}} (J_t - K_t H)^T \quad (55)$$

$$\Sigma_{t+1}^e = (J_t - K_t H) \Sigma_t^e (J_t - K_t H)^T + \Sigma_t^{\Phi}. \quad (56)$$

Iterating Equations 53–56 forward in time gives $\Sigma_T^{\tilde{\mathbf{x}}}$, from which we extracted $\Sigma_T^{\tilde{x}}$ and Σ_T^v (the top left 4×4 submatrix of $\Sigma_T^{\tilde{\mathbf{x}}}$). The expected cost associated with the control law was computed by inserting $\bar{\mathbf{x}}_T$, $\bar{\mathbf{v}}_T$, $\Sigma_T^{\tilde{x}}$, and Σ_T^v into Equation 50. For each condition of the experiment, we used a coordinate, conjugate gradient descent algorithm to solve for the optimal control law associated with the target button shape in which we iteratively performed conjugate gradient descent on \mathbf{l}_\bullet and \mathbf{L}_\bullet . We iterated the descent algorithm until the expected cost reached an asymptote (the expected cost remained constant to within a proportional tolerance of 0.00001 between the line search steps of the conjugate gradient descent algorithm).

In the simulations reported in this paper, we assumed that subjects' fixated the average endpoints of the trajectories. Finger position was expressed in retinal coordinates with the origin at the fixation point. We initialized the covariance matrices to the following:

$$\Sigma_t = \Sigma_t^e = \begin{bmatrix} 0.4^2 I_{2 \times 2} & 0_{2 \times 26} \\ 0_{26 \times 2} & 0_{26 \times 26} \end{bmatrix} \quad (57)$$

to reflect an initial 4 mm scatter in starting positions around the start point. The simulated start position of the finger was at $[0, 50]^T$ and the targets were assumed to be centered around $[-30, 50]^T$ (all units in centimeters)—expressed in shoulder coordinates. Note that these initial conditions effectively assume that the subject begins with an initial estimate of the finger at the start position, but with the true position of the finger scattered around the start with a SD of 4 mm.

We simulated the performance of the optimal controllers in both noiseless and noisy conditions (with sensory and motor noise set to the levels specified by the model parameters)—the former to compute average endpoints and feedback corrections in different experimental conditions and the latter to compute measures of the temporal responses of the controllers to perturbations comparable with those of subjects (i.e., perturbation influence functions). Data from the noisy simulations were also used to cross-check the covariance matrices computed using the linear approximation to the forward dynamics. In all conditions, the average endpoints, corrections, and endpoint covariances estimated from repeated runs of noisy simulations were almost exactly equal to those estimated by iterating the update equations of the linearized system forward in time.

To compute the average endpoints on unperturbed trials, we iterated the nonlinear state update equation in Equation 26 forward in time from the initial position with the noise term, Φ_t , set to 0, as described above for computing \bar{X}_t . To compute the average endpoints for perturbed trials, the model was run forward in time and the sensory noise term was set to $\omega_t = 0$ for $t < t_{\text{pert}}$ and to $\omega_t = [\pm 1, 0, 0, 0]^T$ or $\omega_t = [0, \pm 1, 0, 0]^T$ for $t \geq t_{\text{pert}}$, where t_{pert} is the time at which the finger first reappears from behind the occluder (in its perturbed position), as calculated from the mean trajectory. \bar{X}_t was computed using the Kalman update equation given in Equation 37 and the Kalman gains, K_t^{opt} , associated with the optimal control law.

Simulations of noisy movements were run by adding independent motor and sensory noise to the state at each time step drawn from zero-mean, normal distributions with noise covariances given by Σ_t^ϕ and Σ_t^ω computed from the mean positions and velocities of the model finger and running the Kalman filter and state update equations forward in time. Measurement noise in the Optotrak was simulated by adding independent, zero-mean Gaussian noise with a SD of 0.025 mm to the sample trajectories derived from the noisy simulations. The resulting “measured” trajectories were subjected to the same analysis as subjects' data to derive performance metrics. Since the means and covariances of movement endpoints derived from the noisy simulations were within <1% of those computed by propagating the system update equations in the noiseless simulations, the latter are reported in the modeling section. The trajectories generated in the noisy simulations were used to compute perturbation influence functions for the optimal controllers to compare the temporal responses of the controllers with those of subjects.

To model variability in movement duration, we computed optimal controllers for movement times ranging from 680 to 800 ms in steps of 20 ms. We simulated the performance of these controllers and report the average performance metrics (e.g., correction magnitudes) across the seven simulated movement times. For the noisy simulations, we generated 1000 sample trajectories per movement duration and target shape (200 each for the no-perturbation and the four perturbation conditions) and used these to estimate perturbation influence functions for the optimal controllers.

Results

Experiment 1

Subjects' mean reaction time from the start signal to the initiation of movement was 244.9 ± 12.7 (SE) ms. Movement times aver-

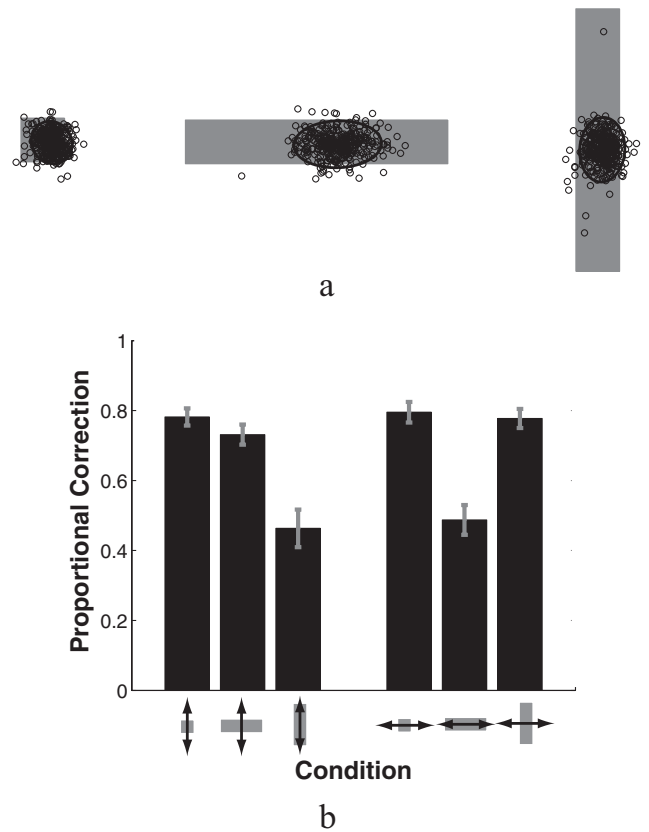


Figure 3. *a*, The scatter of endpoints on unperturbed trials for a representative subject in experiment 1 for each of the three targets. The coordinate frame is aligned with the axis between starting and target positions. Ellipses represent twice the SD of endpoint positions. *b*, Average magnitudes of corrections to 1 cm perturbations of the virtual finger for the three button shapes. Corrections shown are in the direction of the perturbation. Perturbations parallel to the axis between the start position and the target button are shown as horizontal arrows; perturbations perpendicular to that axis are shown as vertical arrows. Error bars indicate SEMs.

aged 733.1 ± 4.7 (SE) ms. The average SD of subjects' movement times was 63 ± 1.0 (SE) ms. The distributions of finger endpoints on unperturbed trials and subjects' corrective responses to perturbations in the position of the virtual finger both provide information about subjects' sensorimotor control strategies. Figure 3*a* shows scatterplots of a representative subject's fingertip endpoints on unperturbed trials for the three target buttons used in experiment 1. The patterns revealed in the scatterplots were consistent across all eight subjects. Table 1 summarizes the average distributions of endpoints across the eight subjects: the average of subjects' mean endpoint positions, the average of subjects' SDs of endpoints in the horizontal and vertical directions (defined relative to the start-target axis), and the average correlation coefficients between the x and y endpoint positions of the fingertip. As shown in Figure 3, the mean endpoints differed for the three target shapes. The mean endpoint was near the center for the square button, but shifted downward for vertical rectangles and to the right (toward the starting position) for horizontal rectangles. The scatter of endpoints also differed for the three different targets, being elongated in the long dimensions of the rectangular targets. Both of these effects appear in the average data, as shown in Table 1. The mean endpoint was 0.657 cm lower for the vertical rectangle than for the square ($T_{(7)} = 3.52$; $p = 0.01$) and was 1 cm closer to the starting position for the horizontal rectangle than for the square ($T_{(7)} = 4.97$; $p < 0.002$). The SD of endpoint scatter in

Table 1. Means and average covariances of subjects' endpoints on unperturbed trials for the three targets in experiment 1

	$(\mu_x, \mu_y) \pm SE$ (cm)	$(\sigma_x, \sigma_y) \pm SE$ (cm)	$\rho \pm SE$
Square	(0.167, -0.177) \pm (0.027, 0.076)	(0.309, 0.306) \pm (0.029, 0.031)	-0.180 \pm 0.040
Vertical	(0.070, - 0.834) \pm (0.042, 0.241)	(0.315, 0.611) \pm (0.024, 0.083)	- 0.026 \pm 0.017
Horizontal	1.167 , -0.206) \pm (0.205, 0.054)	(0.627, 0.305) \pm (0.034, 0.019)	- 0.011 \pm 0.014

SDs and correlation coefficients are averages across subjects. x is taken to be in the direction of this axis, and y is taken to be perpendicular to the axis. Significant differences between subjects' endpoint distributions for cross targets and subjects' endpoint distributions for vertical and horizontal buttons (means and SDs) are highlighted in bold.

the vertical dimension for the vertical rectangle was approximately twice that for the square ($T_{(7)} = 4.3$; $p < 0.005$). Similarly, the endpoint scatter in the horizontal dimension for the horizontal rectangle was approximately twice that for the square ($T_{(7)} = 8.9$; $p < 0.00005$). No other differences were significant; that is, mean endpoint positions and SDs were not significantly different between rectangles and squares along the dimensions in which the figures were the same size (e.g., the vertical dimension for horizontal rectangles and squares).

