Exp Brain Res (2003) 152:341–352
DOI 10.1007/s00221-003-1525-2

RESEARCH ARTICLE

Jeffrey A. Saunders · David C. Knill

Humans use continuous visual feedback from the hand to control fast reaching movements

Received: 16 March 2002 / Accepted: 8 May 2003 / Published online: 6 August 2003

Abstract How visual feedback contributes to the on-line control of fast reaching movements is still a matter of considerable debate. Whether feedback is used continuously throughout movements or only in the "slow" end-phases of movements remains an open question. In order to resolve this question, we applied a perturbation technique to measure the influence of visual feedback from the hand at different times during reaching movements. Subjects reached to touch targets in a virtual 3D space, with visual feedback provided by a small virtual sphere that moved with a subject's fingertip. Small random perturbations were applied to the position of the virtual fingertip at two different points in the movement, either at 25% or 50% of the total movement extent. Despite the fact that subjects were unaware of the perturbations, their hand trajectories showed smooth and accurate corrections. Detectable responses were observed within an average of 160 ms after perturbations, and as early as 60% of the distance to the target. Response latencies were constant across different perturbation times and movement speed conditions, suggesting that a fixed sensori-motor delay is the limiting factor. The results provide direct evidence that the human brain uses visual feedback from the hand in a continuous fashion to guide fast reaching movements throughout their extent.

Keywords Reaching · Visual feedback · Perturbation · Motor control · Virtual reality

Introduction

A large body of research has emphasized the need for vision of the hand for accurate reaching. Completely removing vision of the hand has been observed to cause greater variability of endpoints (Keele and Posner 1968; Carlton 1981; Zelasnik et al. 1983; Hay and Beaubaton 1986), biases toward undershooting target distance (Prablanc et al. 1979a; Jeannerod 1984; Prablanc and Pélisson 1990), and greater curvature in reaching paths (Sergio and Scott 1998; Goodbody and Wolpert 1999). These effects are generally thought to result from lack of information about the initial hand position for movement planning, and/or visual feedback from the hand during the final phase of movement for endpoint control. The question of whether or not the human brain uses continuous visual feedback throughout reaching movements remains hotly debated (Desmurget and Grafton 2000).

The classical view of reaching movements is that they consist of two phases: a fast initial phase that is primarily ballistic, and a slower adjustment phase under the guidance of sensory feedback (Jeannerod 1988). Two basic observations, one computational and one empirical, seem to lend support to this view. First, the minimum delay needed for sensory information to affect the physical movement of the hand [approximately 80–100 ms (Jeannerod 1988)] would seem to render feedback useless for the short duration of many naturally occurring reaching movements, which are on the order of 350–750 ms. Second, the fast and slow phases appear clearly in the transport kinematics of the hand, which show an initial bell-shaped velocity profile followed by an extended tail with, often, a number of re-accelerations.

Recent modeling work has challenged the computational concern by showing that internal forward models (Hoff and Arbib 1991; Jordan 1996; Miall and Wolpert 1996) can compensate for sensory delays in feedback processing by using efferent information to predict the current state of the hand from delayed sensory information (for a good review, see Desmurget and Grafton 2000). Several behavioral studies have provided evidence for the existence of internal forward models for prediction (Wolpert et al. 1995, 1998; Blakemore et al. 1998; Witney et al. 1999). Other studies, which measure corrective

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
Tel.: +1-585-2755400
Fax: +1-585-2713043

responses to perturbations in the location of a visual target, support the hypothesis that the early phases of reaching movements are not entirely ballistic. When target position is changed during a reaching movement (either perceptibly or subliminally during orienting saccades), subjects smoothly and accurately correct their hand movements to compensate for the changes, despite the fact that they are unaware of the perturbations (Soechting and Lacquaniti 1983; Goodale et al. 1986; Pélisson et al. 1986; Prablanc and Martin 1992).

The results of target perturbation studies are generally taken to implicate some form of on-line control of reaching; however, other possibilities exist. Corrections could theoretically derive from changing the final equilibrium point of the arm (Flanagan et al. 1991) or a linear superposition of the trajectories planned for the initial target position and for the perturbed position (Flash and Henis 1991). Accepting the on-line control explanation still leaves open the question of whether and how visual feedback from the hand contributes to the control. Corrections evidenced in target perturbation studies could be based on proprioceptive feedback loops, visual feedback loops, internal feedback loops that rely on feedforward estimates of hand location and velocity (efferent copy), or some combination of the three. Available psychophysical evidence suggests that sensory feedback from the hand is not necessary for on-line corrections. Target perturbations have a similar effect with or without vision of the hand (Goodale et al. 1986; Prablanc and Martin 1992) and even when proprioceptive feedback is further eliminated, subjects are able to compensate for target perturbations (Bard et al. 1999).

In order to test the contribution of visual feedback to the on-line control of reaching movements, one must somehow manipulate the feedback and measure the resulting changes in movement kinematics. This has traditionally been done by comparing movements with and without vision of the hand (closed-loop vs. open-loop). Most early studies showed that completely eliminating vision of the hand affects both accuracy and movement kinematics, e.g., leading to longer movement times (Prablanc et al. 1979a; Carlton 1981; Jeannerod 1984; Prablanc and Pélisson 1990; Jakobson and Goodale 1991; Berthier et al. 1996). Researchers have isolated three potential reasons for these effects: vision of the static position of the hand prior to movement onset may be important for accurate planning, visual feedback from the hand may contribute to on-line control continuously throughout a movement, or visual feedback from the hand may contribute to on-line control only in the end-phase of a movement, when the hand and target both appear in or near the fovea.

A number of recent psychophysical experiments strongly support the importance of visual information about the position of the hand prior to movement onset for motion planning. In open-loop trials, subjects' performance improves significantly when vision of the initial position of the hand is made available prior to movement onset (Prablanc et al. 1979b; Elliott and Madalena 1987; Rossetti et al. 1994; Desmurget et al. 1997; Vindras et al. 1998). In regards to the feedback question, Carlton (1981) showed that while removal of visual feedback from the hand throughout a movement significantly degraded movement accuracy, elimination of visual feedback from the hand in the first 75% of a movement had no effect on subjects' performance. This result is commonly cited as evidence that the contribution of visual feedback to on-line control is limited to the final phase of the movements. In another study consistent with this conclusion, Rizzo and Darling (1997) showed that a subject with cerebral tunnel vision (with only foveal vision) performed a simple reaching task with an accuracy and with kinematics (e.g., movement duration) that were within normal bounds. They concluded that visual feedback from the hand early in movements, when the hand is in the peripheral visual field, does not contribute significantly to on-line control.

Relating the results of visual deprivation studies to questions about the role of visual feedback in movement control has a number of difficulties. The dependent measures (accuracy, duration of movement, etc.) are necessarily gross and therefore relatively insensitive measures of the effects of visual feedback; thus, negative effects must be interpreted with a great deal of caution. More importantly, changing viewing conditions changes task demands, which can significantly affect movement kinematics (Zelasnik et al. 1983; Jakobson and Goodale 1991; Connolly and Goodale 1999), and often changes global visual conditions, such as contrast or the availability of reference frame information for target position (Connolly and Goodale 1999). Connolly and Goodale controlled for these factors and found little difference in pointing accuracy between open-loop and closed-loop conditions, when continuous vision of the target was available (Connolly and Goodale 1999). They did, however, find significant kinematic differences early in the pointing movements between open-loop and closed-loop trials. This suggests a possible role of visual feedback early in the movements. As Connolly and Goodale note, however, the kinematic differences may have resulted from differences in motor plans, since they used a blocked design.

