Exp Brain Res (2005) 162: 458–473
DOI 10.1007/s00221-004-2064-1

RESEARCH ARTICLE

Jeffrey A. Saunders · David C. Knill

Humans use continuous visual feedback from the hand to control both the direction and distance of pointing movements

Received: 4 March 2004 / Accepted: 6 July 2004 / Published online: 8 March 2005
© Springer-Verlag 2005

Abstract Vision of the hand during reaching provides dynamic feedback that can be used to control movement. We investigated the relative contributions of feedback about the direction and distance of the hand relative to a target. Subjects made pointing movements in a 3-D virtual environment, in which a small sphere provided dynamic visual feedback about the position of their unseen fingertip. On a subset of trials, the position of the virtual fingertip was smoothly shifted by 2 cm during movement, either (1) in the direction of movement, which would require adjustments to the distance moved, or (2) orthogonal to the direction of movement, which would require adjustments to the direction moved. Despite not noticing the perturbations, subjects adjusted their movements to compensate for both types of visual shifts. Corrective responses to direction perturbations were observed within 117 ms, and response latencies were invariant to movement speed and perturbation onset time. Initial corrections to distance perturbations were smaller and appeared after longer delays of 130–200 ms, and both the speed and magnitude of responses were reduced for early onset perturbations. Simulations of a feedback control model that optimally integrates visual information over time show that the results can be explained by differences in the sensory noise levels in the visual dimensions relevant for direction and distance control.

Keywords Hand movements · Motor control · Sensorimeter · On-line control · Vision

Introduction

The question of whether movement direction and distance are controlled jointly or independently has been the topic of considerable study. Evidence for the independence of these dimensions has been observed for movement planning (Gordon et al. 1994; Chieffi and Allport 1997; Sainburg et al. 2003), and for sensorymotor adaptation (Krakauer et al. 2000). Independent control of direction and distance is also observed in responses to unexpected target perturbations (Turrell et al. 1998). Whether these movement dimensions are processed differently for on-line visual feedback control remains unknown, however. In previous studies addressing control of direction and distance, visual feedback from the hand was either absent (Turrell et al. 1998) or maintained a constant relationship with the physical hand across many trials (Krakauer et al. 2000). In the latter case, it is difficult to distinguish between effects due to movement planning and on-line feedback control.

We recently reported evidence that humans use continuous visual feedback from the hand to guide reaching, even for relatively fast movements (Saunders and Knill 2003). Figure 1 illustrates the perturbation method used in that experiment and in the one reported here. Subjects performed reaches in a virtual environment (Fig. 1a) in which a sphere was rendered at the visual position of the fingertip. On a proportion of the trials, the visual position of the fingertip was smoothly perturbed by a small amount (Fig. 1b), unnoticed by subjects. In the previous study, we found corrective responses that were detectable within 160 ms of perturbation onset. The response latency was invariant to both the speed of movement and to the timing of the perturbation (early versus midway during the movement), suggesting that visual feedback from the hand is used continuously during the movement.

In the previous study, the fingertip perturbations were applied perpendicular to the path of movement. Sarlegna et al. (2003) recently performed a similar manipulation of visual feedback from the hand, but in the principle direction of movement. They did not observe fast corrective responses. On the perturbed trials in their study, visual feedback about the location of the hand

J. A. Saunders (✉) · D. C. Knill
Center for Visual Science, University of Rochester,
Meliora Hall 246, Rochester, NY 14627–0270, USA
E-mail: saunders@cvs.rochester.edu

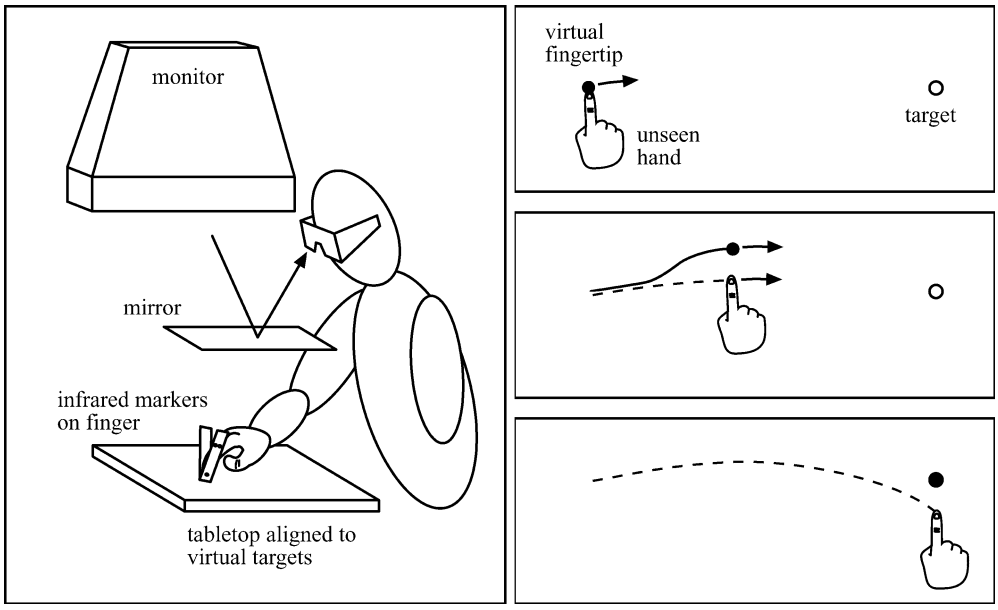

Fig. 1 *Left*: experimental apparatus. Subjects were presented with stereo images using shutter glasses and a monitor viewed through a mirror. The mirror prevented the hand from being viewed, and allowed artificial visual feedback to be presented in the subjects' manual workspace. Subjects moved their fingers to visual targets that were aligned with a tabletop. An Optotrak 3020 system tracked infrared markers attached to a subject's finger, and this data was used to render a virtual sphere representing the subject's fingertip (virtual finger). *Right*: illustration of the task. Views show a top-down projection onto the tabletop. At the start of a trial (*top panel*), a target and the virtual fingertip appeared. The target was positioned 26 cm away from the initial position of the finger, and the target's direction was randomly chosen within a range of $\pm 15°$ around the horizontal midline. The subject's task was to move their virtual fingertip to touch the target. During initial movement, the virtual fingertip coincided with the actual location of the subject's unseen finger. On perturbed trials, the position of the virtual fingertip was smoothly shifted at some point during the movement, so that it no longer matched the actual finger (*middle panel*). To successfully reach the target with the perturbed virtual fingertip, subjects would have to compensate for the displacement (*bottom panel*)

was displaced relative to the hand during subjects' initial orienting saccades to the target. Responses during the primary movement phase only partially compensated (~20%), and began after relatively long latencies, averaging 470 ms. The difference between the two results suggests that the visuomotor system may rely on visual feedback about the direction of the hand relative to the target, but not its distance from the target, at least early in a movement.

There are a number of reasons to expect different contributions from feedback about movement direction and distance. First, as mentioned earlier, there is considerable evidence that direction and distance are processed separately for other aspects of movement control, such as planning and adaptation. Second, perturbations in different directions applied to visual feedback from the hand produce effects that likely differ in perceptual salience. Since hand movements tend to follow straight paths, uncertainty about hand position would be expected to be greatest along the axis of movement, such that there would effectively be more noise in feedback about target distance than direction. Similarly, the human visual system is much more sensitive to changes in velocity orthogonal to an objects' motion path than to changes in velocity along the path (Mateef et al. 2000). Finally, the control problem for these dimensions differs. Once movement is initiated, the direction of movement is relatively constant, and even small errors are diagnostic of future errors in the final position of the hand. In contrast, target distance and hand speed both change markedly throughout the movement. If initial kinematics are only weakly predictive of final movement extent, as was observed by Messier and Kalaska (1999), there might be less incentive for early online feedback control of movement extent, and consequently greater need for later feedback control.

On the other hand, there is evidence that the visual motor system is at least capable of fast online adjustments to both direction and distance movement parameters. Perturbations of the visual target of a reach have been tested both in the direction of movement (Pelisson et al. 1986), requiring adjustment to movement distance, and in a nearly orthogonal direction (Prablanc and Martin 1992), requiring adjustment to movement direction. In both experiments, similar fast corrective responses were observed. While this does not imply that analogous perturbations to visual feedback from the hand would have similar effects, it does demonstrate that both movement direction and distance are potentially subject to online control. Furthermore, the results of the Sarlegna et al. study cannot be directly compared to those of the Saunders and Knill study due to methodological differences. Perturbed feedback occurred earlier in movements in the Sarlegna et al. experiment, and pointing movements were directed to points in space rather than to targets on a solid surface. Finally, as we

will describe later, the analysis method used by Sarlegna et al. was significantly less sensitive than that used by Saunders and Knill.

In the experiment reported here, we used the perturbation method illustrated in Fig. 1 to explicitly compare the contributions of online visual feedback about movement direction and extent. We distinguished these components by comparing responses to perturbations along two different axes relative to the movement path. In the *direction-perturbation* condition, the virtual fingertip was displaced orthogonal to the main direction of movement (Fig. 2a), requiring subjects to adjust the direction of movement. This was the type of perturbation used in our previous study (Saunders and Knill 2003). In the *distance-perturbation* condition, the virtual fingertip was displaced in the direction of movement (Fig. 2b), requiring subjects to adjust the extent of movement.