The difference in mean endpoints would seem to reflect a shape-dependent difference in the ballistic component of subjects' motor plans, that is, the points to which they effectively aimed. Moving to a nearer endpoint position for the horizontal rectangular target shortens the total movement distance and hence is energetically less costly, without causing significantly more misses. The lower endpoint for the vertical rectangular targets may result from a similar principle, as subjects' movement trajectories did curve toward the observer and then back up again (when projected onto the table), so that a lower endpoint may be less energetically costly. We will take up these points again in the modeling section, where we show that an optimal controller shows qualitatively similar behavior.

The difference in endpoint scatter could result from two causes. First, as hypothesized here, subjects may have applied less feedback control in dimensions in which endpoint errors are less costly. This would have led to increased variance in endpoints in those dimensions. Alternatively, they may have allowed more noise in endpoint planning, so the shaping of endpoint scatter does not unambiguously support the hypothesis that humans shape their feedback control laws to the spatial accuracy constraints of a task. The stronger prediction of the adaptive feedback control hypothesis is that subjects would have corrected less for perturbations in the visual position of the finger along the spatial dimension in which the rectangular targets are elongated.

Table 2 summarizes the full two-dimensional correction vectors for each of the perturbations and target shapes, as given by Equations 1 and 2. An easy way to visualize the vectors is as subjects' average correction vector to a negative perturbation (a vertical perturbation "down" or a horizontal perturbation to the "left" toward the target). Flipping the sign of the vector would give the average correction for a positive perturbation. As expected, by far the largest component of the correction vectors is in the direction of the perturbation. Figure 3*b* shows a summary plot of subjects' endpoint corrections in the direction of the perturbations for the three different target shapes.

The only significant differences between subjects' corrections for perturbations with the rectangular targets and those with the square targets are in the long directions of the rectangular targets. Subjects corrected less for vertical perturbations of the finger when the target is a vertical rectangular than when the target is square ($T_{(7)} = 6.98$; $p < 0.0005$). Corrections to horizontal perturbations for the two targets do not differ significantly. By contrast, subjects corrected less for horizontal perturbations of the finger when the target was a horizontal rectangle than when it was

Table 2. Mean corrections to perturbations for the three different target buttons

	Horizontal perturbations $(\mu_{\Delta x}, \mu_{\Delta y}) \pm SE$ (cm)	Vertical perturbations $(\mu_{\Delta x}, \mu_{\Delta y}) \pm SE$ (cm)
Square	(0.795, 0.099) \pm (0.029, 0.007)	(0.026, 0.782) \pm (0.015, 0.025)
Vertical	(0.777, 0.147) \pm (0.027, 0.026)	(0.009, 0.463) \pm (0.007, 0.054)
Horizontal	(0.487, 0.031) \pm (0.043, 0.012)	(0.056, 0.731) \pm (0.020, 0.029)

Horizontal perturbations refer to perturbations parallel to the axis between the starting position and the target position; vertical perturbations refer to perturbations perpendicular to this axis. x is taken to be in the direction of this axis, and y is taken to be perpendicular to the axis. Significant differences between the corrections for rectangular targets and the square targets are highlighted in bold.

square ($T_{(7)} = 6.76$; $p < 0.0005$). Corrections to vertical perturbations did not differ significantly. Similarly, subjects' corrections to perturbations in the direction of the long axis of the rectangles were significantly different from their corrections to perturbations in the direction of the short axis (for vertical rectangles, $T_{(7)} = 7.11$, $p < 0.0005$; and for horizontal rectangles, $T_{(7)} = 5.73$, $p < 0.001$). No other differences were significant. Treating these values as a measure of the overall feedback gain of the system, we find that subjects' reduced the feedback gain in the direction of lowered accuracy demands by $\sim 40\%$ relative to the feedback gain in the direction of higher accuracy demands (the short axis of the rectangles).

The results so far reflect the overall effect of subjects' changes in feedback control strategies as reflected by the endpoints of their movements. They do not indicate whether subjects adjusted their feedback control laws for different shapes throughout their movements or simply at the end, where the corrections were most prominent. Figure 4 shows example kinematic data for two subjects in the square target/horizontal perturbation condition. Figure 4, *a* and *b*, shows the two subjects' velocities measured in the principal direction of movement—parallel to the line between the start and target positions (the same direction as the perturbations)—for each of the no-perturbation, positive-perturbation, and negative-perturbation conditions. Little significant effect of the perturbations can be discerned from this data. Velocity profiles perpendicular to the main movement direction for vertical perturbation conditions appear to be more reliable indicators of subjects' responses to the perturbations (Fig. 5).

Using kinematic data like that shown in Figures 4 and 5 to analyze the temporal dynamics of subjects' responses to perturbations in the main direction of movement has a number of significant problems. First, even mean velocity profiles show a significant amount of scatter across conditions, because of the small numbers of trials per condition, the large variance across trials in individual subjects' movement kinematics, and significant differences in timing created by variance in movement duration and variance in when the start of a movement was triggered. These effects are particularly deleterious to the velocity profiles parallel to the path between the start and the target positions. Making the timing analysis more difficult is that perturbations were not initiated at fixed times relative to movement start, but rather became visible to subjects at a fixed distance from the target—the timing of which varied from trial to trial.

To avoid the problems involved in averaging subjects' kinematics over trials, we developed a more sensitive method for analyzing the temporal dynamics of subjects' corrective responses to perturbations in the visual stimulus (Saunders and Knill, 2004). The fundamental observation behind the method is that the strong temporal correlations in subject's movements allow one to linearly predict the position of the finger from previous positions. Because the predictive filter is relatively invariant to small shifts in time and small changes in temporal scale, one can derive an accurate predictive filter from a collection of movements that may be time shifted (e.g., because of variance in estimates of start time) or scaled (e.g., because of variance in overall speed of movements) in unknown ways relative to one another. Moreover, the predictive filter works well on movements with very different overall trajectories—it predicts positions to within the accuracy limits of the Optotrak. To analyze subjects' corrective responses to perturbations, we first temporally shifted a subject's trajectories so that they lined up in time at the time at which the finger first appears from behind the occluder. This ends up shifting trajectories by ± 3 Optotrak frames (25 ms SD). We then derived the best linear predictive filter from trials containing no perturbations. Beginning with the time at which subjects' fingers appear from behind the occluder, we applied the predictive filter to trajectories on perturbation trials and correlate the residual errors in the predictions at each time step with the perturbations on those trials. The result is what we refer to as a "perturbation influence function" (for details, see Materials and Methods).

Figure 6 shows perturbation influence functions calculated from individual subjects' kinematic data shown in Figures 4 and 5. Particularly for the horizontal perturbations, the influence functions provide a much more sensitive measure of the subjects' responses to the perturbations. Figure 7 shows functions for all six experimental conditions averaged across subjects. The shapes of the influence functions for rectangles and squares are remarkably similar for perturbations in the direction of the short axis of the rectangle, which is equal to the width of the square. They differ, however, for perturbations in the direction of the long axis of rectangles (Fig. 7*a*, red curve; *b*, green curve). This reflects the smaller corrections made to those perturbations. Visual inspection suggests that the differences between corrective responses to rectangles and squares shows up early in the response for perturbations perpendicular to the axis of movement and later for perturbations parallel to the axis of movement. Figure 7, *c* and *d*, shows subjects' average influence functions for the two different rectangle targets with SE bars, and Figure 7, *e* and *f*, shows the average differences between the two influence functions with SE bars.

We estimated subjects' response times to correct for perturbations from the grouped set of influence functions in each of the six target shape/perturbation direction conditions. We marked the response time as the time at which the average perturbation influence function exceeded 0 by 2 SEMs and stayed above that