The conclusions supported by studies comparing open and closed loop pointing and prehension movements are somewhat limited. They clearly implicate a role for visual information about the static position of the hand prior to motion onset in motion planning. They also support the hypothesis that visual feedback contributes to motor control in the final phase of movements, when the hand is near the target. What is left open by the existing data is the question of whether or not the visuo-motor system uses visual feedback from the hand early in hand movements. The experiment reported here was designed as a direct test of the hypothesis that the visuo-motor system uses visual feedback to control movements throughout a movement.

Our strategy was to directly perturb the visual feedback from the hand during reaching movements in a virtual environment. Subjects viewed both a virtual target and a virtual representation of their moving fingertip. At different points during the reach, we perturbed the position of the virtual fingertip in a direction perpendicular to the movement path. By making perturbations small and smooth we were able to insure that subjects did not consciously detect them. Unlike earlier perturbation studies, in which the position of the target is changed subliminally during a movement, perturbations of the visual signal from the hand unambiguously probe the role of visual feedback in control, since corrective responses would necessarily depend on the visual signal from the virtual fingertip.

The experimental hypothesis is that subjects' hand movements are under the control of continuous visual feedback throughout a movement. In any implementation of feedback control, the speed of corrective response would be limited by some minimal sensori-motor delay. A continuous feedback controller would be one for which this is the only source of latency in response. In the simplest form of the continuous control hypothesis, the sensori-motor delay is approximately constant; thus, the hypothesis predicts that the response latencies for perturbation corrections to appear in movement kinematics would be invariant to the time of perturbation. Moreover, because the delay in the hypothetical feedback signal should not scale with task parameters such as the speed or length of a movement, the hypothesis predicts that the delay will be constant over variations in these parameters as well. The amount of correction could vary greatly depending on the particular control strategy and task parameters (as well as on arm dynamics), but for continuous control the delay before initial corrections would remain constant.

An alternative is that continuous visual guidance is applied only at the end of a reach. There are a number of possible reasons for this limitation; effective use of visual feedback might require the hand to be moving slowly, or to be positioned near the target, or in foveal vision (Paillard 1982). Such limits on visual feedback would predict that subjects' responses to perturbations would depend on finger position along a reach. Consequently, early perturbations would result in larger response latencies than later perturbations, and slow movements would result in larger latencies than fast movements.

To distinguish between these two hypotheses (continuous vs. end-stage feedback control), we measured responses to hand perturbations at two different times during a movement and for movements with two different speeds.

Materials and methods

Design and conditions

Figure 1 shows a schematic of the virtual environment used in the experiment. Subjects viewed stereoscopically rendered images reflected through a horizontal mirror, which also served to prevent view of the hand. Visual feedback about the position of a subject's hand was provided by a small sphere that moved in real-time along with the subject's actual finger—a virtual fingertip. A subject's task was to move their finger to touch a sequence of virtual targets, which were rendered to appear to be at the same depth as a tabletop placed under the mirror.

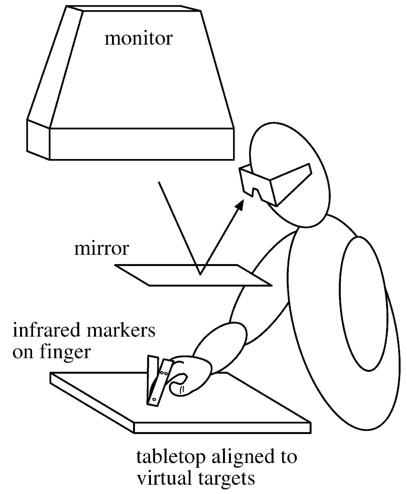

Fig. 1 Illustration of the experimental apparatus. Subjects were presented with stereo images using shutter glasses and a monitor viewed through a mirror. The mirror prevented view of the hand and allowed artificial visual feedback to be presented in the subjects' manual workspace. Subjects moved their 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 control a sphere representing the subject's fingertip (virtual finger)

Figure 2 illustrates the main stimulus manipulation used in the experiment. The virtual fingertip moved in one of three ways—coincident with a subject's fingertip throughout the movement, coincident for a fixed proportion of the movement and then perturbed "up" in a direction perpendicular to the path between the initial hand position and the target, or conversely perturbed "down" at that point. The presence and direction of perturbations were randomized across trials. Perturbations were small (2 cm) and smooth, appearing as a blurred step that extended over 5 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.

Perturbations occurred at one of two points along the movement. In an Early perturbation condition, perturbations occurred when the finger had moved 25% of the distance to the target, while in a Mid-reach perturbation condition, perturbations occurred when the position of the fingertip had moved 50% of the distance to the target. These two perturbation time conditions were randomly interleaved.

Movement speed was manipulated by giving subjects feedback to train them to complete their movements within ±25% of a goal time. Two movement speed conditions were tested in separate experimental sessions. In the Fast condition, subjects were trained to reach the target within ~450 ms of initiation, and in the Slow condition, subjects were trained to reach the target within ~600 ms of initiation. These speeds spanned the range of what pilot subjects subjectively reported as natural movement speeds for the distances used.

Apparatus and display

Visual displays were presented in stereo from 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

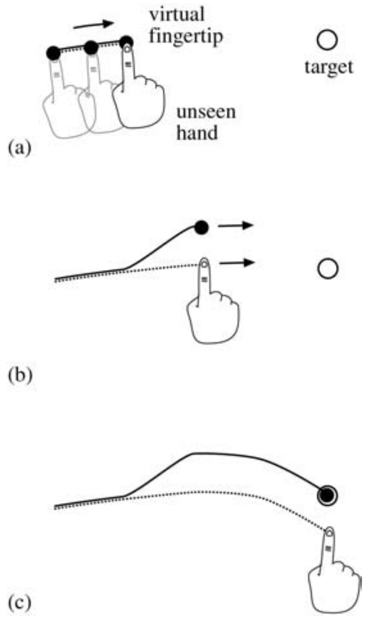

Fig. 2a–c Illustration of a perturbed trial. **a** At the start of movement, the virtual fingertip (*solid circle*) coincided with the actual location of the subject's unseen finger. **b** During the movement, a perturbation was gradually added to the position of the virtual fingertip, so that it was offset relative to the actual finger. **c** To successfully reach the target with the perturbed virtual fingertip, the subject would have to compensate for the perturbation, so that the actual finger ended at a position displaced from the target (shifted in the direction opposite the perturbation)

and right eye views were accurate perspective renderings of the simulated environment. In stereo mode, the monitor had a resolution of 1024×768 and a refresh rate of 120 Hz, or 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 inter-ocular cross-talk. Targets were aligned with an unseen table (~55 cm from the eyes), which provided consistent haptic feedback when a subject's finger touched the virtual target.