We observed online corrections to both direction and distance perturbations, but with differing magnitudes and latencies. To test whether these differences could be attributed to the amounts of visual noise in the direction and distance feedback signals, we simulated a feedback control model that optimally integrates noisy visual information during movements, and compared its performance to human data for the conditions tested in the experiment. Under plausible noise assumptions, mostly determined from human psychophysics on motion and position acuity, the model exhibited qualitatively similar effects to the experimental subjects.

Methods

Design and conditions

Subjects performed a simple pointing task in a virtual environment (Fig. 1a), with visual feedback provided by a small sphere that moved in real-time with their unseen fingertip. On half of the trials, the virtual fingertip was perturbed during the reach, as illustrated in Fig. 1b. On these trials, the virtual finger initially coincided with the actual finger position, but at some point along the movement it was shifted in one of four directions. Perturbations were small (2 cm) and smooth, appearing as a blurred step that extended over 4 cm of the movement. Subjects uniformly reported that they were unaware of any unusual changes in the perceived movement of the virtual finger during the experiment, even when told later about the perturbations.

The main experimental manipulation was the axis of perturbation displacements, which was either aligned with the direction between the initial hand location and the target (*distance perturbations*, Fig. 2a) or orthogonal to that direction (*direction perturbations*, Fig. 2b). We also varied the time of perturbation onset and the movement speed. In the *early onset* condition, perturbations occurred when the finger had moved 25% the distance to the target, while in the *mid-reach onset* condition, perturbations occurred when the position of the fingertip had moved 50% of the distance to the target. We ran subjects in one of two speed conditions: a *fast* condition, in which subjects were given feedback to keep their movement durations within a range around 450 ms, and a *slow* condition, in which subjects were given feedback to keep their movement durations in a range around 600 ms. These speeds spanned the range of what pilot subjects subjectively reported as natural movement speeds for the distances used. Each type of perturbation was presented in both possible directions (up or down, forward or backward). Direction and distance perturbations were presented in separate blocks, and the two movement speeds were tested in separate subject sets. Other parameters, including the presence, onset and sign of perturbations, were randomized across trials within blocks.

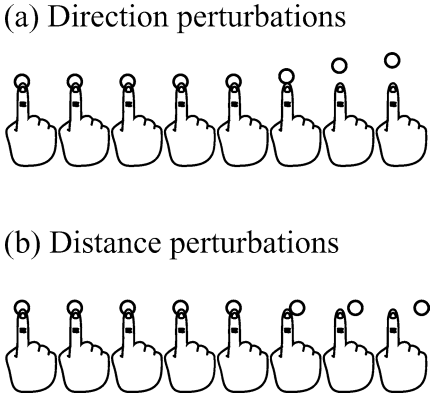

Fig. 2a–b The two types of perturbations. **a** Direction perturbations: the virtual fingertip is displaced in a direction perpendicular to the main direction of movement. **b** Distance perturbations: the virtual fingertip is displaced along the direction of movement. Both examples show positive perturbations. In the experiment, perturbations in the opposite direction were also tested, and the sign was randomized across trials

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 left and right eye views were accurate perspective renderings of the simulated environment. In stereo mode, the monitor had a resolution of 1024×768 pixels and a refresh rate of 150 Hz, or 75 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 inter-ocular cross-talk.

The horizontal mirror prevented a subject's hand from being visible. Visual feedback about hand position was instead provided by a 1 cm diameter virtual sphere that moved in real-time along with the subject's actual finger. The targets of the subjects' movements were 1 cm

diameter circles rendered to be aligned with an unseen tabletop (~55 cm from the eyes).

An Optotrak 3020 system recorded the time-varying position of a subject's finger at 150 Hz. The data was 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 an additional latency on visual feedback of ~10 ms. When computing the rendered position of the virtual fingertip, we compensated for this delay by linearly extrapolating from the latest marker position, using the position from previous frames to estimate velocity. We checked the effectiveness of this extrapolation subjectively by viewing the actual and virtual finger simultaneously (using a half-silvered mirror). Residual errors were not readily apparent.

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 at different depth planes. The combined responses for both eyes were used to determine a globally optimal combination of 3-D 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 to be less than 1–2 mm. Otherwise, the calibration procedure was repeated.

Procedure

The task for each subject was to move their finger back and forth to touch a series of visually-presented targets. At the start of a trial, a new target would appear on the opposite side of the workspace as the current hand position. The target onset was the cue to move. Subjects were instructed to move to touch the target in a fast and "natural" manner. Upon reaching the target, they were to keep their finger in place until the next target appeared. Data collection started once a subject's finger moved 0.25 cm, and continued for 3 s. The recording period was followed by a delay of 800 ms with a blank screen, after which a new target appeared initiating the next trial. The time between trials varied depending on how quickly subjects initiated movement, averaging approximately 4 s.

Subjects were instructed to try to move at about the same speed on each trial. Feedback was provided throughout the experiment to train subjects to reach the target within a range of time around a goal time, which was 450 ms for subjects in the fast condition and 600 ms for subjects in the slow condition. Subjects received positive feedback for movement durations within 75% to 125% of the goal time (a larger circle appeared around target) and negative feedback otherwise (an "X" through the target indicating too fast, no change indicating too slow). For purposes of feedback, reaching the target was defined as bringing the virtual finger within 1 cm of the center of the target. Prior to the main experiment blocks, subjects were allowed 20–40 unperturbed practice trials to familiarize themselves with the timing constraint.

The targets varied in location within two ellipses on the table, which were 10 cm wide, 16 cm tall and separated by 24 cm. The target position for a trial was chosen relative to the end location of the previous trial. For experimental trials, the target distance was always 26 cm (~24° of visual angle), while the direction from the starting position to the target randomly varied between $-15°$ and $+15°$ relative to the left/right axis along the tabletop, subject to the constraint that the endpoint lay within the target ellipse. Randomly interspersed with the experimental trials were a smaller number of trials (20%) with targets at a distance of 22 cm. These trials were included to add some variation in the distance moved, but were not analyzed. Occasionally, subjects moved their hands toward the center of the workspace too early (contrary to instructions), such that it was impossible to satisfy the criteria. In these cases, a random position within the opposite ellipse was chosen as target location, and the trial was discarded from analysis.

On perturbed trials, 2 cm offsets between the actual and visual position of the finger were phased-in gradually over a 4 cm transition range. The amount of displacement on a given frame was a function of the position of a subject's fingertip along the axis from the starting position to the target. On trials with early perturbations, offsets were phased-in when the fingertip was between 4.5 cm and 8.5 cm away from the starting position, as measured along the target axis. On trials with mid-reach perturbations, the transition region was from 11 cm to 15 cm. The centers of these two transition ranges corresponded to where a subject's fingertip was 25% or 50% of the way to the target. Within the transition ranges, the proportion of the perturbation applied was computed using a scaled logistic function. The base function was $f(x) = 1/(1 + 3^{-x})$, where x was the difference between the current position of the fingertip along the target axis and the midpoint of the transition region.

This function was linearly re-scaled to span from 0 to 1 over the 4 cm transition range, $f'(x) = (f(x) - f(-2))/(f(2) - f(-2))$, and the result was multiplied by the magnitude of the eventual displacement (2 cm) to determine the partial displacement. The vector direction of the perturbations was either aligned with the axis from starting position to the target, in the case of distance perturbations, or aligned with the perpendicular axis along the tabletop, in the case of direction perturbations.

The perturbation transitions caused transient discrepancies between the velocity of the virtual fingertip and the velocity of the actual fingertip, which was maximal at the midpoint of the transition. In the case of distance perturbations, the maximal differences in speed were approximately 35–45 cm/s, which corresponds to 60–70% of the actual finger speed, and the differences in the directions of motion caused by the perturbations were negligible. In the case of direction perturbations, maximal differences in speed were 10–12 cm/s, corresponding to 15–25% of the actual finger speed, and the maximal differences in the directions of motion averaged 34°.

Prior to analysis, data were filtered to remove various types of irregular trials. These included trials with incomplete data due to occluded markers, trials with mistimed recordings due to false starts, trials in which subjects began moving before the target appeared, and trials in which motion was not complete within ±40% of the goal time. Some trials contained isolated single frames with occluded markers. If there were no more than four such frames on a trial, the trial was not excluded. Instead, the data from missing frames was filled in by linearly interpolating between adjacent frames. Trials with adjacent missing frames or with more than four isolated missing frames were excluded from analysis.

Subjects participated in two experimental sessions on separate days. Each session began with calibration of the virtual environment, followed by practice trials to familiarize a subject with the task and movement speed, and then two blocks of experimental trials separated by a brief break. Subjects performed 300 trials in each block, of which 240 were experimental trials. This yielded a total of 120 trials per perturbed condition (x-axis onset) and 240 unperturbed trials for each subject.

Perturbations were applied on half of the trials. Perturbed and unperturbed trials were randomly intermixed within blocks, as were the onsets and signs of perturbations. Direction and distance perturbations were blocked by experimental session, with order counterbalanced across subjects. Fast and slow movements were tested in separate subjects.