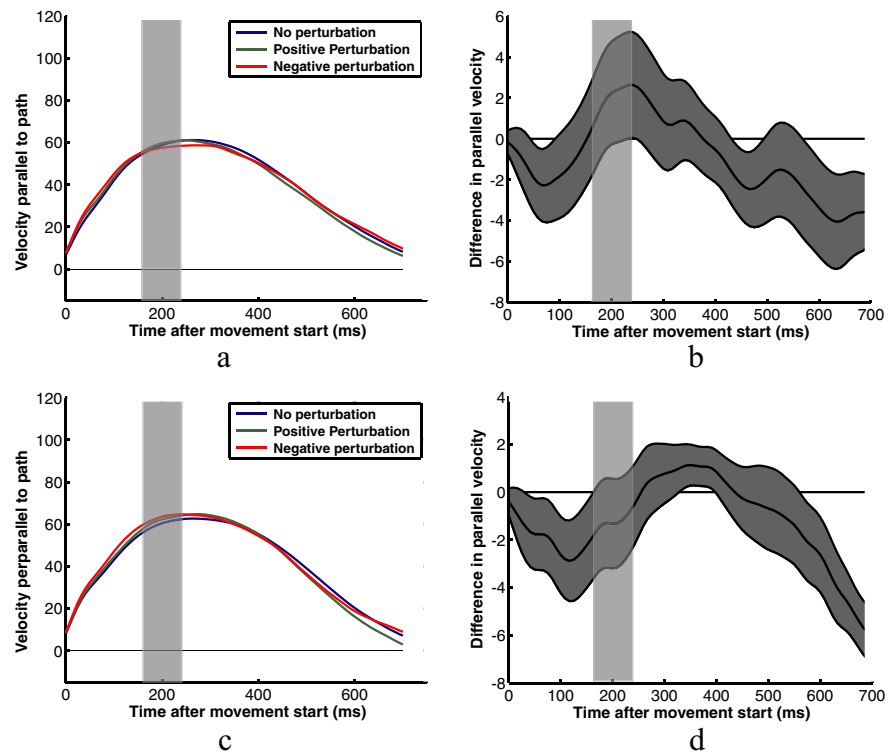


Figure 4. Example velocity profiles for two subjects pointing to the square target. *a* and *c* show subjects' average fingertip velocities for three perturbation conditions—no perturbation; positive 1 cm perturbation in the principal direction of movement, parallel to the path between the start position and the target; and negative 1 cm perturbation in the same direction. Velocities shown are the velocity components parallel to the path. *b* and *d* show the average differences in velocities for positive and negative perturbations. The gray area represents the SEM difference. The transparent gray rectangles represent the time that subjects' fingers were behind the occluder for the average trajectory. The true occlusion times varied from trial to trial.

threshold. Figure 8*a* shows response times estimated for each of the conditions in the experiment. Response times varied from 117 ms for vertical perturbations and square targets to 192 ms for vertical perturbations and vertical targets. We applied the same analysis to the differences in influence functions for vertical and horizontal rectangular targets to find the time at which subjects' corrective responses diverged for the two targets. Figure 8*b* shows the results—167 ms for vertical perturbations and 283 ms for horizontal perturbations. Both the reaction time of subjects' responses to perturbations and the time at which differences in corrections appear are necessarily conservative estimates, in part because of the low slope of the initial response relative to the noise and partly because reaction times are relative to the time at which the very tip of the finger first appears from behind the occluder, when visual information from the emerging finger is likely to be masked by the occluder. The results show that for vertical perturbations, subjects adjusted their feedback control strategies for a majority of the movement duration to match the accuracy demands imposed by differently shaped targets. For the vertical perturbations, the earliest significant corrective responses (for square targets) appear ~ 352 ms after movement initiation, less than halfway in time through the movement (the finger re-emerges from behind the occluder 240 ms after movement onset, on average, and subjects' average movement time was 750 ms). Moreover, the time at which subjects' corrective responses to vertical perturbations for horizontal and vertical rectangles begins to diverge significantly is only 8 ms after the earliest time that the responses themselves become significant. Subjects' average influence functions diverge much later for horizontal

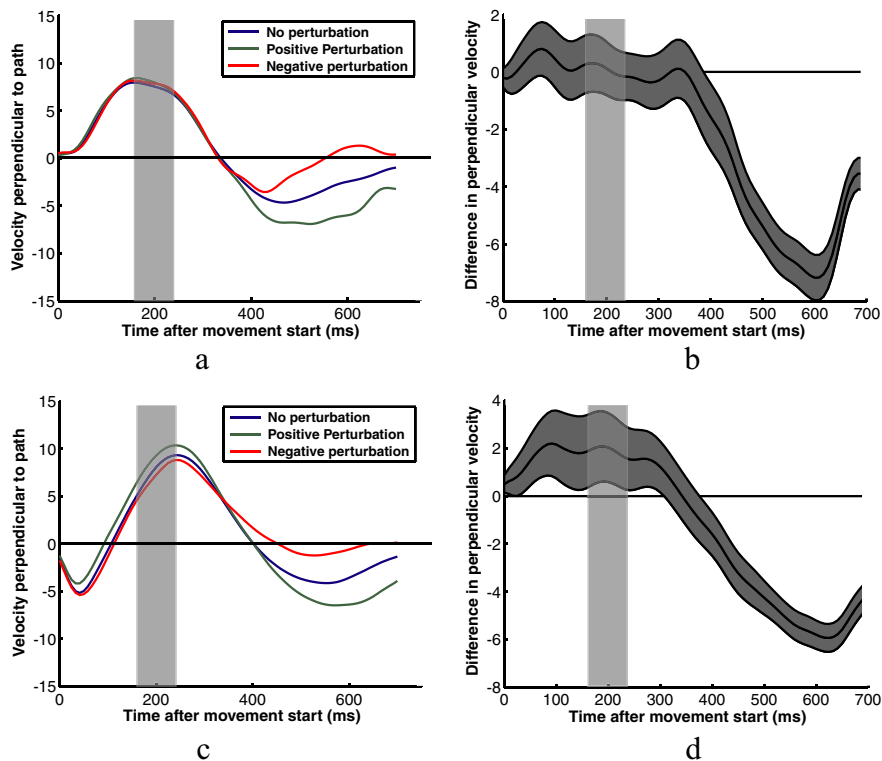


Figure 5. Example velocity profiles for two subjects pointing to the square target. *a* and *c* show subjects' average fingertip velocities for three perturbation conditions—no perturbation, positive 1 cm perturbation perpendicular to the path between the start position and the target, and negative 1 cm perturbation perpendicular to the path. Velocities are in the direction perpendicular to the path. *b* and *d* show the average differences in perpendicular velocities for positive and negative perturbations. The gray area represents the SEM difference. The transparent gray rectangles represent the time that subjects' fingers were behind the occluder for the average trajectory. The true occlusion times varied from trial to trial.

perturbations (~ 110 ms after the initial corrections became significant).

To test the hypothesis that differences in subjects' corrections appear soon after subjects begin correcting for perturbations, we compared the average values of subjects' perturbation influence functions between 150 and 200 ms after its reappearance across different target shapes. Figure 9 shows the results for both vertical perturbations (Fig. 9*a*) and horizontal perturbations (Fig. 9*b*). A one-way repeated-measures ANOVA showed a significant effect of target shape on the average values of the influence functions for vertical perturbations ($F_{(2,7)} = 9.444$; $p < 0.0025$) but not for horizontal perturbations ($F_{(2,7)} = 2.492$; $p = 0.119$).

Experiment 2

The second experiment was designed to test whether or not the feedback control law used by the CNS to control pointing movements is approximately linear. A linear feedback controller coupled to linear dynamics give rise to endpoint corrections to perturbations in visual feedback in one direction that are expressible as a predictable linear combination of the corrections to perturbations in other directions. Although dynamics of the arm are not linear—the state of the fingertip is a nonlinear function of the muscle commands (in our simplified model, it is a nonlinear function of torque commands sent to the joints), a feedback controller that linearly maps feedback error signals to corrective motor commands will appear approximately linear in the corrections evidenced in the fingertip kinematics. This is because the corrections that appear in response to perturbations are small changes overlaid on the underlying feedforward control signal

(or on the average control signal) and the forward dynamics of the arm are smooth enough to be treated as locally linearly. Thus, the hypothesis that the feedback control law is linear leads to the following behavioral prediction. If we know the average corrections to perturbations in directions \mathbf{e}_1 and \mathbf{e}_2 , average corrections to a perturbation in a direction, $\mathbf{e}_3 = w_1\mathbf{e}_1 + w_2\mathbf{e}_2$, a linear controller will lead to corrections Δ_1 with the linear relationship, $\Delta_3 = w_1\Delta_1 + w_2\Delta_2$. Furthermore, this relationship should hold for the perturbation influence functions (which are linear functions of the fingertip kinematics and the perturbations).

Based on the results of experiment 1 and of the modeling, humans optimize their feedback controller so as to decrease the feedback gain in directions that accommodate larger endpoint errors. For the cross target, one could, in theory, design a nonlinear feedback control law that applies small corrections in the vertical and horizontal directions (where corrections are less needed), but large ones in the two diagonal directions. This would be impossible with a linear control law. The results of experiment 2 indicate that subjects did not use such a complex adaptive control law strategy, but rather are approximately linear, both for cases in which a linear control law may be optimal and the case (the cross target) when it is not.

Figure 10*a* shows subjects' corrections to the three different finger perturbations for each of the three targets. The figure only shows the component of the two-dimensional correction vector along the axis of the perturbation. Table 3 summarizes the full two-dimensional correction vectors. Notably, corrections to all three perturbations are high for the cross, although the horizontal component of subjects' corrections to horizontal perturbations is significantly smaller than the diagonal component of their corrections to diagonal perturbations ($T_{(7)} = 6.517$; $p < 0.0005$). The vertical component of subjects' corrections to vertical perturbations were not significantly different from the diagonal component of their corrections to diagonal perturbations ($T_{(7)} = 2.310$; $p > 0.05$). The horizontal component of subjects' corrections to horizontal perturbations for horizontally oriented rectangles was significantly less than for crosses ($T_{(7)} = 6.070$; $p < 0.001$). Similarly, the vertical component of subjects' corrections to vertical perturbations for vertically oriented rectangles was significantly less than for crosses ($T_{(7)} = 8.116$; $p < 0.0001$). This mimics the differences found between rectangles and squares in experiment 1.

The overall picture that emerges is that subjects shaped their control laws for the rectangles as they did in experiment 1—decreasing the feedback gain in the direction of least required accuracy. For the cross targets, they appear to have shaped the control law to correct slightly less in the horizontal direction than the vertical direction—almost treating the cross as a slightly elongated, horizontal rectangle. Looking at projections of subjects' corrections onto the axes of perturbation, however, does not give a full picture of their corrective behavior. By looking at the full

two-dimensional correction vectors, we can better see whether corrections to perturbations in the diagonal direction reflect linear superpositions of corrections to the component perturbations in the x and y directions.