An Optotrak 3020 system recorded the time-varying position of a subject's finger at 480 Hz, and the recorded position was used to dynamically render a 1-cm diameter sphere representing the subject's 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. For each refresh cycle of the display, the most recently recorded finger position was used to update the rendered 3D position of the virtual fingertip. Visual feedback based on marker readings was slightly delayed due to the latency of monitor refresh (60 Hz per eye). Though the delay would be less than ~16 ms, at maximum hand velocity during a reach (~75 cm/s in our experiment) even this small delay could lead to errors in position as large as 1 cm. To compensate for the delay in visual presentation, the rendered position of the virtual finger was extrapolated from the latest recorded position, using previous recent positions to estimate the distance moved in a refresh cycle. 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.

The Optotrak sampled finger position at 480 Hz, but due to difficulty synchronizing with the 60-Hz display loop, not all samples were recorded. Missing frames occurred on about 2% of display loops, and represent about 0.3% of total possible samples. Gaps in the time series consisting of three or fewer missing frames were filled in by interpolating between adjacent frames. Thus, the effective sampling rate for finger position varied between 120 and 480 Hz.

Calibrating the virtual environment required knowing 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 3D 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 a rendered dot moved appropriately. Calibration was deemed acceptable if deviations were less than 1–2 mm. Otherwise, the calibration procedure was repeated.

For perturbed trials, a 2-cm offset was added between the actual finger and virtual finger during the course of the movement. The direction of displacement was perpendicular to the line connecting the initial finger location and the target location and was within the plane of the table. Thus, in the coordinate frame defined by the movement direction (X-axis in the target direction, Z-axis normal to the table), the perturbations were ±2 cm in the Y-axis direction, corresponding to about 1.7° of visual angle. Perturbations were added gradually over a 5-cm range along the distance to the target, as a function of the position of the finger along the axis between initial and target locations (X-axis). The gradation followed a logistic function, with a 50% point at the center of the 5-cm range and with 25% and 75% points 1 cm away from the center. In the Early perturbation condition, transition regions were centered at the point 25% of the distance to the target, and in the Mid-reach perturbation condition, transition regions were centered at the point 50% of the distance to the target. For simplicity, we will refer to the perturbation as "occurring" at the midpoint of its gradual introduction.

Subjects' direction of gaze was recorded throughout the experiment using an ASL model 5000 eye tracker, sampling at 60 Hz, and analyzed offline to determine relative eye-hand coordination. The eye-tracker was mounted on the chin-rest in which subjects rested their heads during the experiment.

Procedure

A subject's task 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 0.5 cm and continued for 1 s. The recording period was followed by a 750-ms delay with a blank screen, after which a new target appeared initiating the next trial. This resulted in trials that lasted approximately 2.5 s.

The targets varied in location within two ellipses on the table, which were 6 cm wide, 12 cm tall and separated by 28 cm. The target position for a trial was chosen relative to the end location of the previous trial. For the experimental trials, the target distance was always 28 cm, while the direction from the starting position to the target randomly varied between −15° and +15° relative to the horizontal midline of the table, subject to the constraint that the endpoint lay within the target ellipse. 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.

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. The goal times were 450 ms in the Fast condition, and 600 ms in the Slow condition. These two movement speed conditions were blocked and tested in separate experimental sessions. 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 ("x" through target indicating too fast, "+" indicating too slow). For purposes of feedback, movement was considered completed when both the speed of the virtual finger was less than 10 cm/s and its X-Y projected position was within a 2-cm-wide by 4-cm-high ellipse centered around the target. Note that because the spatial criterion is large in the y direction, subjects would not have to compensate for perturbations in order to receive positive feedback about timing. Prior to the main experiment blocks, subjects performed 40 unperturbed practice trials to familiarize themselves with the timing constraint.

Two thirds of the trials in each session were baseline trials, and one third were perturbed. Perturbed and baseline trials were randomly intermixed, with the constraint that at least one baseline trial separated any two perturbed trials. This constraint was included to reduce the potential for interference across sequential perturbed trials. For the same reason, we also excluded from analysis all baseline trials that were preceded by perturbed trials. This reduced the amount of baseline trials available for analysis by half.

Prior to analysis, data were filtered to remove irregular trials. These included trials with incomplete data due to occluded markers, trials with mistimed recording due to a false start, trials in which subjects began moving before the target appeared, and trials in which movement time was much shorter or longer than the goal time (±40%).

Subjects participated in two experimental sessions on separate days, one with feedback consistent with Fast movements and one with feedback consistent with Slow movements, with order counterbalanced across subjects. The experimental sessions began with calibration of the eye tracker and virtual environment, followed by practice trials to familiarize with movement speed, and then two blocks of experimental trials separated by a brief break. Subjects performed 432 trials per block, which yielded a total of approximately 72 total trials per perturbed condition (perturbation time × direction of perturbation × movement speed) and 576 baseline unperturbed trials for each movement speed.

Subjects

Twelve naïve, paid subjects participated in the experiment. 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

Movement kinematics

Figure 3 shows some sample finger trajectories for baseline unperturbed trials (a) and trials with either positive or negative perturbations (b and c). The raw finger position data from each trial are expressed in a coordinate frame aligned with the table and movement direction, in which the X-axis is the direction between the starting and target positions, the Y-axis is perpendicular to this in the plane of the tabletop (away from the subject), and the Z-axis is normal to the tabletop. In this coordinate frame, perturbations were along the Y-axis, and target locations were the same for all trials within a condition. The plots clearly show that subjects corrected for the perturbations. The endpoints for perturbed trials were shifted, by an amount approximately equal to the magnitude of the perturbations, in the opposite direction, i.e., downward for positive perturbations (b) and upward for negative perturbations (c).

(a) Unperturbed

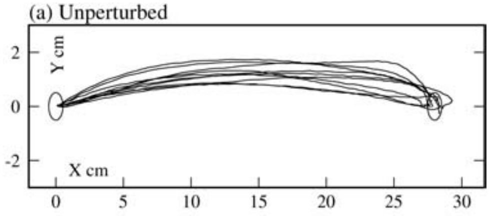

(b) Perturbation upward

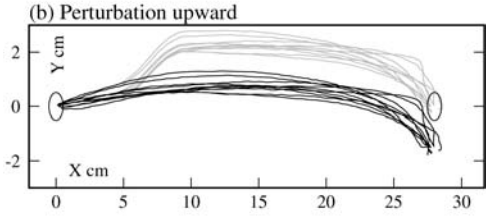

(c) Perturbation downward

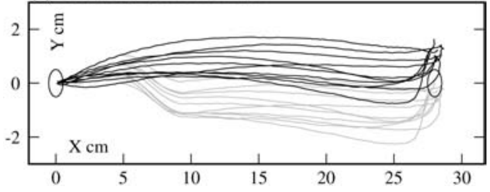

Fig. 3a–c Sample finger trajectories from a single subject's data. The graphs plot X-Y position, as projected onto the table top, in coordinates that have been aligned with the movement direction, such that the normalized position of the target is constant [oval at point (28,0)]. Positive Y values correspond to increasing distance from the subject. The three graphs show **a** baseline trials with no perturbation, **b** trials with positive perturbations, and **c** trials with negative perturbations. For the perturbed trials, the position of the virtual fingertip has also been plotted in gray. The *downward* and *upward curves* in b and c indicate corrective responses to the perturbations

Mean trajectories were computed by parameterizing the normalized X-Y coordinates by arc length, and then averaging across trials in the same condition. Fig. 4 shows X-Y plots of average finger movements for two sample subjects. Perturbed and unperturbed trials with corresponding movement speed and perturbation onset are plotted together on the same graphs. Corrective responses are clearly evident in the averaged finger position plots as well. Trajectories from trials with positive and negative perturbations diverge as the finger approaches the target, and average endpoints are shifted by the appropriate amount.