Analysis

For the analysis, finger positions were represented in a normalized coordinate frame, with the x-axis being the direction from initial position to target, and the y-axis being the direction in the tabletop plane orthogonal to the x-axis. We ignored the vertical (z-axis) components of the data series, treating the finger positions as 2-D points in the xy-plane.

Movements showed large variability in overall trajectories, but were also very smooth, so much of the variability could be due to differences in initial conditions, motor plans, goal paths, and so on. To derive a sensitive measure of the effect of perturbations, we temporally de-correlated the data series using linear filters fit to the data (Saunders and Knill 2003).

We assumed that, for unperturbed trials, the normalized position of the fingertip at time t was a linear function of the n previous positions:

$$p(t) = w_1(t) \cdot p(t-1) + w_2(t) \cdot p(t-2) \ldots w_n(t) \cdot p(t-n) \quad (1)$$

The weights $w_1(t)$, $w_2(t)$, ... $w_n(t)$ were allowed to vary as a function of time, but were assumed to be constant across trials. To align the data series temporally, we defined $t = 0$ to be the midpoint of the perturbation transition, so that t represents approximate time after perturbation. Since there was variability in the speed and time-courses of the trials, the assumption that weights are the same across trials corresponds to assuming that the weights are relatively constant over small shifts in t.

If the data series are smooth and temporally-correlated, the linear model given by Eq. 1 should closely fit the data and account for most of the variability across trials. Responses to perturbations can be distinguished as changes in trajectory that would not be predicted based solely on the previous path. Thus, we augmented the model with an additional term representing the effect of perturbations:

$$p(t) = w_1(t) \cdot p(t-1) + w_2(t) \cdot p(t-2) \ldots w_n(t) \cdot p(t-n) + w_{\text{pert}}(t)\Delta_{\text{pert}} \quad (2)$$

where Δ_{pert} is $+1$ for positive perturbations and -1 for negative perturbations. The function $w_{\text{pert}}(t)$ represents the evolving influence of the perturbation on the subjects' movements, and this should be zero for times before subjects show any corrective response to perturbations. We refer to this function as the perturbation influence function.

For each subject, we applied linear regression to the baseline trials to estimate the autoregressive weights in Eq. 1. We set $n = 6$ for the analysis, but found that the results were insensitive to the exact value for $n > 6$. To compute influence functions for each perturbation type $w_{\text{pert}}(t)$, we correlated the sign of the perturbations $(+1, -1)$ on perturbation trials with the error between the hand position at time t and the position predicted by the autoregressive model. For trials with distance perturbations, perturbation influence functions were computed using only the x-component of the kinematic data (approximately the movement extent), while for trials with direction perturbations, influence functions were

computed using only the *y*-component of the data (approximately equivalent to movement direction).

Response times were computed by locating the first point in time at which the mean perturbation influence functions, averaged across subjects, were significantly different from zero. The raw mean influence function was first smoothed with a one-sided (causal) exponential filter, $f(t) = \exp(t/\lambda)$ for $t < 0$, with $\lambda = 35$ ms, producing a time-weighted cumulative measure of the perturbation influence. The filtered result, $f(t)^* w_{\text{pert}}(t)$, was compared at each time t to a statistical threshold derived from re-sampling baseline trials, and the earliest super-threshold time was taken to be the response latency. The re-sampling procedure consisted of repeatedly applying the analysis to randomly-chosen baseline trials (sampling with replacement) and using the results to estimate the null-model noise distribution of filtered $w_{\text{pert}}(t)$.

To determine error bounds on the response time estimates, we re-sampled across subjects to compute mean influence functions, and used these samples to compute a

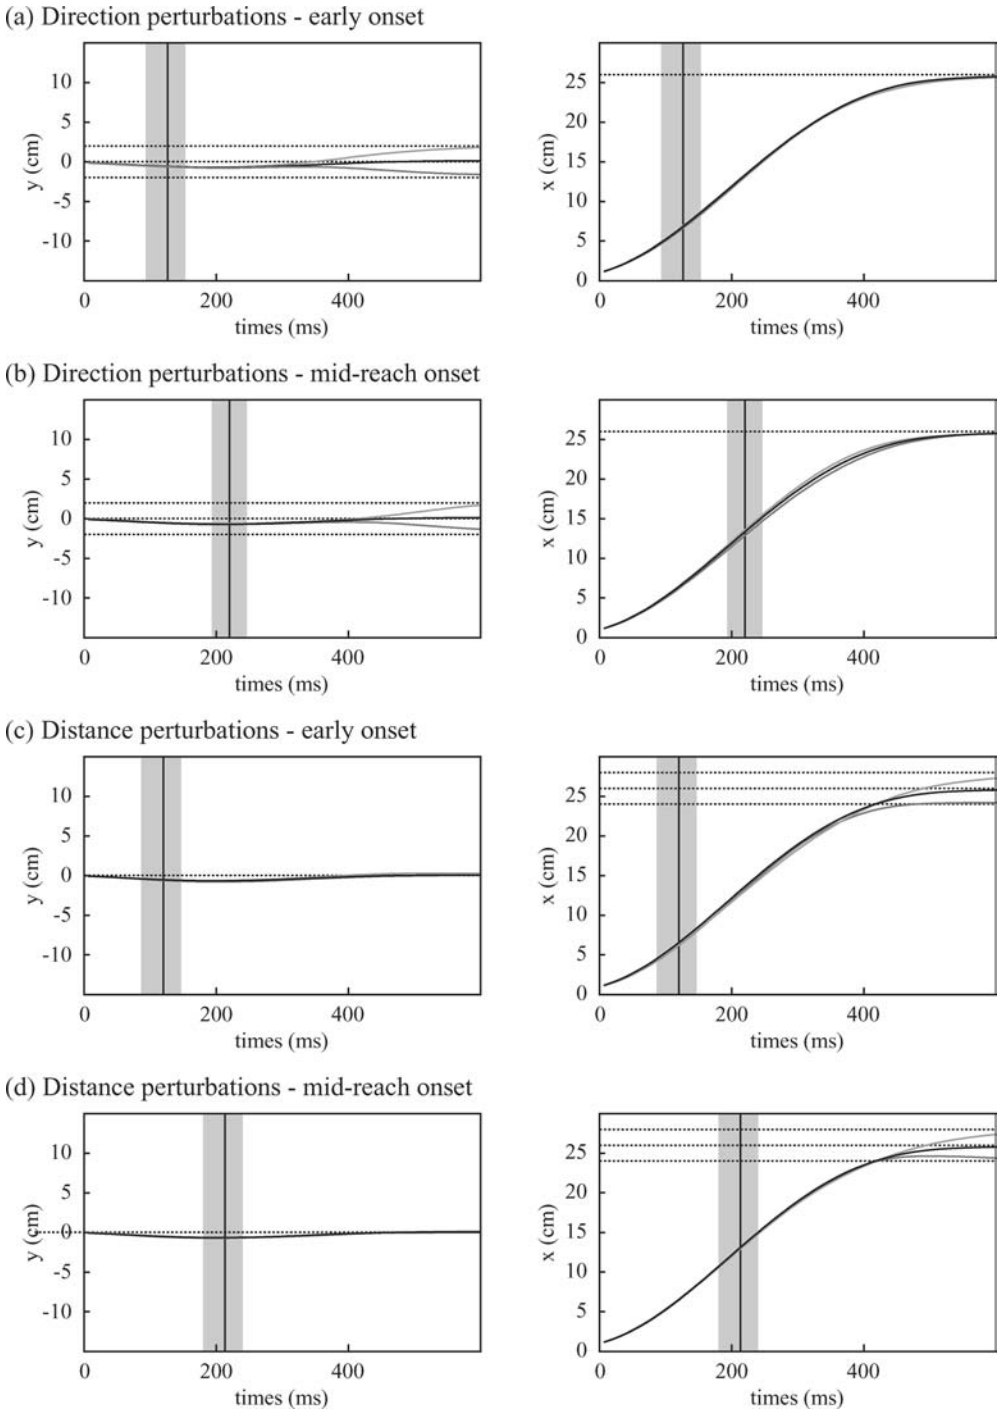

Fig. 3a–d Mean kinematic trajectories for a representative subject in the slow movement condition. The three *black lines* correspond to mean data from trials with positive perturbations (*solid*), negative perturbations (*dotted*), or no perturbations (*dashed*). The *shaded areas* denote the perturbation transition phases, during which the perturbation was gradually applied. The *gray horizontal lines* depict the expected final finger positions if subjects fully compensated for the perturbations. **a** Direction perturbations, early onset; **b** direction perturbations, mid-reach onset; **c** distance perturbations, early onset; **d** distance perturbations, mid-reach onset. For direction perturbations (**a** and **b**), corrections appear in the *y*-axis component of finger position (*left*), and for distance perturbations (**c** and **d**), corrections appear in the *x*-axis component (*right*)

distribution of response times. The filtering and threshold comparison applied to the samples was the same as applied to the actual data. The error bounds plotted in the figures correspond to 5% and 95% points in the sample histograms. Comparisons across conditions were based on the distribution of differences between response times. For a given pair of conditions, the difference in response times was computed for each sample, and the resulting set was taken as representing the distribution of actual differences for purposes of statistical comparisons.