Figure 10*b* shows plots of subjects' average correction vectors as measured for the diagonal perturbations for all three target shapes (Table 3 summarizes subjects' average corrections for the different target shapes and perturbation directions). Since the diagonal perturbations were in the direction $(0.707, 0.707)$ and $(-0.707, -0.707)$, a linear control law would predict that the two-dimensional correction vector obtained from Equation 1 for the diagonal perturbations would be related to the correction vectors for vertical and horizontal perturbations by the relationship, $\Delta\mathbf{x}_{\text{diag}} = 0.707\Delta\mathbf{x}_{\text{horiz}} + 0.707\Delta\mathbf{x}_{\text{vert}}$. Superimposed on the graph are plots of the predicted correction vectors. The figure clearly shows no significant difference in the measured and predicted corrections. It also shows that corrections to diagonal perturbations are not all concentrated along the 45° diagonal of the perturbation axis. In particular, subjects' average correction vectors for horizontally oriented rectangles and crosses are rotated toward the vertical relative to the average correction vectors for vertically oriented targets. This follows from the apparent linearity of the control law—subjects corrected more for the horizontal component of the diagonal perturbations than for the vertical component for vertically oriented rectangles. Similarly, for horizontally oriented rectangles and crosses, subjects corrected more for the vertical component of the perturbations.

Subjects' behavior as reflected by their overall corrections measured at the end of their movements was reflected in their online behavior as well. Figure 11, *a* and *b*, shows perturbation influence functions computed in the x and y directions for the horizontal and vertical perturbations, respectively (as shown in Table 3, very small proportional corrections appear in the orthogonal directions for these perturbations—these are not shown here). The perturbation influence functions are qualitatively similar to those in experiment 1, with subjects responses to perturbations for the cross appearing similar to their responses to perturbations for the square in experiment 1. The one difference, reflecting the pattern in the endpoints, is that subjects corrected somewhat less online for horizontal perturbations for the cross than for the vertical rectangle. Similarly, the linearity in the endpoint corrections appears to hold for the whole pattern of online corrections. Figure 11, *c* and *d*, shows subjects perturbation influence functions computed from their corrections in the x and y directions for diagonal perturbations along with the influence functions predicted from a linear sum of the influence functions derived from the horizontal and vertical perturbations.

The scatter of subjects' endpoints on unperturbed trials reflects a pattern that would appear from effectively treating the crosses as slightly elongated horizontal rectangles. Figure 12 shows a representative example of one subjects' endpoint scatter

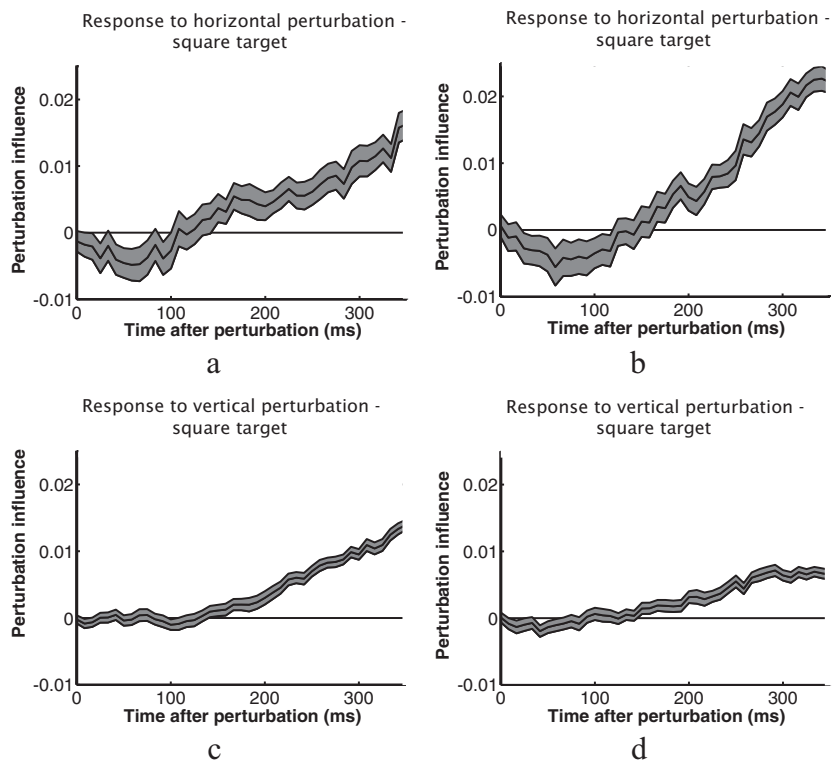


Figure 6. Perturbation influence functions calculated for the same subjects and conditions whose kinematics are shown in Figures 5 and 6. SEs shown in gray were calculated by bootstrap. We resampled the trials used to compute each influence function and calculated the SD of the resulting bootstrapped estimates of the influence functions. *a, b*, Influence functions calculated for horizontal perturbations and square targets from the kinematic data shown in Figure 5. *c, d*, Influence functions calculated for vertical perturbations and square targets from the kinematic data shown in Figure 6.

for the three target buttons. Table 4 shows the average endpoints and SDs of subjects' scatter in the x and y dimensions. Of particular note is that the scatter is elongated horizontally for the horizontal rectangle (as in experiment 1), similarly, but less so for the cross, and vertically for the vertical rectangle. The differences in endpoint scatter that were significantly different from the cross are shown in bold in Table 4. Although the average horizontal scatter appears somewhat larger for the cross than the vertical rectangle, the difference did not reach significance. This raises the possibility that the CNS could potentially use a nonlinear control law of the type that would be optimal for the cross—if sufficiently pressed to do so—but subjects are unable or simply choose not to optimize a complex cost function shaped to the accuracy demands of the cross. In the context of the current experiment, it may be that, although the CNS can efficiently adapt to different target shapes on a trial-by-trial basis by linearly combining simple control laws, the task demands preclude computing and using a more complex control law matched to the crosses. Given the low motor variability demonstrated by subjects (relative to target size), the associated reduction in cost would likely be minimal.

Modeling

To understand how task demands shape feedback control strategies in the experimental task, we simulated the performance of optimal controllers derived for the different target shapes used in experiment 1 for a simplified model of the human arm. Optimal controllers are defined as those that minimize the expected (average) value of a cost function. The cost function coupled with the dynamical system model and the noise assumptions for the system determine the form of the optimal controller for a task

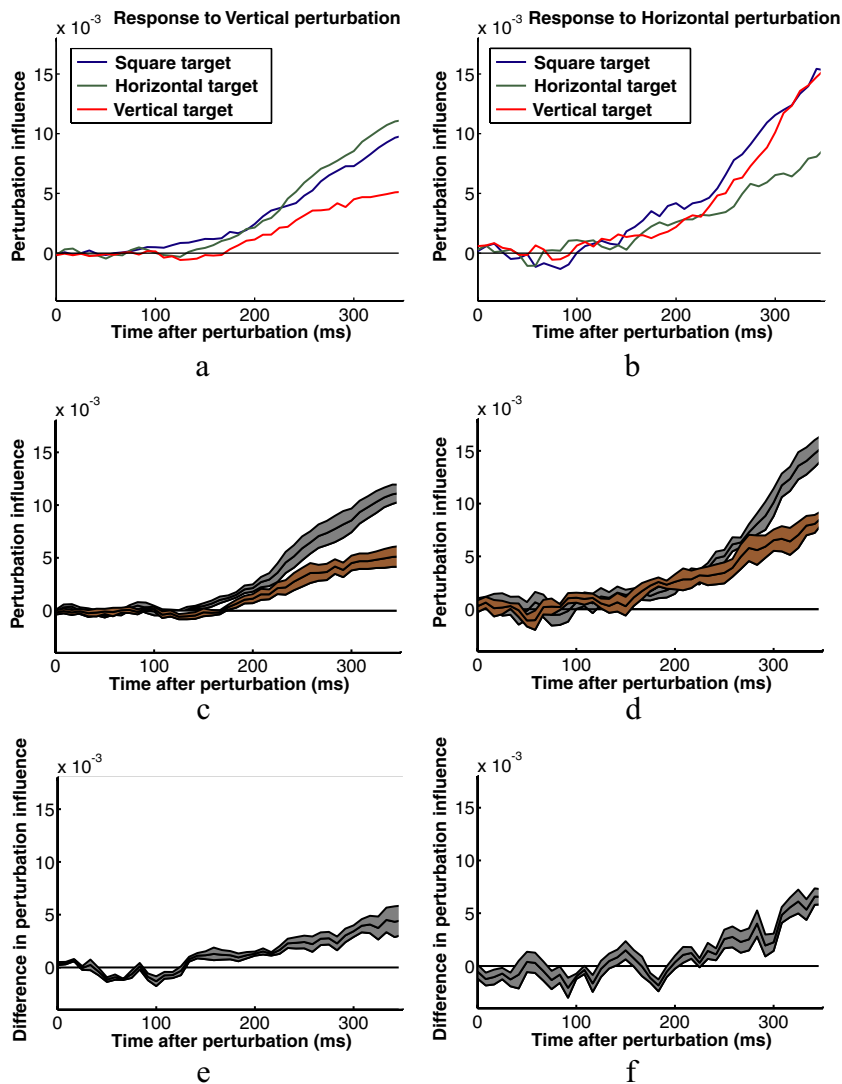


Figure 7. Perturbation influence functions averaged across the eight subjects for the three button shapes. *a*, Vertical perturbations (perturbations perpendicular to the axis between the start and target positions). *b*, Horizontal perturbations (perturbations parallel to the axis between the start and target positions). *c*, *d*, Same plots for horizontal and vertical rectangle conditions with error bars (SEs of the between subjects' means) plotted in gray. *e*, *f*, Average within subject differences in the influence functions shown in *c* and *d* with error bars in gray.