As is commonly found with lateral pointing movements, an overall curvature is evident in the mean trajectories, which varied in magnitude across subjects.

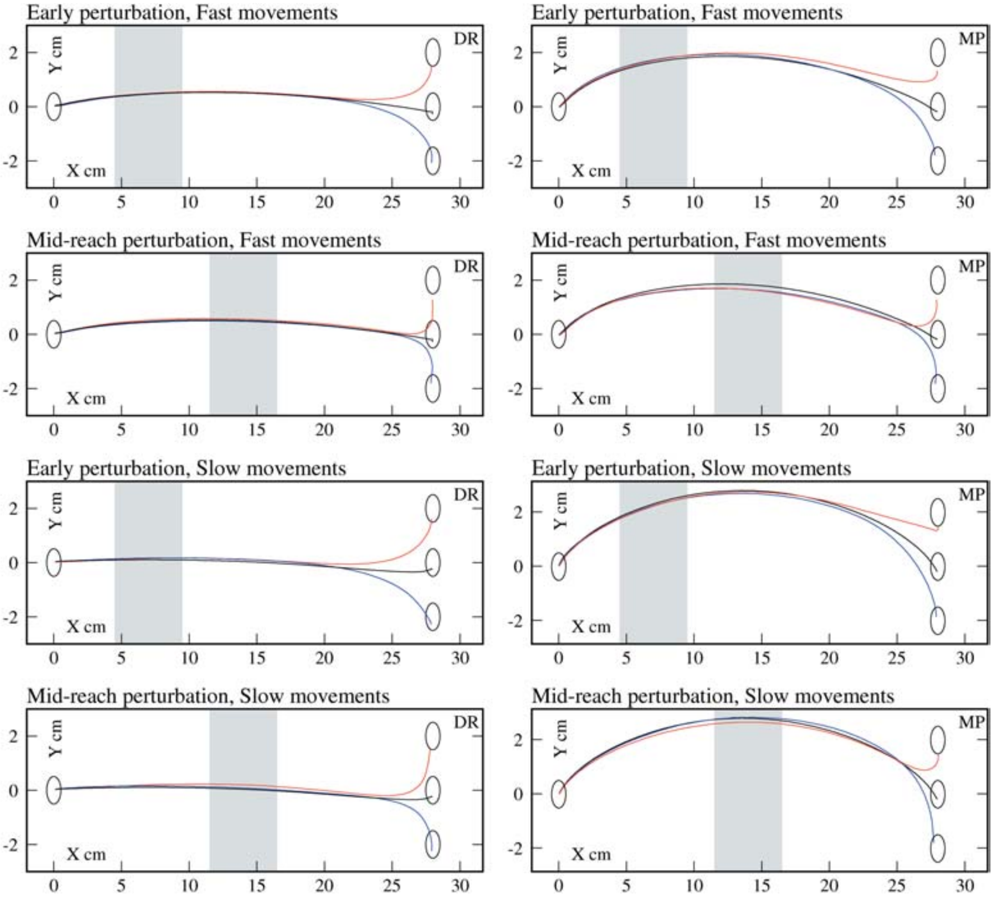
Early perturbation, Fast movements

Mid-reach perturbation, Fast movements

Early perturbation, Slow movements

Mid-reach perturbation, Slow movements

Fig. 4 Mean X-Y finger trajectories in each of the four perturbation time × movement speed conditions, for two sample subjects. The three mean trajectories in each graph correspond to matching trials with positive perturbations (*blue*), negative perturbations (*red*) and no perturbations (*black*). The ovals show endpoint locations that correspond to successfully moving the virtual fingertip to the target. The *gray regions* in each of the graphs show the regions of the smooth perturbation onset

Subjects appeared to be rotating their arm around their body when transporting the hand, to varying extents, resulting in trajectories that curve outward from the subject.

The final positions of the virtual finger were very close to the target, indicating that by the end of the trials subjects had fully compensated for the perturbations. However, some component of the total corrections could be secondary movements, as opposed to on-line adjustments. To distinguish these, we computed the average position of the finger relative to the target at the end of initial movements, which was defined as being when both the Z component of normalized finger position (height above the table) and finger speed fell below threshold values. Overall, subjects' corrections during initial movements compensated for an average 60% of the perturbation displacements, corresponding to a mean endpoint error of 0.8 cm. The amount of initial correction was affected by both perturbation time ($F_{(1,33)}$=40, P< 0.001) and movement speed ($F_{(1,33)}$=66, P< 0.001), with the proportion of correction being larger for slow movements than fast movements, and larger for early perturbations than for late perturbations (Table 1).

Movement kinematics are illustrated in Fig. 5. The plots show X and Y normalized coordinates of finger position separately as a functions of time (*left*), and the X and Y components of finger velocity as functions of time (*right*), averaged across trials from a sample subject and condition (no time normalization). Corrective responses can be seen in the Y components of the data, which for perturbed trials diverge away from zero. The X components appear similar across conditions, and there was no indication that the perturbations affected velocity profiles or timing of peak velocity.

Table 1 Proportion of corrections during initial movement

	Fast	Slow
Early perturbations	59%	73%
Mid-reach perturbations	45%	62%

We were principally interested in measuring and comparing response latencies across experimental conditions. Doing so from the raw trajectories posed two problems. First, the perturbations were triggered by the finger reaching a certain position along the trajectory rather than a certain time; thus, the time of the perturbations varied from trial to trial. This adds considerable noise, as well as potential biases, to estimates of response latency drawn from comparing mean perturbed and unperturbed trajectories. Second, the raw trajectories have a large variance, reducing the sensitivity of any method based on simple point-by-point comparisons of averaged trajectories. In order to compute more reliable