In addition to response times, we report a measure of the amount of correction in response to perturbations. We wanted this measure to reflect only corrections that occurred during the main portion of movements, and not any additional secondary corrections, which requires a criterion for identifying the end of initial movements. We defined this to be the first minimum in the speed profile after the deceleration phase had begun. To avoid detecting spurious re-accelerations, we smoothed the time series with a 10 ms Gaussian filter prior to computing minima, and further required that the hand speed be below 6 cm/s. These criteria were ad hoc, chosen to produce subjectively accurate and reliable measures of the end of initial movements, and small variations in the parameter values did not significantly affect the results.

Subjects

Sixteen naïve, paid subjects participated in the experiment—eight in the fast movement condition and eight in the slow movement condition. Subjects 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.

Results

Figure 3 shows mean kinematics of finger position for a representative subject in the fast movement condition. The four rows correspond to early and mid-reach direction perturbations (a and b) and early and mid-reach distance perturbations (c and d). To generate the mean data series, individual trials were first normalized with respect to target direction so that the main axis of movement was transformed to the horizontal x-axis, and then projected onto the tabletop (ignoring the height off the table). The left panels plot the y-axis component of finger position as a function of time, while the right panels plot the x-axis component. Each graph contains three plots, corresponding to mean kinematics of trials with positive perturbations (solid), negative perturbations (dotted), or no perturbations (dashed). The subject clearly made corrective responses to both direction and distance perturbations. In both cases, the mean finger position functions for positive and negative perturbed trials diverged from that of unperturbed trials, in the direction consistent with bringing the perturbed visual finger to the target. In the case of direction perturbations, this can be seen in the y-axis components, while for distance perturbations, the divergence is in the x-axis plots. The performance of the other subjects was similar.

The amount of correction in response to perturbations, averaged across subjects and conditions, was 1.5 cm, corresponding to 75% of the perturbation magnitude. Figure 4 plots mean percent correction for the different movement speed and perturbation onset sub-conditions. Subjects showed greater overall correction during slow movements than for fast movements, and more correction in response to early perturbations than to late perturbations. Thus, magnitude of correction increased with the time available for correction, as would be expected from continuous feedback control if the post-perturbation periods were too short to fully compensate.

As described in the "Methods" section, we derived perturbation influence functions for each experimental condition to characterize the influence of the perturbation on subjects' finger movements as a function of time after the perturbation. Figure 5 plots mean perturbation influence functions for direction perturbations (solid) and distance perturbations (dashed) for each of the four combinations of movement speed (slow versus fast) and perturbation onset (early versus mid-reach). In the case of direction perturbations, the latencies of initial responses were identical across sub-conditions; the corrective responses evident in the aggregate data began within an average of 117 ms. In comparison, responses to distance perturbations were later and smaller in magnitude than responses to equivalent direction perturbations. There were also differences across sub-conditions: mid-point perturbations produced apparently earlier responses than early perturbations, and fast movements produced apparently earlier responses than

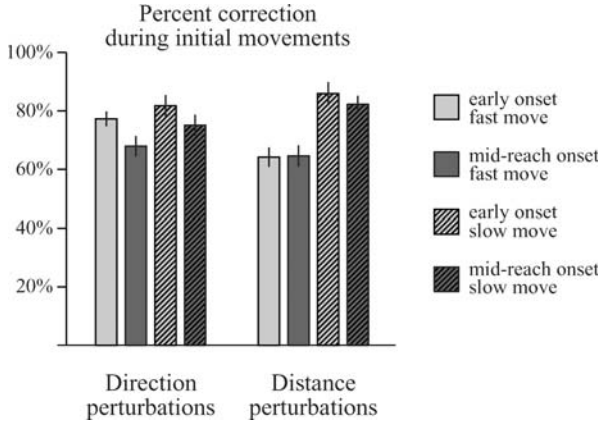

Fig. 4 The bars show the average amount of correction during subjects' initial movements in response to perturbations of different types, expressed as a percentage of the perturbation magnitude (so full compensation is 100%). The *vertical lines* depict standard errors

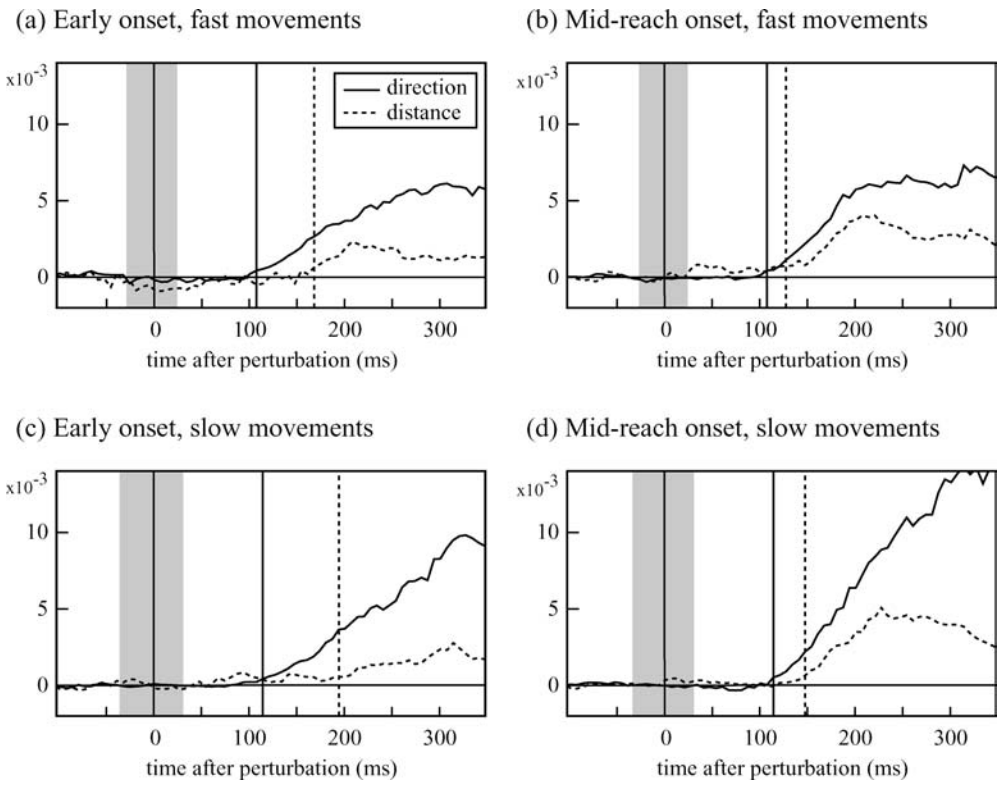

Fig. 5a–d Mean perturbation influence functions, averaged across subjects, for direction perturbations (*solid*) and distance perturbations (*dashed*). The *shaded areas* denote transition phase, during which the perturbation was gradually applied. The *vertical lines* mark the time of the earliest detectable response. **a** Early onset, fast movements; **b** mid-reach onset, fast movements; **c** early onset, slow movements; **d** mid-reach onset, slow movements

slow movements. Figure 6 plots response times for each of the direction and distance perturbation conditions.

Discussion

Response times for direction perturbations

Our finding of constant response times for direction perturbations is in agreement with the results of our previous study (Saunders and Knill 2003), which tested conditions identical to the direction perturbations used here. The only apparent discrepancy is the overall magnitude of response latencies: 117 ms in the present experiment versus 160 ms in our previous study. Most of this difference can be attributed to the increased sensitivity of our current analysis. The method of marking the point at which the perturbation influence function exceeds a fixed significance level has an inherent positive bias, which will depend on the noise in the measured influence functions. In the present experiment, response times were computed from the mean influence functions averaged across subjects, rather than separately for individual subjects. This, by itself, increases sensitivity and reduces bias. Additionally, because the mean influence functions are less noisy, less smoothing is required to compute stable estimates of when responses occurred. Decreasing the temporal extent of the smoothing kernel further reduces the amount of positive bias in the estimates. To measure the effect of these changes, we applied our previous analysis to the direction perturbation conditions in the current experiment, and found mean response times of 140 ms. The remaining discrepancy between the current results and our previous results can be explained by our recent discovery of an added one frame delay (17 ms) in the rendering of the virtual fingertip using the graphics card used for the previous experiment—a delay that we had not accounted for in the previous estimates. Thus, the present data replicate our previous findings, and provide further evidence that visual feedback from the hand is used for continuous control of movement direction. The shorter reaction

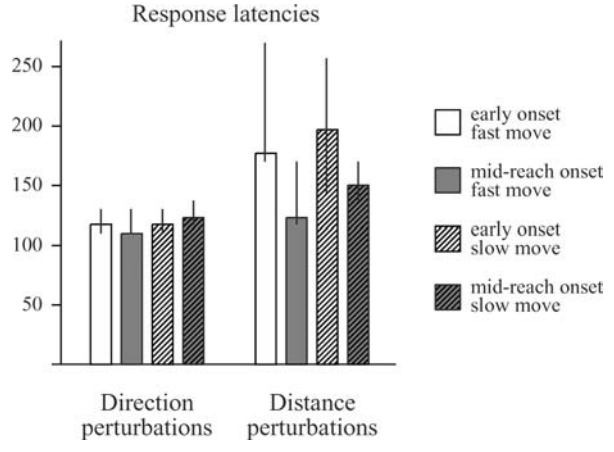

Fig. 6 The *bars* plot the response latencies: the time between the midpoint of perturbation transition and the earliest detectable response in the mean data. The *vertical lines* show 95% confidence intervals for the response time estimates, as determined by resampling

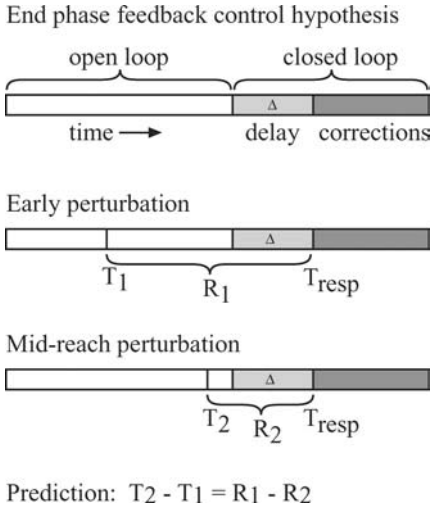

Fig. 7 Illustration of the predictions of an end-phase feedback control model. If movements have an initial open-loop phase, and perturbations are applied during this open-loop phase, the apparent response time would vary. Under this hypothesis, the difference in response latencies (*R1* and *R2*) would be expected to be equal to the difference in perturbation onsets (*T1* and *T2*)

time measured here, 117 ms, further supports this conclusion.