(Todorov and Jordan, 2002). In our simulations, we asked whether a simple performance cost (maximizing the probability of hitting a target) was enough to account for the qualitative behavior shown by subjects; in particular, whether optimal controllers derived for each target shape show the measured shape-dependent differences in feedback gain.

Although most motor control models based on optimal control theory impose some form of energy or smoothness constraint (Flash and Hogan, 1985; Uno et al., 1989) (no. 3220), performance costs alone may be enough to shape control laws to match those found in humans (Harris and Wolpert, 1998). Feedback control can, in theory, induce both kinds of cost. First, in a noisy system, feedback corrections necessarily make individual movements less smooth and add to the force output or energy expenditure of the system. Thus, high feedback gains will add to the subjective cost of a controller. Less obvious are possible deleterious effects created by the proportional noise that accompanies corrective signals. The added noise is generally outweighed by the fact that feedback corrects for the effects of noise in the sys-

tem; however, in a nonlinear plant, noise created by the signals generated to correct for errors in one dimension can leak into the movement in other dimensions. Because hand movements are generated by torques created at rotary joints by muscles, noise created by corrective signals for errors in one spatial dimension will affect the movement of the end-effector in both dimensions. When a task requires greater accuracy in one dimension than another, this nonlinear effect can theoretically create a situation in which feedback corrections for errors that do not need to be corrected (e.g., in the vertical direction for the vertically oriented targets in the experiment) can lead to greater variance in the dimension in which high accuracy is needed (e.g., the horizontal direction for the same targets). The end result is that a controller with high feedback gain for sensory signals of errors in the spatial dimension demanding less accuracy may lead to worse average performance than one with lower feedback gain for those signals.

We derived optimal controllers for each of the targets used in experiment 1 using a cost function that only contained objective performance terms, that is, that penalized misses and high endpoint velocities, but not the force output of the controller. The only free parameters of the model that we adjusted to match human performance were the sensory and motor delays, the coefficient of the proportional motor noise and the coefficient, λ_{stop} , that determined the relative contribution of endpoint accuracy and stopping velocity to the cost. We included a 10 ms motor delay and a 20 ms sensory delay in the models (see Materials and Methods) to match the temporal profile of subjects' performance (discussed in more detail below). We set the coefficient of proportional motor noise to 0.06 to approximately match the variance of subjects' endpoints for the square targets and we set $\lambda_{\text{stop}} = 0.5$. Although somewhat arbitrary, this value gave average endpoint corrections qualitatively similar to subjects. Changing this value has a somewhat complex effect on the models' performance, although the qualitative pattern of endpoint corrections remains the same for a large range of values (e.g., setting $\lambda_{\text{stop}} = 0.1$ leads to slightly greater corrections for square targets and slightly less corrections for rectangular targets).

To capture the variability in movement duration shown by subjects, we computed and simulated the performance of seven optimal controllers for movement durations ranging from 680 to 800 ms in steps of 20 ms (approximately the mean of subjects' movement times ± 1 SD). All performance metrics are averages of those derived from controllers derived for the different movement durations.

Figure 13 shows the main performance features of the controllers derived for each of the three target shapes. The proportional

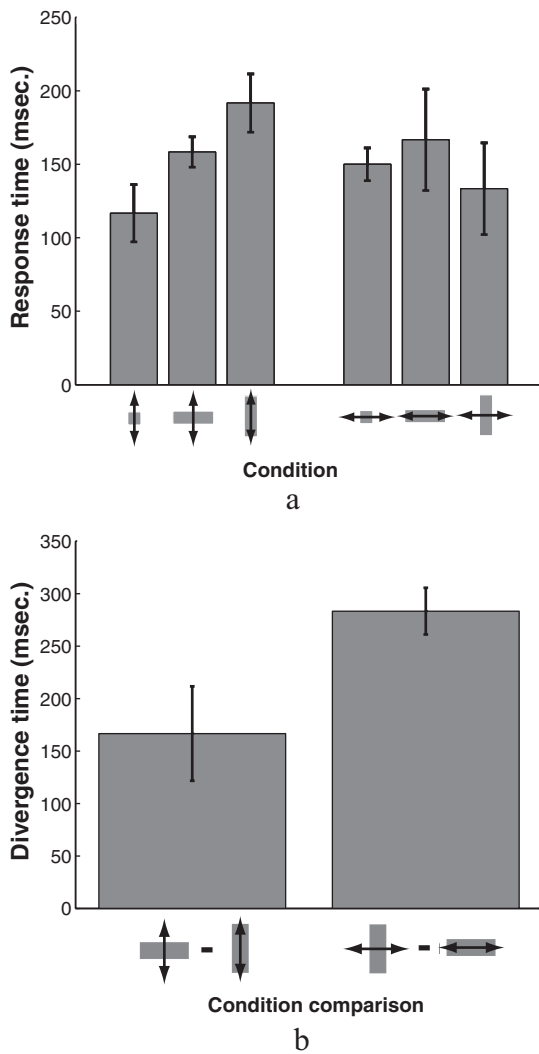


Figure 8. *a*, Response times estimated the collection of all subjects' influence functions for the six target shape and perturbation direction combinations. Response times were estimated as the point at which the average influence functions exceeded 0 by >2 SEs and remained above that threshold (see, for example, Fig. 5*c,d*). *b*, Estimates of the times at which subjects' corrective responses to perturbations for vertical and horizontal rectangles diverged for the two perturbation directions using the same criterion (see Fig. 5*e,f*). Error bars show SEs estimated from a bootstrap procedure in which response times were estimated using the above procedure for 1000 random draws, with replacement, of the eight influence functions derived for each subject.

magnitude of each controller's corrections to vertical and horizontal corrections shows the same qualitative pattern as subjects in experiment 1 (Fig. 3). The optimal controller for the vertical rectangles corrects significantly less for vertical perturbations in sensory feedback than does the optimal controller for horizontal or square targets. Similarly, the optimal controller for horizontal rectangles corrects significantly less for horizontal perturbations of visual feedback than does the optimal controller for vertical or square targets. These changes in feedback gain lead to changes in the error covariance of the models' endpoints that qualitatively mimic those of subjects. The SDs of endpoints in the long dimensions of rectangular targets are approximately twice the SDs of endpoints in the same dimensions for the square targets.

We earlier argued that noise considerations for a system with a nonlinear mapping of motor commands to endpoint position like the human arm could drive the system to decrease the feed-

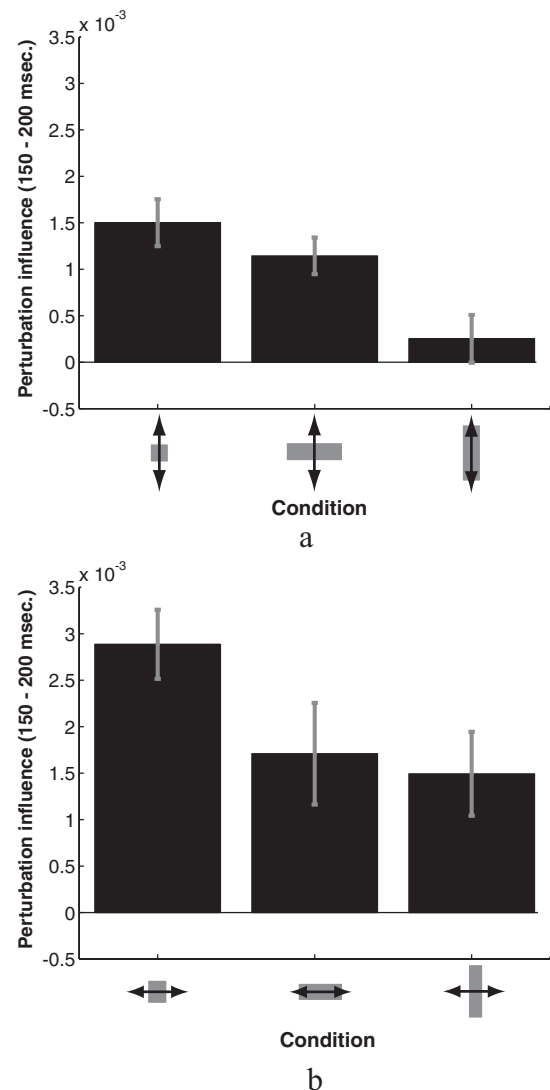


Figure 9. The average magnitude of subjects' perturbation influence functions between 150 and 200 ms after the finger reappeared from behind the occluder. *a*, Results for vertical perturbations. *b*, Results for horizontal perturbations. Error bars indicate SE.

back control gain for sensory error signals in stimulus dimensions that did not require high accuracy. This is based on the intuition that proportional noise in torque commands generated to correct for errors in one dimension will leak into the other dimension. To test whether this in fact occurs, we compared the endpoint covariance for the optimal controller derived for the square target (which maximizes feedback gain in both vertical and horizontal dimensions) with the endpoint covariance of the optimal controllers derived for the horizontal and vertical targets, respectively. Table 5 shows the SDs of endpoints for the three. As predicted, the SDs of endpoints in the direction of the short dimensions of rectangular targets is larger for the optimal controller derived for the square target than it is for the optimal controllers derived for the rectangular target shapes. More generally, as shown in Figure 13*a*, the optimal controllers shape the error covariances of the endpoints to match the shapes of the targets; that is, the optimal controllers for rectangular targets trade off variance in the constrained dimension for variance in the unconstrained dimension.