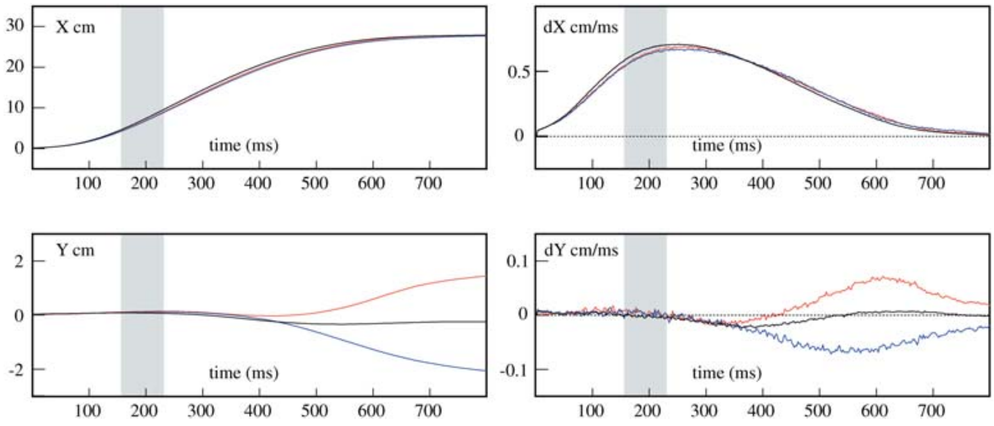

Fig. 5 Kinematic data plotted as functions of time for a sample subject and condition (slow movements, early perturbations). The *left* plots show the mean X and Y components of normalized finger position, and the *right* plots show the mean X and Y components of finger velocity. The three mean trajectories in each graph are for trials with positive perturbations (*blue*), negative perturbations (*red*), and no perturbations (*black*)

estimates of response latencies, we used a regression analysis that takes advantage of the smoothness of the trajectories. Kinematic data were fit by a linear model which expressed finger position as a weighted sum of the previous N positions plus a weight given to the perturbation used on a trial. The method amounts to filtering out temporal correlations in finger position and correlating the residual with the fingertip perturbation. This normalizes for overall variations between trials, and effectively estimates the proportion of residual error that can be accounted for by the perturbations. At times before the feedback has had an effect on the kinematics, the perturbation weight will necessarily be zero. We estimated response latencies to be the first time after the perturbation that the effect of the perturbations became significantly greater than zero.

Regression analysis

For the baseline, unperturbed condition, we used an autoregressive linear model to predict the position of the finger at time, t, from the position of the finger at previous times,

$$y(t) = \sum_{i=1}^{n} w(t,i)y(t-i) \tag{1}$$

The weighting function, $w(t,i)$, varies as a function of time, reflecting the fact that reach kinematics are not time-invariant. To the extent that finger positions are temporally correlated, the linear model should be effective at predicting $y(t)$. In our analysis, we have found that setting the order of the auto-regressive model to n=8 fit the data well, though the value of this parameter did not significantly affect our estimates of response latency.

In order to account for the effect of finger perturbations on subjects' movements, we extended the model to include a time varying term that represents the added influence of the perturbation on finger position,

$$y(t) = \sum_{i=1}^{n} w(t,i) \cdot y(t-i) + w_P(t) \cdot \Delta_P \tag{2}$$

where Δ_P is the perturbation in the Y direction, and $w_P(t)$ is the proportional influence of the perturbation on finger position at time t, after the influence of the preceding trajectory has been accounted for. If at time t a perturbation has had no effect, the auto-regressive component of the model should almost fully predict the current finger position, so the perturbation weight would be zero. Thus, perturbation weights will necessarily be zero for times up to the time of the perturbation plus the sensori-motor delay.

The main advantage of fitting auto-correlation weights along with perturbation weights is that the residual noise is much smaller than the variability in raw position, allowing smaller perturbation effects to be detected. There are other possible ways to factor out temporal correlations. In particular, computing smoothed derivatives also has a temporal de-correlating effect, so comparing velocity or acceleration functions to detect perturbation effects would share some advantages with our method. However, differentiating a discrete time series is equivalent to applying a linear filter, so these methods are actually suboptimal special cases of our method, in which weights are constant (e.g., [−1 1] for unsmoothed velocity) rather than fitted to the data at each time step. Differences between velocity or acceleration functions would therefore be a less sensitive measure than the perturbation weights from our regression model.

To account for the variance in the actual time of perturbations across trials, we made two further assumptions in the model: (1) The autoregression function is invariant over the small range of times that perturbations happen (±30 ms across different trials) (analysis of the unperturbed trials shows that this appears to be true) and

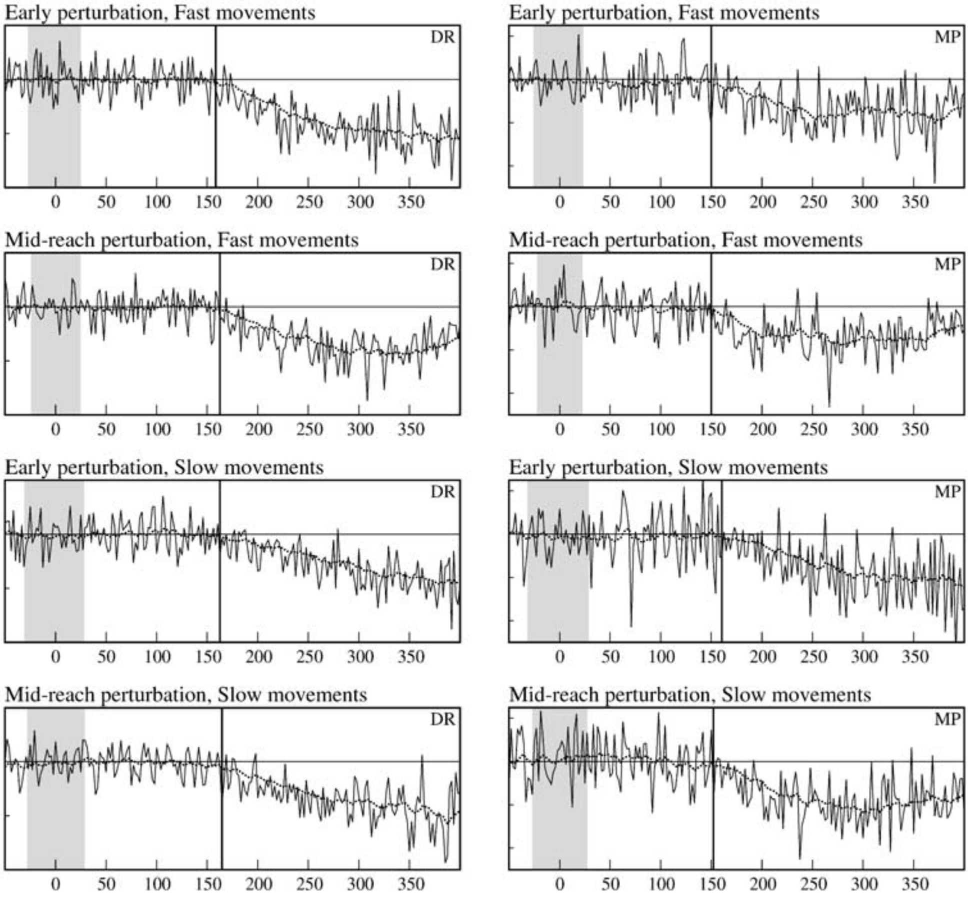

Fig. 6 Perturbation weight functions for the same subjects and conditions shown previously. The *superimposed dashed lines* show the result of smoothing the weight functions by an exponential filter. The *vertical lines* show the earliest significant response to perturbations, and the *shaded regions* show the time window of perturbation onset

(2) The perturbation weight, $w_P(t)$, when expressed as a function of the time after a perturbation, is invariant over the same time window. This allowed us to shift all of the trial trajectories to a common temporal reference frame centered at the time of perturbation prior to fitting the model. Thus, in the upcoming plots, $t=0$ is the time of the perturbation.