There is some ambiguity in the response times, due to the gradual application of the perturbations, which was spread over a period of 50–70 ms. Depending on how sensitive subjects are to the small conflicts present early in the transition period, or to the initial changes in velocity, a perturbed visual signal could have been available as much as 35 ms earlier than the "onset" time we used for analysis. Note, however, that the perturbations were introduced sub-linearly during the transition, so the initial conflicts were small. For example, during the first 15% of the transition, the conflicts remained less than 2 mm. Thus, it is likely that a perturbed signal effectively becomes available only at some time after the start of the transition. On the other hand, there is no reason to expect that the midpoint of the transition, which we treat as the "onset", would accurately correspond to the initial availability of a perturbed signal. Thus, our RT estimates should be interpreted as having additional uncertainty, beyond measurement noise, of up to 35 ms.

Continuous feedback for distance

If visual feedback about the distance of the hand from the target were not continuously used throughout a movement, but rather only during the end phase, one would expect the qualitative pattern of latency differences shown here. However, while the qualitative differences would be consistent with such an account, the quantitative differences are not. If feedback control commenced at some set point during movements, the difference in response latencies for early and mid-reach perturbations would be expected to be equal to the difference in perturbation onset times, as illustrated in Fig. 7. This prediction was not confirmed by the data. For slow movements, the measured response latencies differed by 40 ms, which is significantly less than the 108 ms difference in onset times for early and mid-reach perturbations ($p < 0.02$). The difference in response latencies for fast movements was also lower than predicted by the difference in onset times, 54 ms versus 88 ms ($p < 0.03$).

A further reason to doubt an end phase feedback explanation is that, in the conditions that produced the slowest responses, corrections began relatively early in the movements. In the case of early onset distance perturbations, responses occurred by the time the hand had moved approximately 70% of the distance to the target. Assuming even a minimal sensorimotor delay of 80 ms, these corrections would have to have been initiated while the hand was less than half-way to the target, prior to the point of peak hand velocity. The observation of such early perturbation responses, combined with the smaller-than-expected differences in response latencies, argues strongly against a model in which distance feedback is only used at the end of movements.

While the responses we observed to distance perturbations were slower overall than those to direction perturbations, the difference was much smaller than would have been expected by comparing the Sarlegna et al. and Saunders and Knill results. This could largely be due to the more sensitive analysis applied here to measure response times. Still, the data show significant differences in response times between direction and distance perturbations, as well as differences across distance perturbation subconditions. One possibility is that these effects result from differences in sensory uncertainty of visual feedback about finger direction and distance relative to the target. These two components correspond to different perceptual dimensions (speed versus motion direction) and have differing amounts of variability over time. To test whether sensory uncertainty could explain the differences in responses we observed, we simulated a feedback controller that optimally integrates sensory feedback with an ongoing estimate of hand state to generate online control signals for the hand, and compared its performance to the human data for the same experimental conditions.

Modeling

In Saunders and Knill (2004), we applied a feedback control model to the conditions of a motion direction feedback perturbation study to characterize the expected response of a feedback to a controller limited by noise in the visual coding of finger position and velocity. The model combines an optimal sensory estimator for finger position and velocity with an online control law derived from the minimum jerk criterion (Hoff and Arbib 1993) for generating motor commands to move the finger. The sensory component of the model is an adaptive Kalman

filter that maintains a running estimate of the state of the hand using noisy sensory information and outgoing motor signals. The filter continuously updates its estimate of hand state by combining its current estimate with the information derived from new sensory (visual) signals and an estimate of the motor command that was most recently sent out to the system. To account for sensorimotor delays, the filter runs in a delayed loop, and the delayed estimates of hand state are input into a forward model that uses estimates of the motor commands sent out during the delay period to "predict" the current state of the hand from the delayed estimate.

We adapted our previous model to the present experiments to test whether responses to direction and distance perturbations could be similarly explained in terms of sensory feedback noise. As before, we incorporated both signal-dependent noise on the control signal and state-dependent noise on the sensory feedback. The latter models speed-dependent changes in motion discrimination thresholds and eccentricity-dependent changes in spatial acuity. We also assumed that the sensory estimates of both hand position and velocity were temporally low-pass filtered in concordance with observed psychophysical behavior. Details are given in Appendix A. Most of the parameters for the sensory filters—the noise levels for visual estimates of the direction and speed of finger movement, and the noise levels for visual position estimates—were derived from published psychophysical data, as described in the Appendix. The one noise parameter that we have added to model the current data is designed to capture the uncertainty in position estimates that would be induced by temporal blurring of the visual input. The early visual coding of position is derived from temporally low-pass filtered images—effectively limiting positional acuity for a moving object like the finger by the spatial extent of the temporally blurred image of the object. We modeled this motion-blur induced uncertainty with a noise term whose standard deviation was proportional to the velocity of the finger. The total noise on position estimates is then given by the sum of a noise term that models static positional acuity (incorporating the effects of eccentricity on spatial sampling), and a term that models the positional noise induced by motion blur. The latter noise is zero in the direction perpendicular to the finger's motion and is proportional to finger speed in the direction of motion. The constant of proportionality was set to equate the motion blur noise to what would be expected for a visual signal blurred over a 40 ms window, equal to the time constant used to model the low-pass filtering of the visual signal (see Appendix A). The only free parameters in the model were the noise terms for the control signal and the sensorimotor delay. These were set to the same values used to fit the data in the previously reported study (Saunders and Knill 2004). Thus, all model parameters were fixed prior to running the simulations and none were fit to the current data.

Figure 8 shows example mean trajectories derived from simulating the model's performance in the fast movement condition of the experiment described here. The model's corrections are qualitatively similar to those shown by the human subjects. More illuminating are the results of applying the auto-regressive analysis to estimate perturbation influence functions for the model, shown in Fig. 9. The model replicates the qualitative results found in the human data. In particular, the influence functions are smaller for the distance perturbations than for the direction perturbations, and their

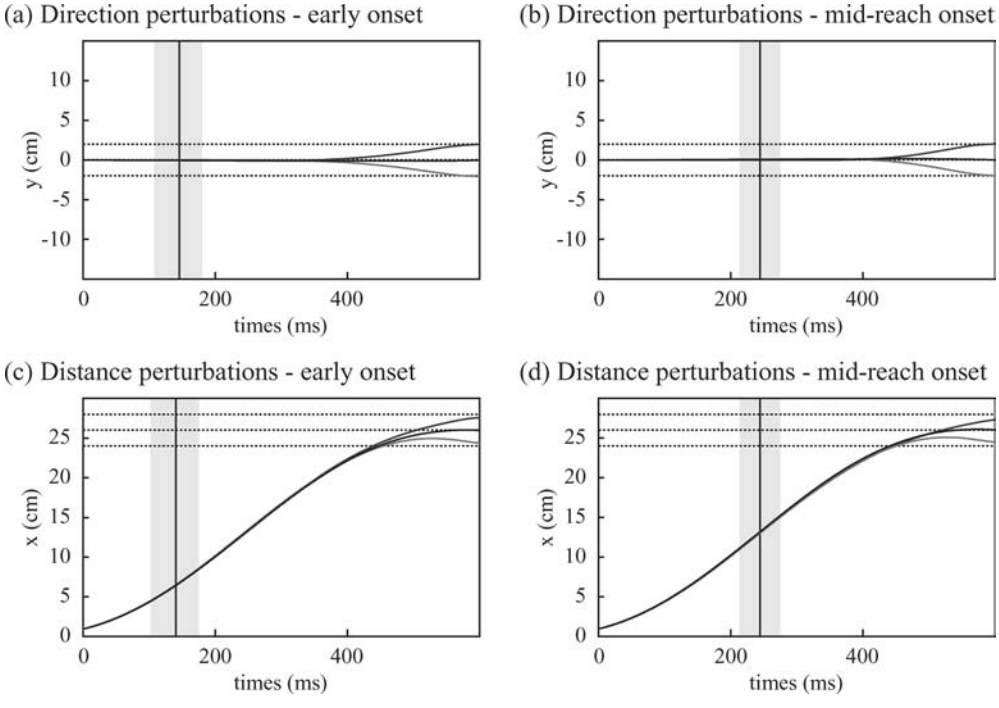

Fig. 8a–d Mean kinematic trajectories derived from simulating the performance of a feedback control model (see "Methods" and "Appendix A") on the experimental conditions tested here. **a** Direction perturbations, early onset; **b** direction perturbations, mid-reach onset; **c** distance perturbations, early onset; **d** distance perturbations, mid-reach onset