To compare the temporal response properties of the model with that of human observers, we simulated running the optimal

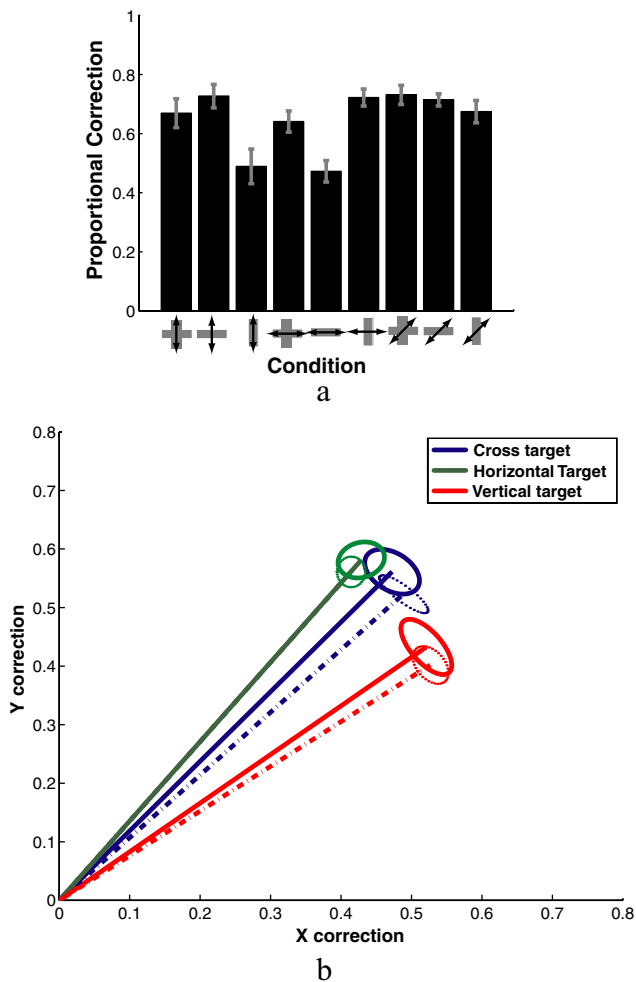


Figure 10. *a*, Average endpoint corrections for perturbations in experiment 2, expressed as a proportion of the perturbation. Corrections shown here are projections of the two-dimensional correction vectors onto the axis of the perturbation. Error bars indicate SE. *b*, Two-dimensional plot of subjects' average correction vectors in response to diagonal perturbations as calculated from Equation 1. This should be read as subjects' corrections to a perturbation of the finger by $[-0.707, -0.707]$ (in centimeters). Since responses to positive and negative perturbations are averaged, the negative of these vectors would reflect subjects' average corrections to perturbations of the finger of $[0.707, 0.707]$. The solid lines show the measured correction vectors. The dashed lines show the corrections vectors predicted by a linear sum of the corrections to vertical and horizontal perturbations. Ellipses centered at the ends of the vectors are SE covariance ellipses.

controllers in the conditions of experiment 1 by running the controllers with motor and sensory noise and “measured” trajectories by adding simulated Optotrak measurement noise to the resulting finger trajectories, as described in Materials and Methods. A total of 1000 trials was run for each target shape and at each of the seven durations simulated (200 for each of the no-perturbation condition and the four perturbation conditions). The resulting trajectories were analyzed in the same way as subjects' trajectories to compute perturbation influence functions for each perturbation and each target shape. Figure 14 shows the resulting influence functions. They replicate subjects' performance (compare with Fig. 7) in several important respects. First, they show the same delays that appear in subjects' influence functions; second, differences in the influence functions appear early in movements; third, the perturbation influence functions for perturbations in the principal direction of movement increase more slowly early in the response than do the influence functions

derived for perturbations perpendicular to the principal direction. This replicates similar behavior found previously (Saunders and Knill, 2005) and occurs despite the fact that the controller contains no structural distinction between control in one direction or another—it arises from the dynamics and statistics of the movements and of the sensory feedback.

Although the optimal controllers show a clear decrease in variance in one spatial dimension as a result of reducing feedback gain in the other, this is not well reflected in subjects' data. Subjects show no significant change in endpoint variance along the constrained dimension when they reduce feedback gain in the other, unconstrained dimension. Thus, they achieved no net performance improvement from the change in feedback policies across shapes. Given the uncertainty in our estimates of endpoint variance (see Table 1), the small changes predicted by the optimal controllers is within or near the limits of uncertainty in our estimates of endpoint variability; thus, it is difficult to say much about this difference in effects. Any small effects present would be further masked by other sources of constant noise in subjects' performance, variance attributable to variance in finger pose and placement of the finger pad at time of contact. Still, the lack of significant effects of the kind predicted provides weak evidence that the CNS is not shaping its feedback control law based on accuracy demands.

The performance metrics of the optimal controllers differ from that of subjects in several other regards. First, the optimal controllers for the square target correct somewhat more for perturbations than do subjects, whereas the optimal controller for the horizontal rectangles corrects less for horizontal perturbations than do subjects. Although such differences are potentially interesting, we should note that the absolute magnitude of corrections is greatly influenced by a number of model parameters about which we have limited information. Increasing sensory noise, for example, decreases the corrections of the models. Increasing the initial uncertainty about the position of the finger and increasing the motor noise have the opposite effect—leading to increased correction magnitudes. All of these effects arise naturally from the behavior of the Kalman filter that estimates finger state from sensory information—it gives more weight to incoming sensory information as the uncertainty of the sensory information decreases or the uncertainty of the internal estimate of hand state increases. Izawa and Shadmehr (2008) showed that humans are sensitive to these kinds of changes—subjects correct less for online shifts in target position when the uncertainty in the sensory information about the final target location is high than when it is low (the effect of online sensory uncertainty) and more for online shifts in target position when the uncertainty in the sensory information about the initial target location is high than when it is low (the effect of initial estimation uncertainty).

Another pattern that appears in human behavior that is stronger in the optimal controllers is the shift in mean endpoints away from the center for the rectangular targets. The optimal controllers shift their mean endpoints for horizontal targets toward the starting position more than subjects did. This results from the stopping cost term, as shorter movements lead to lower endpoint velocity variances. Although we can in theory search for the parameter settings that quantitatively match subjects' data, it would seem premature given the range of possible accounts for the difference (e.g., subjects may effectively incorporate a tighter error bound on endpoints than the actual sizes of the targets). A more striking difference is that the optimal controllers for vertical targets shift their mean endpoints away from the simulated subject,

Table 3. Mean corrections to perturbations for the three different target buttons used in experiment 2

	Horizontal perturbations ($\mu_{\Delta x}$, $\mu_{\Delta y}$) \pm SE (cm)	Vertical perturbations ($\mu_{\Delta x}$, $\mu_{\Delta y}$) \pm SE (cm)	Diagonal perturbations ($\mu_{\Delta x}$, $\mu_{\Delta y}$) \pm SE (cm)
Cross	(0.641, 0.070) \pm (0.036, 0.014)	(0.050, 0.670) \pm (0.025, 0.049)	(0.473, 0.561) \pm (0.028, 0.027)
Vertical	(0.722, 0.074) \pm (0.024, 0.031)	(0.023, 0.489) \pm (0.019, 0.059)	(0.521, 0.433) \pm (0.026, 0.033)
Horizontal	(0.473, 0.067) \pm (0.037, 0.015)	(0.133, 0.727) \pm (0.041, 0.040)	(0.429, 0.581) \pm (0.023, 0.021)

Horizontal perturbations refer to perturbations parallel to the axis between the starting position and the target position; vertical perturbations refer to perturbations perpendicular to this axis. x is taken to be in the direction of this axis, and y is taken to be perpendicular to the axis. Significant differences between the corrections for rectangular targets and the cross targets are highlighted in bold.

whereas subjects show a small shift in the opposite direction. One possible account for this is that subjects incorporate an energy cost, which our simulations show leads to a shift in same direction as subjects; however, other accounts are possible. When we optimized subjects' eye fixation positions along with the control law, we found that the optimal fixation is significantly shifted in the same direction as those shown in subjects' finger endpoints (e.g., toward the starting location in the horizontal target condition). This is correlated with a similar shift in average endpoints for the controllers—an effect that derives from maximizing the acuity of the visual information available about the relative position of the finger and the target. Although an interesting prediction of optimal control, we do not have eye-tracking data to confirm or disconfirm the hypothesis. We should emphasize, however, that when we constrain the mean endpoints to different locations in the vertical target (by incorporating an appropriate term in the cost function), one finds very little difference in the accuracy or stopping costs for the different endpoints (until they approach the target borders); thus, any one of a number of other constraints may be in operation driving subjects to aim below the centers of the vertical targets.

Finally, the behavior of the models reflects a trade-off between the accuracy cost and the stopping cost (Liu and Todorov, 2007). For the nonquadratic cost function used here (probability of a miss), the trade-off is complex—when we explored changing the weight on the stopping cost, it led to changes in both the mean endpoint and feedback gain. These were coupled by the cost function—when we explored reducing the weight to the stopping cost, the optimal controllers for the square targets increased their feedback gain, but the optimal controllers for the rectangular targets decreased their feedback gains in the long dimensions of the target. This resulted from a shift of the mean endpoint toward the middle of the target—a behavior that increased the variance of endpoint velocity but improved accuracy, while reducing the need to correct for errors in the long dimensions of the targets. Although we could in theory adjust the magnitudes of the models' corrective responses by varying parameters like these, and indeed might be able to do so to better fit subjects quantitative data, what does not vary across such changes is the qualitative differences in feedback gain in a particular direction across different shapes.