To estimate response latency, we first smoothed the perturbation weight functions with an exponential filter (time constant = 25 ms), which produces a more cumulative measure of the influence of a perturbation. These filtered weight functions were then compared against thresholds determined by re-sampling baseline, unperturbed trials (see Appendix). A response was inferred when the cumulative influence crossed the threshold and remained past threshold for at least the following 200 ms.

Figure 6 plots the perturbation weight functions for data from the same subjects and conditions plotted earlier, with vertical lines depicting response time. Although the X-Y plots for these data appeared heterogeneous, the response latencies are similar across conditions. Figure 7 shows the mean response latencies across subjects for the four conditions of the experiment. The overall mean latency was 163 ms, and the condition means were all close to the overall mean. An ANOVA found no effect of either time of perturbation ($F_{(1,33)}$=0.40, P=0.53 n.s.) or movement speed ($F_{(1,33)}$=2.44, P=0.13 n.s.), nor any interaction ($F_{(1,33)}$=2.44, P=0.13 n.s.).

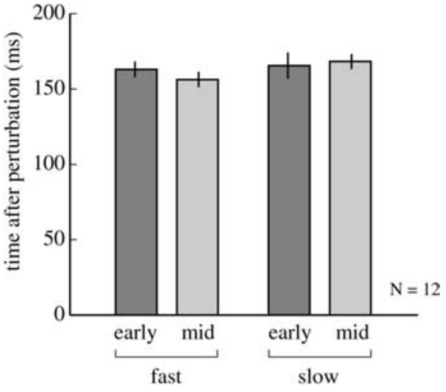

Fig. 7 Mean response latencies for each of the four combinations of perturbation time (early vs. mid-reach) and movement duration (fast vs. slow)

Learning effects

In order to successfully perform the task, it was necessary for subjects to compensate for perturbations on trials when they occurred. One concern is that this method may

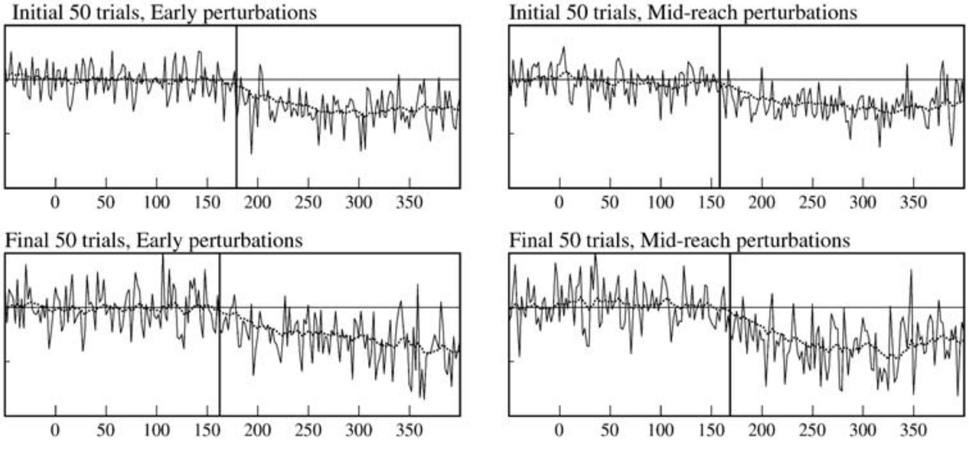

Fig. 8 Test for an effect of learning. Perturbation weight functions computed across subjects using either the first 50 trials (*top left and right*) or the last 50 trials (*bottom left and right*) of each subject's first experimental session. Weight functions were computed separately for early perturbations (*left*) and mid-reach perturbations (*right*)

have biased subjects toward using visual feedback more than they would have otherwise. Was there a learning component to subjects' responses to perturbations?

Two aspects of our method make it unlikely that learning significantly contributed to our results. First, the direction of perturbations was randomized, so subjects could not adapt to any fixed offset. Thus, the only advantage to be gained from training would be a general increased reliance on visual feedback. Second, perturbations were subliminal—subjects did not notice the perturbations being applied. However, a possibility remains that exposure to early perturbation trials, in which end-point error may have been greater than expected, trained subjects to pay more attention to visual feedback. To test for a learning effect, we compared subjects' responses to perturbations that occurred in the first 50 trials and the last 50 trials of each subject's initial session. These subsets of trials were too small to compute reliable perturbation weights for individual subjects, so we pooled across subjects for the regression analysis. The results are shown in Fig. 8. We found no difference in latency or magnitude of response. Since this analysis still pools across a set of trials, we cannot entirely rule out effects of learning. But it does imply that any adaptation that occurred was complete within the first 50 trials, during which subjects were exposed to only 16 perturbed trials, on average.

Eye-hand coordination

The performance of the eye tracker varied greatly across subjects, but complete records of gaze direction were available for at least 50% of trials for all subjects. On the vast majority of trials, subjects made saccades to fixate the target before initiating hand movements. On 91% of trials with gaze direction data, fixation was in the vicinity of the target before any hand motion was detected, and for 50% of the remaining trials, fixation was near the target before the hand had moved 25% of the distance to the target. We specifically looked at the timing of the initial saccade relative to the application of perturbations, and on all but a handful of perturbed trials (~3 per session), saccades were complete before perturbations began. Thus, we can be confident that on most trials the virtual finger was visible throughout the movement, and not obscured by saccadic suppression. There was also no indication that subjects were attempting to track their moving finger during the movement. Informal observation suggests that subjects often made an initial saccade to fixate near where an upcoming target would be expected to appear, and then a smaller secondary saccade after the target appeared, but this cannot be confirmed from our truncated pre-trial gaze data. The average delay between completing the orienting saccade and initiating hand movements cannot be accurately estimated from our data, since for a large proportion of trials, subjects made predictive eye movements to the target area before we started recording gaze direction data (approximately 100 ms before movement initiation).

Discussion

In the present study, we applied a classical perturbation technique to study the control strategy used by the visuomotor system; that is, we induced small perturbations in the input to the system and measured the resulting changes in the output. For the problem at hand, the relevant input was the visual signal from the hand and the output was the position of the hand in space, both expressed as a function of time. The conventional model, in which visual feedback contributes to control only in the late phase of reaching movements, predicts that response latencies to the early perturbations in our experiment would have been significantly longer than to the late perturbations. We found exactly the opposite: no such effect. The near equivalence of response latencies to early and late perturbations of the virtual finger strongly suggests that visual feedback from the moving hand is incorporated into continuous on-line control throughout reaching movements.

In some conditions, detectable corrections in response to perturbations were observed when the hand was less than 60% of the way to the target. Given the speed of the movements and accounting for even minimal estimates of feedback delay, these corrections must have been based on visual feedback from early in the movement, well before the subjects' hands reached peak velocity. The fact that the corrections occurred despite subjects' lack of conscious awareness of the perturbations, and that they appeared in the earliest trials in the experiment, implies that they resulted from automatic feedback control mechanisms that are in operation during normal reaching movements. These results are particularly striking in light of the relative impoverishment of the visual information about the hand (a small red sphere rather than the entire limb).