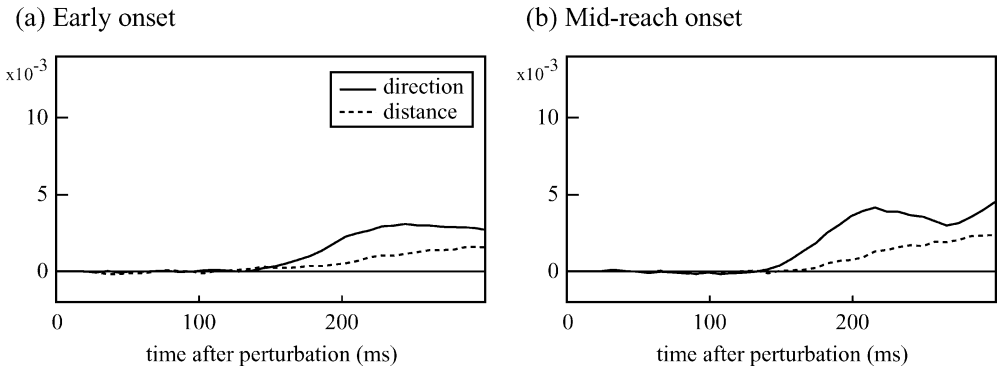

Fig. 9a–b Mean perturbation influence functions computed by applying the same auto-regressive analysis used on the human data to the simulated data of the model. **a** Early onset; **b** mid-reach onset

magnitude depends on the time of the perturbation in the same way that subjects' influence functions do. Most notably, the model shows apparently different delays in the different conditions of the experiment. These differences result principally from the difference in the magnitudes of the influence functions. In essence, the model's responses to late perturbations are stronger and hence reach significance sooner than its responses to early perturbations. As reported previously, the model deviates somewhat from subjects' behavior at the ends of the movements. This is likely due to the fact that the model did not include low-pass filtering of the motor output and supported the production of instantaneous changes in acceleration needed for fast, late corrections—an implausible behavior for a physical system.

In our model, we assume that feedback control is based on the system's estimate of the current state of the hand relative to the target. Under this assumption, perturbations evoke responses because they change the observed position and velocity of the hand relative to the target. It is also possible that subjects reacted to the conflict between observed and expected hand position on perturbed trials. For example, if subjects attempted to guide their hand along a predetermined kinematic trajectory, the relevant control signal would be the difference between observed and desired current position, which would not directly involve comparisons between hand and target. The present results cannot distinguish these possibilities, since perturbations alter both the visual position of the hand relative to the target and the visual position of the hand relative to its expected, physical position. Thus, there is some ambiguity as to the specific control variable that drives subjects' responses to the perturbations tested here. In another related study (Saunders and Knill 2004), we tested conditions that did allow greater comparison of control strategies. The observed results were inconsistent with a strategy of steering the hand along a predetermined trajectory, which is one natural strategy that would depend solely on the error between observed and expected state, rather than target-relative state. Note also that our model is not insensitive to conflicts between observed and expected hand state. However, its performance is not directly driven by differences between observed and expected hand state; rather, these variables affect performance via their contributions to the integrated estimate of instantaneous hand state.

Conclusions

The data show evidence for online feedback control of both movement direction and distance. For both dimensions, perturbations produced responses after relatively short latencies: 117 ms for direction perturbations, and 130–200 ms for distance perturbations. In the case of early onset perturbations, responses were observed by the time the hand was 50–70% of the distance to the target, implying that the corrections were initiated prior to the hand reaching peak velocity.

For direction perturbations, response latencies were the same for early and mid-reach perturbations, and for fast and slow movements. This replicates our previously-reported findings (Saunders and Knill 2003), and further supports the hypothesis that visual feedback is used continuously for control of reaching movements. The estimate of feedback delay computed from the direction perturbation data, 117 ms, is close to the reaction times observed for corrections in response to target perturbations (Pelisson et al. 1986; Prablanc and Martin 1992; Brenner and Smeets 1997). It also closely corresponds to the feedback delay we used to model results from another experiment, in which visual feedback was perturbed in ways that selectively altered motion and position feedback signals (Saunders and Knill 2004). These consistencies support our current, shorter estimate of feedback latency. As discussed above, we believe that our previous estimate of feedback latency had larger positive bias, and that the response time computed from the present data more accurately reflects the feedback delay.

The responses to distance perturbations were not invariant to onset time and movement speed. However, the differences in response times were smaller than would be expected under a hypothesis of end-phase feedback control. A control model that relies on an optimal Kalman filter for updating estimates of finger state from sensory feedback replicates all of the main patterns found in the data. The sensory noise parameters used for model simulations were drawn from

psychophysical data on position and velocity acuity and were the same as those used to model previous data (Saunders and Knill 2004). The one piece that we added to the model here was a position noise term that models the effects of motion blur on position estimates. The effects of motion-blur on position estimates effectively gives rise to the differences between the model's responses to direction and distance perturbations. The model simulations demonstrate that the differences across conditions observed in the human data are consistent with reasonable assumptions about the amount of perceptual noise present in the relevant control dimensions. Thus, despite some variations in measured reaction times, the results for distance perturbations remain consistent with the hypothesis of continuous feedback control of movement distance.

A conventional view is that visual feedback contributes to control of reaching movements only during an end-stage, when the hand is close to the target. As described above, our results clearly contradict this model. One reason that feedback might not be expected to contribute during the initial portion of a movement is because visual feedback about the state of the hand relative to the target would be poor. However, under the conditions tested here, the hand could still be visually compared to the target throughout the movement, even if subjects fixate the target. The quality of information about the hand state does strongly depend on its eccentricity relative to the target, which is one reason that we included eccentricity-dependent noise parameters in the model. In fact, the model performance demonstrates that varying amounts of sensory noise are sufficient to explain apparent differences in sensorimotor reaction times. Note that visual feedback from the hand is not strictly necessary to perform reaching tasks, as evidenced by our ability to pick up objects with our eyes closed. The accuracy of our model similarly decreases only modestly with the removal of visual feedback. Thus, it is not surprising that the initial influence of visual feedback would be small and difficult to detect, particularly if accuracy were used as a measure. However, the perturbation responses observed here demonstrate that the human sensorimotor system does take advantage of visual feedback about the hand throughout movements, even for the minimal feedback provided in our experiment (a small dot representing the fingertip).

The current results seem to conflict with those of Sarlegna et al. (2003), who found only weak responses to perturbations of visual feedback from the hand, and after a much longer latency of 470 ms. While our results demonstrate that there is indeed a difference between responses to perturbations affecting movement direction and distance, we still observed relatively fast responses to distance perturbations as well as direction perturbations. Response latencies ranged from 130–200 ms depending on perturbation onset and movement speed. While these latencies are longer than those observed for direction perturbations, in all cases they are much shorter than the latencies reported by Sarlegna et al. The cumulative amount of correction was also greater in the present study.

The analysis method used here largely accounts for the differences. Were we to have marked responses as beginning when the average finger trajectories from perturbation trials deviated significantly from those from unperturbed trials, we would have found similar response times to Sarlegna et al., in that responses would only have appeared at the end of movements. This overestimation results from the high variance in subjects' raw trajectories. Only by using the autoregressive model to account for correlations within trajectories is one able to obtain a high-enough signal-to-noise ratio to measure response times with any degree of sensitivity. Even with this improvement, our estimates are likely to be high, as demonstrated by the model results: when the same method was applied to the simulated data, response times were significantly higher than the sensorimotor delay built into the model. Furthermore, in the Sarlegna et al. study, the ''onset'' of perturbations was immediately after initial saccades, while the perturbations used here occurred during the main transport portion of movements, well after initial saccades. There could be some refractory period after saccades during which feedback cannot effectively be used, resulting in somewhat longer apparent feedback delays. Other methodological factors could account for the remaining differences in our findings (for example in proportional correction). In the Sarlegna et al. experiment, movements were directed to positions in space, such that the virtual target provided no impact resistance or haptic feedback. This is less natural than the pointing movements used in our experiment, which were directed toward virtual targets coincident with a physical tabletop, and might have somehow altered subjects' strategies. Finally, we marked the ends of movements differently than Salegna et al., in part because contact with the tabletop effectively stopped movements in our experiments. Small differences in this variable could easily have led to the differences we found in the proportional amount of correction for perturbations.