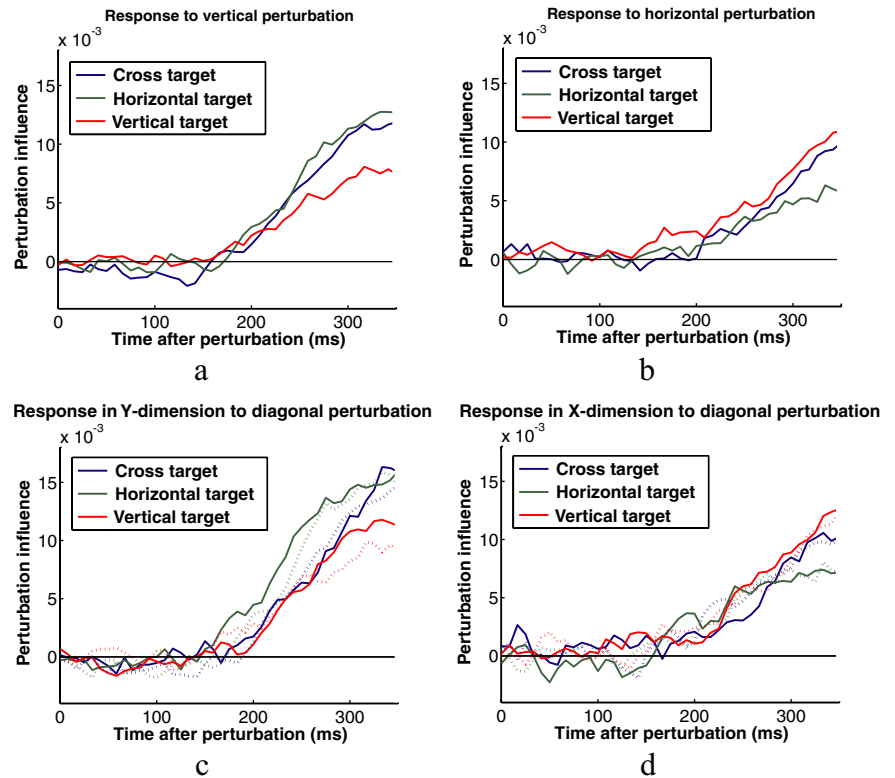


Figure 11. Perturbation influence functions for the three button shapes. *a*, Influence functions computed from the y component of subjects' trajectories in trials with vertical perturbations. *b*, Influence functions computed from the x component of subjects' trajectories in trials with horizontal perturbations. *c*, *d*, Perturbation influence functions derived from the y and x components, respectively, of subjects' trajectories computed from trials containing diagonal perturbations. The dashed curves show the perturbations influence functions predicted from a linear superposition of the perturbation influence functions computed from the trials with vertical and horizontal perturbations (*a*, *b*).




Figure 12. The scatter of endpoints on unperturbed trials for a representative subject in experiment 2 for each of the three targets. The coordinate frame is aligned with the axis between starting and target positions.

Discussion

The experimental results show that humans adjust how they use visual feedback from the hand on a trial-by-trial basis to fit the accuracy demands of individual movements. The variability of subjects' endpoints was in qualitative agreement with these results—endpoint variability was higher in spatial dimensions requiring less accuracy and with lower feedback gain. This kind of

Table 4. Means and average covariances of subjects end-points on unperturbed trials for the three targets in experiment 2

	$(\mu_x, \mu_y) \pm SE$ (cm)	$(\sigma_x, \sigma_y) \pm SE$ (cm)	$\rho \pm SE$
Cross	$(0.224, -0.234) \pm (0.079, 0.033)$	$(0.414, 0.319) \pm (0.037, 0.012)$	-0.240 ± 0.062
Vertical	$(0.065, -\mathbf{0.524}) \pm (0.065, 0.079)$	$(0.304, \mathbf{0.450}) \pm (0.020, 0.034)$	-0.020 ± 0.034
Horizontal	$(\mathbf{1.077}, -0.117) \pm (0.234, 0.054)$	$(\mathbf{0.578}, 0.278) \pm (0.055, 0.014)$	-0.036 ± 0.043

SDs and correlation coefficients are averages across subjects. x is taken to be in the direction of this axis, and y is taken to be perpendicular to the axis. Significant differences between subjects' endpoint distributions for cross targets and subjects' endpoint distributions for vertical and horizontal buttons (means and SDs) are highlighted in bold.




Figure 13. *a*, Endpoint covariance ellipses for optimal controllers derived for each of the three target shapes individually. Ellipses represent twice the SD of endpoint positions. *b*, The average magnitude of correction for 1 cm perturbations in horizontal and vertical directions for each of the three controllers (the controller is indexed by the target shape shown on the x -axis).

Table 5. Performance of three different controllers (each row) shown as the endpoint SDs in x and y

Controller	σ_x (cm)	σ_y (cm)
Square	0.30	0.34
Vertical	0.28	0.87
Horizontal	0.44	0.27

variance pattern, in which variance is highest in kinematic dimensions less relevant to a task, has been shown before (Scholz and Schöner, 1999; Domkin et al., 2002); however, in previous studies, one could not discount the possibility that changes in feedforward rather than feedback control shaped subjects' motor variability. The simulations of the optimal controller for the task described here show clearly that significant differences in endpoint variability result from differences in the feedback control law.

Although it would seem natural that changes in feedback gain should lead to concomitant changes in variance, the data of the current experiments cannot clearly distinguish between variance attributable to motor planning and online control. More direct evidence that changes in feedback policies reflect themselves in the variance structure of movements comes from studies of task-dependent changes in feedback control policies for bimanual

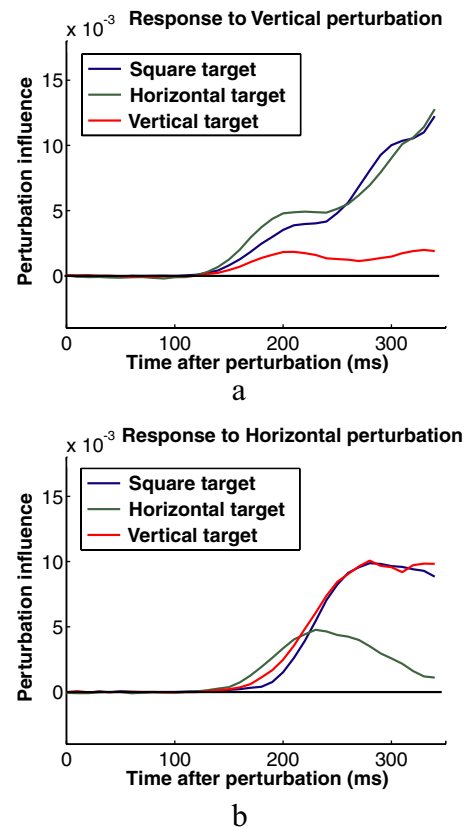


Figure 14. Perturbation influence functions computed for the optimal controllers derived for each shape. *a*, Results for vertical perturbations. *b*, Results for horizontal perturbations.

control. Diedrichsen (2007) showed that differences in feedback coordination between left and right hands for coordinated and independent control tasks cooccur with changes in the covariance of left and right hand movement directions. These differences appear early in movements after a delay of a <200 ms, consistent with the hypothesis that the differences were created by differences in the online control policies across tasks.

Automaticity of control function

Subjects' fast reaction times to the perturbations (typically <160 ms) argue that they result from automatic online corrections to errors signaled by sensory information in what has been described as the "autopilot" mechanism in human motor control. These response times are in qualitative accord with many others reported in the literature, reflecting automatic corrections to visual shifts in target position (Prablanc and Martin, 1992; Brenner and Smeets, 2003; Izawa and Shadmehr, 2008), hand position (Saunders and Knill, 2003, 2004, 2005; Franklin and Wolpert, 2008), cursor position on a screen (Brenner and Smeets, 2003), or background motion (Saijo et al., 2005). The current results show that the magnitude of these corrections is mediated by the accuracy demands of different targets. For visual perturbations per-

pendicular to the principal motion direction, the differences in subjects' responses appeared very soon after the onset of corrective responses, suggesting that they result from differences in the system mediating automatic online control of movements. Moreover, the linearity of subjects' responses suggests a common mechanism subserving all of the responses.

Several other factors argue that a single automatic control mechanism is responsible for the behavior shown here. First, the perturbations were not accompanied by visual transients (masked by the occluder) that might trigger a discrete change detection system. Second, the small magnitudes of the perturbations (1 cm) were within the variance of subjects' finger positions at the point when the finger emerged from the occluder and were little greater than subjects' localization acuity at the eccentricity of the occluder (~ 1.3 cm in tabletop coordinates at an eccentricity of 17° of visual angle). They were therefore within the variability that the automatic feedback controller presumably deals with in normal, unperturbed movements. Finally, not only did subjects not detect the perturbations, but their reaction times to correct for the perturbations were smaller than those found for voluntary changes in movements (>200 ms) (Day and Lyon, 2000; Saijo et al., 2005; Franklin and Wolpert, 2008). It is therefore unlikely that the brain system mediating voluntary changes in movements was engaged in this task.