Besides providing direct evidence for a continuous visual control signal from the moving hand, our experiment provides uniquely strong evidence for continuous on-line control of any sort. Previous data is consistent with either continuous or intermittent control of hand transport. Visual deprivation studies do not address the distinction, and corrective responses to target perturbations are consistent with either continuous or intermittent control strategies. When target perturbations are visible, they provide a transient signal that one might expect to trigger a discrete controlled response. Responses to perturbations made during initial orienting saccades could also result from a system designed to make discrete corrections of an initial motor plan at the time of target acquisition immediately following the orienting saccade.

While the latencies of responses were constant across conditions, the magnitudes of corrective responses were not. Smooth corrections during subjects' initial movements were larger (i.e., more accurate) for early perturbations than for mid-reach perturbations and were larger for slow movements than fast movements. These effects are not surprising if feedback is used continuously during movements, since reducing movement speed and shortening onset time both have the effect of increasing the amount of time that feedback is available. The benefit for early corrections is consistent with results of an experiment by Komilis et al. (1993), which made a similar comparison for target perturbations at different points along a movement.

One limitation of the present experiment is that the visual environment is very impoverished. The temporal consistency between the visual feedback and the actual movement of the finger was high, which enhanced the sense of immersion, but the small sphere representing the finger provided a minimal form of feedback. Moreover, no background visual information was present against which the position and movement of the hand could be measured. These observations render the findings of early feedback contributions more striking. We would expect even stronger effects if movements were made in a more realistic virtual environment.

Displacing prisms provide a more natural means of inducing spatial shifts in the visual feedback from the hand, but suffer from the problem that the shift is fixed and cannot be easily varied from trial to trial. Jakobson and Goodale (1989) avoided this problem by analyzing movement kinematics from only the initial trials of prism exposure in an adaptation experiment. Even when the prism shifts were small enough to go undetected, subjects accurately corrected for the induced shift. Subjects were also more likely to make multiple changes in movement direction on initial post-adaptation trials. While their data did not provide timing information about subjects' corrective responses, the results are qualitatively consistent with continuous feedback control.

The response latencies measured here (~160 ms) are somewhat longer than those previously reported for both detected and undetected target perturbations (Prablanc and Pélisson 1990; Prablanc and Martin 1992), which are in the range of 100–120 ms. This might reflect differences in the signals available for control in the two cases. Targets appear in the fovea after movement onset, while the moving hand appears in the periphery. Moreover, target perturbations elicit fast, corrective saccades which themselves could provide a useful signal to the reach controller. A simpler explanation might lie in the differences in perturbation magnitude between this study (~1.7° of visual angle) and previous studies (on the order of 10–30° of visual angle). In support of this explanation, we have found very similar response latencies to those reported here in a study of corrective responses to small target perturbations, both undetectable (during orienting saccades) and detectable (just after orienting saccades) (McConnell et al. 2001). We should also note that the smoothness of the perturbations adds a degree of uncertainty in the response latency estimates, as they extended over 25–35 ms around the perturbation time, depending on the experimental condition. Finally, the method we used to detect the time at which perturbation weights became significant is inherently biased toward greater times. Appendix details an analysis in which the method described here was applied to weight functions that are noisy copies of a piece-wise linear function. For functions whose parameters are well fit to subjects' data, the method over-estimates response latencies by approximately 25 ms. Correcting for this factor, would bring the latencies measured here down to an average of 138 ms.

Given the long latency for visual feedback to affect the motion of subjects' hands (optimistically, near 138 ms) relative to the short duration of the movements (~450 ms for the fast condition), one must ask how the visuomotor system can effectively use feedback for on-line control. In the introduction, we pointed out that internal feedforward models of the motor system could be used to compensate for delays in sensory feedback (Miall and Wolpert 1996; Desmurget and Grafton 2000). A feedforward model allows the system to estimate the current state of the hand from delayed sensory feedback, which can then be used to generate on-line adjustments to the movement of the hand. That the brain has such forward models available is suggested by evidence from hand localization experiments (Wolpert et al. 1995) and from reaching experiments with deafferented patients (Bard et al. 1999). Our experiment does not directly address the question of whether or not the brain uses a forward model to correct for sensory delays; however, the proportionally long delays in our experiment and the relatively effective corrections exhibited by subjects suggest the strong possibility that the brain incorporates feedforward models in its on-line control strategy.

The results of the current experiment provide the first direct evidence for continuous, on-line visual control of the moving hand that extends throughout the course of reaching movements. We hope that these results will help to settle the long-running debate concerning the role of visual feedback in the control of reaching movements. The technique of perturbing a virtual hand during reaching movements provides a promising tool for further exploring the nature of the visual feedback that the brain uses to control reaching movements.

Acknowledgements We would like to thank Ittai Bushlin, Albin Yeung, and Kavita Kadiwar for their assistance collecting data. Parts of these results were reported at the first annual meeting of the Vision Sciences Society in 2001. This research was supported by NIH grant EY09383.

Appendix

The empirical goal of the paper was to measure the time at which subjects began to show corrective responses to perturbations in the visual position of the virtual fingertip. In our analysis, this shows up as a change in the weight given to the perturbation in a linear model that predicts the position of the subjects' fingertip at each point in time. The true value of this weight function is necessarily zero before subjects show a corrective response and goes negative after they begin to correct. The measurement problem is to find at what time the true weight function changes from zero based on noisy estimates of the weight function (see, for example, Fig. 6). This is a particularly difficult estimation problem when one does not have a model available for the form of the true weight function (a form of blind signal detection). In effect, we have assumed that the perturbation weight function is of the form

$$w(t) = \begin{cases} 0; \ t < \Delta \\ f(t-\Delta); \ t \geq \Delta \end{cases} \tag{3}$$

where Δ is the response time of the subject and *f()* is some unknown function. The problem of estimating subject's reaction time is to estimate Δ from a noisy measurement of *w(t)*.

In the paper, we estimated Δ by finding where the weight function *w(t)* crosses (and stays below) a negative threshold value (the function *f()* is constrained to be negative in our application, since it reflects a corrective response in the direction opposite to the perturbation). Because the noise levels in the estimated weight functions were high relative to the size of the effect we were attempting to detect, we smoothed the weight functions prior to detection. In order to be conservative in our estimate, we used a causal smoothing filter. The necessary effect of this is to create a positive bias in our estimates of Δ.

In order to explore the size of the bias, we modeled *f()* as a linear ramp with a slope that fits well to subjects' average data. Corrupting the model weight function by noise of the same magnitude as measured in subjects' data (using the first 40 estimates of the weight function, during which time we are confident that the true weight function is 0) and applying the estimation technique described in the paper does indeed give positive biases in estimates of Δ. We found that setting a threshold equal to 1 std. deviation of the variance in the smoothed weight function estimates (derived from a re-sampling procedure) gave a positive bias of approximately 25 ms on average across a large number of simulated data sets. When the threshold is set to 0, we found no measurable bias. Unfortunately, the decrease in bias is offset by a rapid increase in variance of the estimated values for Δ across the simulated dataset; therefore, we settled on 1 std. deviation as the threshold to use in the paper (values greater than 1 do not appreciably shrink the variance).