While the perturbation responses reported here clearly demonstrate that visual feedback contributes to online control, the data are ambiguous as to the specific perceptual dimensions that underlie this control. The distance perturbation conditions required subjects to adjust primarily the extent of movement, and the direction perturbation conditions required adjustments primarily in the direction of movement, so we interpret responses to these perturbations as reflecting feedback control of distance and direction respectively. However, there are many possible models for controlling these task dimensions, with differing relevant control variables, which these data cannot distinguish. For example, are corrections based on the current state of the hand relative to the target, or relative to some predetermined kinematic plan? Another question would be whether the representation of hand state used for control includes

only position and velocity, or higher derivatives as well. Such questions are beyond the scope of the present study, but are addressed to some extent in recent related paper (Saunders and Knill 2004).

We have also neglected to consider the role of depth information. Since movements were largely within a single plane, we have treated the task and data as two-dimensional. However, because the tabletop was slanted relative to the viewer, the visual information about the state of the hand relative to the target included a depth component. In particular, direction perturbations displaced the virtual finger in depth as well as in the frontal (projection) plane. For movements that are nearly planar, as we observed, the depth and frontal components of hand kinematics are highly correlated. When we re-analyzed the kinematic data for the direction perturbations using the vertical position of the hand in viewer-centered coordinates (in the image plane) instead of in tabletop coordinates, we found essentially equivalent results. The main complication is that the effective magnitude of the perturbation signal depends on perceptual sensitivity to both depth and frontal dimensions. If observers were much less sensitive to motion in depth than to frontal motion, as one would expect, direction perturbations would have been attenuated somewhat due to foreshortening. In fact, the model has this foreshortening effect built into its model of the sensory signal, which it assumed to be the movement of the finger in the image plane (see details in Appendix A). This quantitative ambiguity does not affect the distinction between direction and distance perturbation conditions. Control of movement direction might be based to varying extents on either 2-D or 3-D position and velocity, but regardless of the role of depth information, the task-relevant dimensions would be different than those for controlling movement extent (distance and speed).

In conclusion, the results provide evidence that visual feedback from the hand is used for online control of both the direction and distance of the hand. The results for direction perturbations strongly implicate continuous feedback control of movement direction. The results for distance perturbations are also consistent with continuous feedback control, particularly when the perceptual noise in relevant control dimensions is considered. While there is evidence that movement direction and distance are controlled independently during movement planning, we found no evidence for a qualitative distinction between these dimensions with regard to online feedback control. Rather, our results are consistent with a model in which available visual information is continuously integrated in an optimal manner throughout movement to estimate hand state and generate corrective adjustments to movements. Apparent differences in subjects' sensitivity to direction and distance perturbations are well accounted for by differences in the sensory noise that corrupts position and speed estimates in the principal direction of movement versus in the perpendicular direction.

Appendix A

For comparison to human data, we implemented a feedback control model that optimally integrates noisy position and velocity feedback signals over time, and simulated the performance of the model for the conditions of our experiments. The model was an expanded version of the feedback control system proposed by Hoff and Arbib (1993), in which the sensory front-end is augmented with a Kalman filter for integrating noisy sensory feedback with ongoing estimates of hand state. We modeled the dynamics of the hand as a linear system driven by a jerk signal (derivative of acceleration) that satisfies the minimum jerk principle, as proposed by Hoff and Arbib. The system was corrupted by two independent sources of added Gaussian noise, one with standard deviation proportional to the jerk signal and the other with a constant standard deviation. The general form for the system dynamics is given by

$$X^h_{t+1} = \mathbf{A}X^h_{t+1} + u_t + \sum_i \varepsilon_{i,t}\mathbf{C}_i u_t + \xi_t \tag{3}$$

where the state vector represents the two-dimensional position, speed and acceleration of the hand in a coordinate frame aligned with the table, as well as the two-dimensional position of the target on the table (we did not simulate the small movements in the perpendicular direction that subjects show). The x-axis was taken to be in the direction of the target from the original hand position and the y-axis was defined as the perpendicular direction within the plane of the table. $\epsilon_{i,t}$ is a standard normal random variable, $\mathbf{C}_t$ represents the proportional motor noise, and ξ_i is a constant noise source with covariance matrix $\mathbf{\Sigma}^{\xi}$. The state vector is given by

$$X^h_t = \left[x_t, x'_t, x''_t, y_t, y'_t, y''_t, Tx, Ty\right]^T \tag{4}$$

where Tx and Ty represent the position of the target. The state transition matrix, $\mathbf{A}$, is given by

$$\mathbf{A} = \begin{bmatrix} 1 & \partial & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & \partial & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \partial & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & \partial & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \tag{5}$$

where ∂ is the time step used to iterate the update equations in the simulations. Given an estimate of the state of the system at time t, the minimum jerk control law is given by (Hoff and Arbib 1993)

$$u_t = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{-60}{(D-t\partial)^3} & \frac{-36}{(D-t\partial)^2} & \frac{-9}{(D-t\partial)} & 0 & 0 & 0 & \frac{60}{(D-t\partial)^3} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{-60}{(D-t\partial)^3} & \frac{-36}{(D-t\partial)^2} & \frac{-9}{(D-t\partial)} & & \frac{60}{(D-t\partial)^3} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \hat{X} \tag{6}$$

In order to realistically model the visual feedback available to the human visual-motor system, we had to account for two properties of the visual feedback. First, the uncertainty in both position and motion information, as reflected in discrimination thresholds, varies as a function of the state of the hand. Positional acuity is inversely proportional to eccentricity; thus it is well modeled by a noise source with a standard deviation proportional to the 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 fitted well 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 standard deviation is proportional to the speed of the motion, and a constant noise component (Orban et al. 1985; De Bruyn and Orban 1988). Second, estimates of position and velocity are effectively low-pass filtered by the visual system. Psychophysical sensitivity to periodic modulations in both speed and direction of motion is fitted well by a model in which velocity estimates are filtered through a second-order filter with a time constant of 40 ms (Werkhoven et al. 1992). While similar estimates are not available in the literature for position estimates, we assumed that the same filter applies to position estimates.

In order to incorporate the state-dependent nature of visual noise and the temporal filtering of visual feedback, we found it convenient to incorporate the sensory parameters in an augmented state vector along with the hand and target state parameters. The full state vector was

$$X_t = [X_t^{hT}, X_t^{sT}]^T \tag{7}$$

where the sensory sub-vector is augmented by dummy variables used to implement the second-order filter,

$$X_t^s = \left[\tilde{v}_t, \tilde{v}'_t, \tilde{w}_t, \tilde{w}'_t, v_t, v'_t, w_t, w'_t\right] \tag{8}$$

with update equations

$$\begin{matrix} \tilde{v}_{t+1} = e^{-\delta/\tau}\tilde{v}_t + x_t, & \tilde{v}'_{t+1} = e^{-\delta/\tau}\tilde{v}'_t + x'_t, \\ v_{t+1} = e^{-\delta/\tau}\tilde{v}_t + e^{-\delta/\tau}v_t, & v'_{t+1} = e^{-\delta/\tau}\tilde{v}'_t + e^{-\delta/\tau}v'_t \end{matrix}^h \tag{9}$$

Similar equations apply to the w parameters which represent the y-position and the y-component of velocity. ∂ is the time step used in the simulation to iterate the update equations and $\tau = 40$ ms is the time constant of the filter. The update equations in Eq. 9 are modeled by appropriate entries in an augmented update matrix **A.**

In order to accommodate the state-dependent noise in the sensory signal, the update equation for the full system takes the form

$$X_{t+1} = \mathbf{A}X_t + \mathbf{L}_t\hat{X}_t + \sum_i \varepsilon_{i,t}\mathbf{C}_i\mathbf{L}_t\hat{X}_t + \sum_i \eta_i D_i X_t + \xi_t \tag{10}$$

The second term is the control law given in Eq. 6 above, the third term represents the signal-dependent motor noise, and the fourth term represents the state-dependent noise used to model state-dependent uncertainty in the sensory estimates of position and velocity.

The sensory signal is modeled by the equation

$$Z_t = \mathbf{H}X_t + \omega_t \tag{11}$$

where **H** is a matrix that simply peels off the temporally-filtered, noisy estimates of position and velocity (v_t,v_t',w_t,w_t'), and ω_t is a constant noise source with covariance matrix Σ^ω.

The optimal, adaptive Kalman filter for this system is given by the time update equations

$$\hat{X}_{t+1} = (\mathbf{A} + \mathbf{L}_t)\hat{X}_t + \mathbf{K}_t(Z_t - \mathbf{H}\hat{X}_t) \tag{12}$$

where the Kalman gain matrix, $\mathbf{K}_t$, is given by

$$\mathbf{K}_t = \mathbf{A}\Sigma_t\mathbf{H}^T(\mathbf{H}\Sigma_t\mathbf{H}^T + \Sigma_t^\omega)^{-1} \tag{13}$$

Σ_t is the error covariance matrix for the state estimate at time t. The error covariance matrix is updated according to the time update equation

$$\begin{aligned} =&\Sigma_{t+1}\mathbf{A}\Sigma_t\mathbf{A}^T - \mathbf{K}_t\mathbf{H}\Sigma_t\mathbf{A}^T + \sum_i \mathbf{C}_i\mathbf{L}_t\hat{X}_t\hat{X}_t^T\mathbf{L}_t^T\mathbf{C}_i^T \\ &+ \sum_i \mathbf{D}_i(\Sigma_t - \mathbf{K'}_t\mathbf{H}\Sigma_t)\mathbf{D}_i^T \\ &+ \sum_i \mathbf{D}_i\left[\hat{X}_t + \mathbf{K'}_t(Z_t - \mathbf{H}\hat{X}_t)\right] \\ &\times \left[\hat{X}_t + \mathbf{K'}_t(Z_t - \mathbf{H}\hat{X}_t)\right]^T\mathbf{D}_i^T + \Sigma^\xi, \end{aligned} \tag{14}$$

where **K**' is the modified Kalman gain (used to predict the state of the system at time t rather than to predict the state of the system at time $t+1$), given by

$$\mathbf{K'}_t = \Sigma_t\mathbf{H}^T\left(\mathbf{H}\Sigma_t\mathbf{H}^T + \Sigma_t^\omega\right)^{-1}. \tag{15}$$

The Kalman filter runs delayed by Δ ms, so the delayed state estimates are propagated forward through the state update equations using the remembered values of the motor commands, u_t, to arrive at an estimate for the current hand state.