What brain systems are involved in using visual feedback specifically from the moving hand to control hand movements? It is often assumed that a common system is involved in correcting for target changes that occur during movements and correcting for visually signaled errors in hand movements (Desmurget et al., 1999). A strong candidate for this area is the posterior parietal cortex—a brain region thought to be involved in integrating sensory signals with feedforward information from internal models (Desmurget et al., 1999; Desmurget and Grafton, 2000; Shadmehr and Krakauer, 2008) and shown to be involved in fast online corrections to target shifts (Desmurget et al., 1999; Tunik et al., 2005). It remains an open question as to whether sensory feedback from the hand and sensory signals about a target engage common corrective mechanisms.

Task demands shape feedback control

Several other groups have recently reported task-dependent changes in online control of hand movements. Liu and Todorov (2007) showed that endpoint stability constraints can lead to changes in how much subjects correct for target perturbations. Subjects correct more when they are allowed to hit a target than when they have to stop at the target (a constraint varied by changing the impedance on a robot arm holding the target). Diedrichsen and colleagues have shown that the coupling of sensory feedback from individual hands to control both hands in bimanual tasks depends on the coupling required by the task (Diedrichsen, 2007; Diedrichsen and Gush, 2009). As found here for corrections in different spatial dimensions, bimanual control appears to be linear in the sense that subjects' corrective responses to force perturbations of the two hands are linear superpositions of the corrections to perturbations to each hand individually (Diedrichsen and Dowling, 2009). Finally, Franklin and Wolpert (2008) have shown that subjects can learn to disregard transient visual perturbations of a cursor moved by the hand when the perturbations do not affect endpoint accuracy.

Whereas most previous work manipulated the nature of the task performed (e.g., touching vs hitting or independent bimanual control of multiple cursors vs common control of one cursor), the current task remained constant across conditions. Only

the accuracy demands changed. Subjects could have performed well in the pointing task using a common control strategy (one designed to minimize errors in all directions). Nevertheless, subjects varied their online control policies to fit the accuracy demands of each target. Perhaps the most important feature of the current study, however, is that target shapes changed randomly from trial to trial. In all of the previous experiments, task demands were blocked. The current results, therefore, show a remarkable degree of flexibility in the online control policies of the CNS, as they vary from trial to trial as a result of changing accuracy demands. This suggests that similar trial-to-trial flexibility would be shown when the nature of the task changes, since in many of those cases the costs of not changing control policies is quite a bit higher than here.

Optimal control models

The principal purpose of deriving and simulating an optimal control model was to address the question of whether accuracy demands by themselves would, in principle, drive changes in optimal control policies. The answer to this question is clearly yes—a result of the fact that feedback control signals in torque space generate noise that is distributed across different spatial dimensions. This does not necessarily imply that energy constraints do not also shape online control strategies. In the context of the current task, both accuracy and energy costs lead to a diminution of feedback gain and so are difficult to disentangle experimentally. O'Sullivan et al. (2009), in a very different task setting, have found that subjects give a greater weight to minimizing effort than to minimizing variability. This supports the common assumption that the CNS takes into account energy when planning and executing movements. It remains to be seen how much such subjective, "effort" costs vary across different tasks, in which the costs of motor variability change.

The two-dimensional arm model is constrained in ways that subjects were not, but it contains the major feature of the arm movement plant that creates the spreading of control noise to multiple dimensions in Euclidean space. The modeling results provide a "proof of concept" that accuracy constraints alone can drive changes in feedback gains that arise from changes in accuracy demands; however, one should take care in drawing inferences from fine quantitative matches (or mismatches) between model and human performance. For this reason, we hesitate to make overly much out of some of the finer detailed comparisons of human and model performance.

Conclusion

We have shown that subjects adjust their feedback control strategies in a fast, flexible, "on-demand" way to match the task constraints imposed by the shapes of target objects in a pointing task. Moreover, the changes shown by subjects mimic those that would be required to optimize accuracy in the pointing task. The results raise an important challenge for computational models—how does the CNS adjust its feedback control strategies so quickly to match the changing task demands of individual movements? Clearly, it does not perform the kind of optimization required to derive optimal controllers for each target on each trial. One possibility is that, over the course of development, the CNS has learned a library of basis control laws that it combines in a flexible way to match the task demands of individual movements. How this might work is a subject of future theoretical work in optimal control.

References

- Brenner E, Smeets JB (2003) Fast corrections of movements with a computer mouse. *Spat Vis* 16:365–376.
- Burbeck CA (1987) Position and spatial frequency in large-scale localization judgments. *Vision Res* 27:417–427.
- Burbeck CA, Yap YL (1990) Two mechanisms for localization? Evidence for separation-dependent and separation-independent processing of position information. *Vision Res* 30:739–750.
- Day BL, Lyon IN (2000) Voluntary modification of automatic arm movements evoked by motion of a visual target. *Exp Brain Res* 130:159–168.
- De Bruyn B, Orban GA (1988) Human velocity and direction discrimination measured with random dot patterns. *Vision Res* 28:1323–1335.
- Desmurget M, Grafton S (2000) Forward modeling allows feedback control for fast reaching movements. *Trends Cogn Sci* 4:423–431.
- 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, Dowling N (2009) Bimanual coordination as task-dependent linear control policies. *Hum Mov Sci* 28:334–347.
- Diedrichsen J, Gush S (2009) Reversal of bimanual feedback responses with changes in task goal. *J Neurophysiol* 101:283–288.
- Domkin D, Laczko J, Jaric S, Johansson H, Latash ML (2002) Structure of joint variability in bimanual pointing tasks. *Exp Brain Res* 143:11–23.
- Flash T, Hogan N (1985) The coordination of arm movements—an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703.
- Franklin DW, Wolpert DM (2008) Specificity of reflex adaptation for task-relevant variability. *J Neurosci* 28:14165–14175.
- Goodale MA, Pelisson D, Prablanc C (1986) Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature* 320:748–750.
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784.
- Hollerbach MJ, Flash T (1982) Dynamic interactions between limb segments during planar arm movement. *Biol Cybern* 44:67–77.
- Izawa J, Shadmehr R (2008) On-line processing of uncertain information in visuomotor control. *J Neurosci* 28:11360–11368.
- Körding KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247.
- Levi DM, Klein SA, Yap YL (1988) “Weber’s law” for position: unconfounding the role of separation and eccentricity. *Vision Res* 28:597–603.
- Li W, Todorov E (2004) Iterative linear-quadratic regulator design for nonlinear biological movement systems. Paper presented at First International Conference on Informatics in Control, Automation and Robotics, Setúbal, Portugal, August.
- Li W, Todorov E (2007) Iterative linearization methods for approximately optimal control and estimation of non-linear stochastic system. *Int J Control* 80:1439–1453.
- Liu D, Todorov E (2007) Evidence for the flexible sensorimotor strategies predicted by optimal feedback control. *J Neurosci* 27:9354–9368.
- Mateeff S, Dimitrov G, Genova B, Likova L, Stefanova M, Hohnsbein J (2000) The discrimination of abrupt changes in speed and direction of visual motion. *Vision Res* 40:409–415.
- Orban GA, Van Calenbergh F, De Bruyn B, Maes H (1985) Velocity discrimination in central and peripheral visual field. *J Opt Soc Am A* 2:1836–1847.
- O’Sullivan I, Burdet E, Diedrichsen J (2009) Dissociating variability and effort as determinants of coordination. *PLoS Comput Biol* 5:e1000345.
- Pélisson D, Prablanc C, Goodale MA, Jeannerod M (1986) Visual control of reaching movements without vision of the limb. 2. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. *Exp Brain Res* 62:303–311.
- Prablanc C, Martin O (1992) Automatic control during hand reaching at undetected two-dimensional target displacements. *J Neurophysiol* 67:455–469.
- Saijo N, Murakami I, Nishida S, Gomi H (2005) Large-field visual motion directly induces an involuntary rapid manual following response. *J Neurosci* 25:4941–4951.
- Sarlegna F, Blouin J, Bresciani JP, Bourdin C, Vercher JL, Gauthier GM (2003) Target and hand position information in the online control of goal-directed arm movements. *Exp Brain Res* 151:524–535.
- Saunders JA, Knill DC (2002) How is visual feedback from the hand used to control reaching movements? *Perception* 31:144.
- Saunders JA, Knill DC (2003) Humans use continuous visual feedback from the hand to control fast reaching movements. *Exp Brain Res* 152:341–352.
- Saunders JA, Knill DC (2004) Visual feedback control of hand movements. *J Neurosci* 24:3223–3234.
- Saunders JA, Knill DC (2005) Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements. *Exp Brain Res* 162:458–473.
- Scholz JP, Schöner G (1999) The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126:289–306.
- Shadmehr R, Krakauer JW (2008) A computational neuroanatomy for motor control. *Exp Brain Res* 185:359–381.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.
- Todorov E, Li W (2004) A generalized iterative LGQ method for locally-optimal feedback control of constrained nonlinear stochastic systems. Paper presented at 43rd IEEE Conference on Decision and Control, Atlantis, Paradise Island, Bahamas, December.
- Toet A, Koenderink JJ (1988) Differential spatial displacement discrimination thresholds for Gabor patches. *Vision Res* 28:133–143.
- 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.
- Uno Y, Kawato M, Suzuki R (1989) Formation and control of optimal trajectory in human multijoint arm movement—minimum torque-change model. *Biol Cybern* 61:89–101.
- Whitaker D, Latham K (1997) Disentangling the role of spatial scale, separation and eccentricity in Weber’s law for position. *Vision Res* 37:515–524.