Not only do the simulations give a rough estimate of the bias in our estimates of Δ, they also provide a measure of the inherent uncertainty of the estimate, caused by noise in the estimated weight functions. We found that for noise equal to the average noise in the estimated weights, the method gave rise to a standard deviation of 13 ms in estimates of Δ across the simulated data set. This compares favorably with the 16 ms value we found for the variance in estimated response times across all subjects and conditions in the experiment. Thus, almost all of the variance (approximately 66%) of estimated response times across subjects is accounted for by the inherent uncertainty induced by the noisiness of individual subjects' data. The remaining 34% of the variance is primarily accounted for by individual differences between subjects, as reflected in the lack of significant condition effects found in the ANOVA.

References

Bard C, Turrell Y, Fleury M, Teasdale N, Lamarre Y, Martin O (1999) Deafferentation and pointing with visual double-step perturbations. Exp Brain Res 125(4):410–416

Berthier NE, Clifton RK, Gullapalli V, McCall DD, Robin D (1996) Visual information and object size in the control of reaching. J Mot Behav 28:187–197

Blakemore SJ, Wolpert DM, Frith, CD (1998) Central cancellation of self-produced tickle sensation. Nat Neurosci 1(7):635–640

Carlton LG (1981) Processing visual feedback information for movement control. J Exp Psychol Hum Percept Perform 7:1019–1030

Connolly JD, Goodale MA (1999) The role of visual feedback of hand position in the control of manual prehension. Exp Brain Res 125:281–286

Desmurget M, Grafton S (2000) Forward modeling allows feedback control for fast reaching movements. Trends Cog Sci 4(11):423–431

Desmurget M, Rossetti Y, Jordan M, Meckler C, Prablanc C (1997) Viewing the hand prior to movement improves accuracy of pointing performed toward the unseen contralateral hand. Exp Brain Res 115(1):180–186

Elliott D, Madalena J (1987) The influence of premovement visual information on manual aiming. Q J Exp Psychol A 39(3):541–559

Flanagan JR, Feldman AG, Ostry DA (1991) Equilibrium control vectors subserving rapid goal-directed arm movements. In: Requin J, Stelmach GE (eds) Tutorial in Motor Neuroscience, Kluwer Academic Publishers, The Netherlands, pp 357–367

Flash T, Henis E (1991) Arm trajectory modifications during reaching towards visual targets. J Cog Neurosci 3:220–230

Goodale MA, Pélisson 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(6064):748–750

Goodbody SJ, Wolpert DM (1999) The effect of visuomotor displacements on arm movement paths. Exp Brain Res 127:213–223

Hay L, Beaubaton D (1986) Visual correction of a rapid goal-directed response. Percept Mot Skills 62(1):51–57

Hoff B, Arbib MA (1991) A model of the effects of speed, accuracy, and perturbation on visually guided reaching. In: Caminiti R, Johnson PB, Burnod Y (eds) Control of Arm Movement in Space: Neurophysiological and Computational Approaches. Exp Brain Res Series 22:285–306

Jakobson LS, Goodale MA (1989) Trajectories of reaches to prismatically-displaced targets: evidence for "automatic" visuomotor recalibration. Exp Brain Res 78(3):575–587

Jakobson LS, Goodale MA (1991) Factors affecting higher-order movement planning: a kinematic analysis of human prehension. Exp Brain Res 86:199–208

Jeannerod M (1984) The timing of natural prehension movements. J Mot Behav 16:325–254

Jeannerod M (1988) The neural and behavioural organization of goal directed movements. Clarendon Press, Oxford

Jordan MI (1996) Computational aspects of motor control and motor learning. In: Heuer H, Keele SW (eds) Handbook of Perception and Action: Motor Skills, vol.2. Academic, San Diego, pp 71–118

Keele SW, Posner MI (1968) Processing of visual feedback in rapid movements. J Exp Psychol 77:155–158

Komilis E, Pélisson D, Prablanc C (1993) Error processing in pointing at randomly feedback-induced double-step stimuli. J Mot Behav 25(4):299–308

McConnell D, Grudic G, Knill DC, Kumar V (2001) Reach Corrections To Unnoticed Target Perturbations. Poster presented at the annual meeting of the Vision Sciences Society, Sarasota, FL

Miall RC, Wolpert DM (1996) Forward models for physiological motor control. Neural Netw 9:1265–1279

Paillard J (1982) The contribution of peripheral and central vision to visually guided reaching. In: Ingle DJ, Goodale MA, Mansfield RJW (eds) Visually Oriented Behavior. MIT Press, Cambridge, pp 367–385

Pélisson 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):303–311

Prablanc C, Martin O (1992) Automatic control during hand reaching at undetected two-dimensional target displacements. J Neurophysiol 67(2):455–469

Prablanc C, Pélisson D (1990) Gaze saccade orienting and hand pointing are locked to their goal by quick internal loops. In: Jeannerod M (ed) Attention and performance XIII: motor representation and control. Lawrence Erlbaum Associates, Hillsdale, pp 653–676

Prablanc C, Echallier JF, Jeannerod M, Komilis E (1979a) Optimal response of eye and hand motor systems in pointing at a visual target I: Spatio-temporal characteristics of eye and hand movements and their relationships when varying the amount of visual information. Biol Cybern 35:183–187

Prablanc C, Echallier JF, Komilis E, Jeannerod M(1979b) Optimal response of eye and hand motor systems in pointing at a visual target II: static and dynamic visual cues in the control of hand movement. Biol Cybern 35:113–124

Rizzo M, Darling W (1997) Reaching with cerebral tunnel vision. Neuropsychologia 35(1):53–63

Rossetti Y, Stelmach G, Desmurget M, Prablanc C, Jeannerod M (1994) The effect of viewing the static hand prior to movement onset on pointing kinematics and variability. Exp Brain Res 101:323–330

Sergio LE, Scott SH (1998) Hand and joint paths during reaching movements with and without vision. Exp Brain Res 122:157–164

Soechting JR, Lacquaniti F (1983) Modification of trajectory of a pointing movement in response to a change in target location. J Neurophysiol 49:548–564

Vindras P, Desmurget M, Prablanc C, Viviani P (1998) Pointing errors reflect biases in the perception of the initial hand position. J Neurophysiol 79(6):3290–3294

Witney A, Goodbody SJ, Wolpert DM (1999) Predictive motor learning of temporal delays. J Neurophysiol 82:2039–2048

Wolpert DM, Ghahramani Z, Jordan MI (1995) An internal model for sensorimotor integration. Science 269(5232):1880–1882

Wolpert DM, Goodbody SJ, Husain M (1998) Maintaining internal representations: The role of the superior parietal lobule. Nat Neurosci 1(6):529–533

Zelasnik HN, Hawkins B, Kisselburgh L (1983) Rapid visual feedback processing in single-aiming movements. J Mot Behav 15:217–236