In order to simulate the model, we had to set parameters for the visual noise in position and velocity.

We chose values that fit with observed psychophysical data, as described below. We fit the motor noise parameters and the sensorimotor delay by hand to subjects' data.

Position noise

We used results from two-point interval discrimination studies to set the parameters for eccentricity-dependent visual noise in position estimates (Burbeck 1987; Burbeck and Yap 1990; Whitaker and Latham 1997). The data from these studies is consistent with a Weber fraction of 0.05 in position estimates beyond several degrees away from the fovea for stimuli viewed for 250 ms. This value is invariant to a large number of properties of the target (Burbeck 1987; Toet and Koenderink 1988). Since subjects' finger movements were almost entirely along the axis between the starting position and the target (the x-axis), we modeled the standard deviation of the eccentricity-dependent noise in subjects' visual estimates of hand position (for 250 ms viewing) in tabletop coordinates to be

$$\begin{aligned}\sigma_x &= 0.05x + 0.05 \\ \sigma_y &= \sqrt{2}(0.05x + 0.05)\end{aligned} \tag{16}$$

where the $\sqrt{2}$ factor in the y-direction corrects for perspective foreshortening (subjects viewed the tabletop from an angle of $\sim$45°). The constant additive term models a minimum standard deviation in position estimates of 0.5 mm in the center of the fovea (3′ arc), but has little impact on the behavior of the model, since the hand is outside the fovea for most of the movement. We used a small angle approximation in treating the position in tabletop coordinates as proportional to visual angle.

Since position estimates are necessarily derived from temporally blurred images of the finger, the uncertainty in position estimates should scale with the velocity of the finger. This is accounted for by adding a third term to the position noise model,

$$\begin{aligned}\sigma_x &= 0.05x + 0.05 + kv_x \\ \sigma_y &= \sqrt{2}(0.05x + 0.05) + kv_y\end{aligned} \tag{17}$$

k is determined by the time constant for the motion blurring effect. For the simulations described in the paper, we assumed a time constant of 40 ms, equivalent to the time constant of the low-pass filter used to model temporal blurring of the position and velocity signals. For simplicity, we assumed uniform blurring over a 40 ms window. This leads to a constant, $k = 0.00116$ s (a uniform distribution over the interval $[0,x]$ has a standard deviation $s = 0.29x$). In order to parameterize the noise model, we multiplied the standard deviations by a factor of 10.2 so that a sensory version of the Kalman filter run on stimuli for 250 ms gave estimates with the standard deviations listed above (thus matching the static position acuity thresholds measured psychophysically).

The constants of proportionality in Eq. 17 determine the appropriate elements of the $\mathbf{D}_i$ matrices that model the state-dependent noise in position. The constant noise term is the standard deviation of the corresponding elements of the constant noise vector, ξ. The noise in visual estimates of target position was modeled as being constant, since the target is stationary and assumed to be fixated during movement. We assumed a standard deviation of 0.5 mm (equivalent to a visual angle of 3′ of arc).

Modeling motion noise

Results from speed and direction discrimination studies show a somewhat more complicated behavior than position perception. Up to speeds of 64 deg/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 (Mateef 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 fitted well by a mixed constant and proportional noise model in which the standard deviation of visual estimates of speed is given by

$$\sigma_{v_x} = 2 \text{ deg/s} + 0.08v_x \tag{18}$$

where we have assumed that the speed of motion is equivalent to its velocity along the x-axis (approximately true for most of the duration of subjects' movements). Using a small angle approximation to convert this to units of distance along the tabletop (assuming an average viewing distance of 52 cm) gives

$$\sigma_{v_x} = 1.8 \text{ cm/s} + 0.08v_x \tag{19}$$

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 eight. For the standard deviation of velocity estimates perpendicular to the direction of motion, therefore, we have

$$\sigma_{v_y} = 0.25 \text{ deg/s} + 0.01v_x \tag{20}$$

Converting this to tabletop coordinates and adjusting for perspective foreshortening, we have for the standard deviation

$$\sigma_{v_y} = 0.35 \text{ cm/s} + 0.014v_x \tag{21}$$

Again, we approximated the principle direction of motion to be in the x-direction. We scaled the parameters by a constant factor of 14.3 so that a sensory version of the Kalman filter run on static stimuli for 500

ms would give velocity estimates with the standard deviations listed above. The constants of proportionality in Eqs. 20 and 21 determine the appropriate diagonal elements of the $\mathbf{D}_i$ matrices that model the state-dependent noise in position. The constant noise terms are the standard deviations of the corresponding elements of the constant noise vector, ξ.

Other model parameters

The motor noise used to fit the human data was set to

$$\sigma_u = 1.5 \text{ cm/s}^2 + 0.05u \quad (22)$$

The coefficient of proportionality (0.05) for the proportional noise determined the diagonal elements of the C_1 and C_2 matrices corresponding to the motor commands (the jerk signals in x and y), while the constant noise term determined the standard deviation of the same elements of the constant noise vector, ξ. We simulated the model with a sensorimotor delay of 115 ms, for a duration of 750 ms, iterating the system with a time step of $\partial = 2$ ms. Perturbations were induced at $t = 214$ ms for the early perturbation and $t = 326$ ms for the mid-reach perturbation. The duration of movement was set based on estimates for the average time that subjects' fingers came to rest for the slow movement conditions, and the time of perturbation was set based on the average time after the start of movement that subjects' fingers reached a quarter or half of the distance to the target.

References

Brenner E, Smeets JBJ (1997) Fast responses of the human hand to changes in target position. J Motor Behav 29(4):297–310

Burbeck CA (1987) Position and spatial frequency in large-scale localization judgments. Vision Res 27:417–428

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

Chieffi S, Allport DA (1997) Independent coding of target distance and direction in visuo-spatial working memory. Psych Res 60(4):244–50

De Bruyn B, Orban GA (1988) Human velocity and direction discrimination measured with dot patterns. Vision Res 28:1323–1336

Gordon J, Ghilardi MF, Ghez C (1994) Accuracy of planar reaching movements. I. Independence of direction and extent variability. Exp Brain Res 99(1):97–111

Hoff B, Arbib MA (1993) Models of trajectory formation and temporal interaction of reach and grasp. J Motor Behav 25:175–192

Krakauer JW, Pine ZM, Ghilardi MF, Ghez C (2000) Learning of visuomotor transformations for vectorial planning of reaching trajectories. J Neurosci 20(23):8916–24

Levi DM, Klein SA, Yap YL (1988) Weber's law for position: unconfounding the role of separation and eccentricity. Vision Res 28:597–603

Mateef 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 40:409–416

Messier J, Kalaska JF (1999) Comparison of variability of initial kinematics and endpoints of reaching movements. Exp Brain Res 125(2):139–52

Orban GA, Van Calenbergh F, De Bruyn B, Maes H (1985) Velocity discrimination in central and peripheral vision. J Opt Soc Am A 11:1836–1847

Pelisson D, Prablanc C, Goodale MA, Jeannerod M (1986) Visual control of reaching movements without vision of the limb. II. Evidence of fast unconscious processes correcting the trajectory of the hand to the final position of a double-step stimulus. Exp Brain Res 62(2):303–11

Prablanc C, Martin O (1992) Automatic control during hand reaching at undetected two-dimensional target displacements. J Neurophysiol 67(2):455–69

Sainburg RL, Lateiner JE, Latash ML, Bagesteiro LB (2003) Effects of altering initial position on movement direction and extent. J Neurophysiol 89(1):401-415

Saunders JA, Knill DC (2003) Humans use continuous visual feedback from the hand to control fast reaching movements. Exp Brain Res 152(3):341–352

Saunders JA, Knill DC (2004) Visual feedback control of hand movements. J Neurosci 24(13):3223–3234

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

Toet A, Koenderink JJ (1988) Different spatial displacement discrimination thresholds for Gabor patches. Vision Res 28:133–143

Turrell Y, Bard C, Fleury M, Teasdale N, Martin O (1998) Corrective loops involved in fast aiming movements: effect of task and environment. Exp Brain Res 120(1):41–51

Werkhoven P, Snippe HP, Toet A (1992) Visual processing of optic acceleration. Vision Res 32:2313–2329

Whitaker K, Latham K (1997) Disentangling the role of spatial scale, separation and eccentricity in Weber's law for position. Vision Res 37:515–524