Context Estimation for Sensorimotor Control

PHILIPP VETTER AND DANIEL M. WOLPERT
Sobell Department of Neurophysiology, Institute of Neurology, University College London,
London WC1N 3BG, United Kingdom

Vetter, Philipp and Daniel M. Wolpert. Context estimation for sensorimotor control. *J Neurophysiol* 84: 1026–1034, 2000. Human motor behavior is remarkably accurate and appropriate even though the properties of our own bodies as well as the objects we interact with vary over time. To adjust appropriately, the motor system has to estimate the context, that is the properties of objects in the world and the prevailing environmental conditions. Here we show that to determine the current context the CNS uses information from both prior knowledge of how the context might evolve over time and from the comparison of predicted and actual sensory feedback. We show that these two sources of information may be modeled within the CNS and combined to derive an accurate estimate of the context which adjusts motor command selection. This provides a novel probabilistic framework for sensorimotor control.

INTRODUCTION

Our ability to generate accurate and appropriate motor behavior relies on tailoring our motor commands to the prevailing movement context. This context embodies parameters of both our own motor system, such as the level of muscle fatigue, and the outside world, such as the weight of a bottle to be lifted. As the consequence of a given motor command depends on the current context, the CNS has to estimate this context so that the motor commands can be appropriately adjusted to achieve accurate control.

In general, the CNS does not have direct access to the true context but can gain information about it through sensory feedback. The estimation problem arises because our sensors are noisy, time-delayed, and often provide incomplete information about the context. Also the context is not static but tends to evolve in a structured but often uncertain way. For example the weight of a bottle tends to decrease as we pour from it, and muscles fatigue with use. Thus as the context evolves over time, sensory feedback can provide information about the current context. In the presence of an unknown and possible changing context, the CNS can continuously estimate the current context using two types of processes. First, the CNS could compare the actual sensory feedback with the feedback predicted for each possible context. This allows the CNS to estimate the likelihood of each context being present given a particular sensory feedback signal. Second, the CNS could model how the context is likely to evolve over time. Given these two types of processes, the CNS can then estimate the probability of being in a particular context (see METHODS for details).

Previous studies have shown that the motor system can adapt, at least partially, to multiple contexts, such as altered relationships between actual and perceived finger location (Kravitz and Yaffe 1972; Welch et al. 1993) and is able to select the appropriate behavior when the context is cued by gaze direction (Hay and Pick 1966; Kohler 1950; Shelhamer et al. 1991), body orientation (Baker et al. 1987), arm configuration (Gandolfo et al. 1996), or an auditory tone (Kravitz and Yaffe 1972). When the cue given is intermediate between two cues each associated with a different context, the behavior seen is consistent with mixing of the learned behaviors (Ghahramani and Wolpert 1997).

Whereas these previous studies have shown how prior cues determine movement selection in multiple task learning or how sensory feedback during movement is used within a single context (e.g., Elliott et al. 1995), we have sought to study how sensory information is used on-line to estimate the current context and thereby control movement. To achieve this, subjects received visual feedback under one of two contexts during pointing. The contexts were two different relationships between perceived and actual finger position. This relationship was generated experimentally using virtual reality.

First, we explored how visual feedback limited to brief instances during a pointing movement affected pointing behavior (rotation + translation experiments) after learning to point under two different contexts. The timing of feedback and which context the feedback was generated by was systematically varied. We assessed how such feedback affected pointing. In the first, the rotation experiment, the two contexts were chosen so that the difference between them increased into the movement, while in the second, the translation experiment, it remained fixed.

Our results indicated that recent feedback has a greater effect on pointing than earlier feedback. To study the mechanism by which more recent feedback may be made more salient than older feedback, we conducted a third, transition experiment. Here we tested two hypotheses that made qualitatively different predictions about pointing behavior when visual finger feedback is withheld for several movements. The results of all three experiments are then placed within a probabilistic framework of context estimation.

METHODS

Apparatus

Subjects made unconstrained movements to visual targets under stereoscopic visual feedback (Fig. 1*A*) (for full details of the system,

---

Address for reprint requests: P. Vetter, Sobell Department of Neurophysiology, Institute of Neurology, University College London, Queen Square, London WC1N 3BG, United Kingdom (E-mail: p.vetter@ucl.ac.uk).
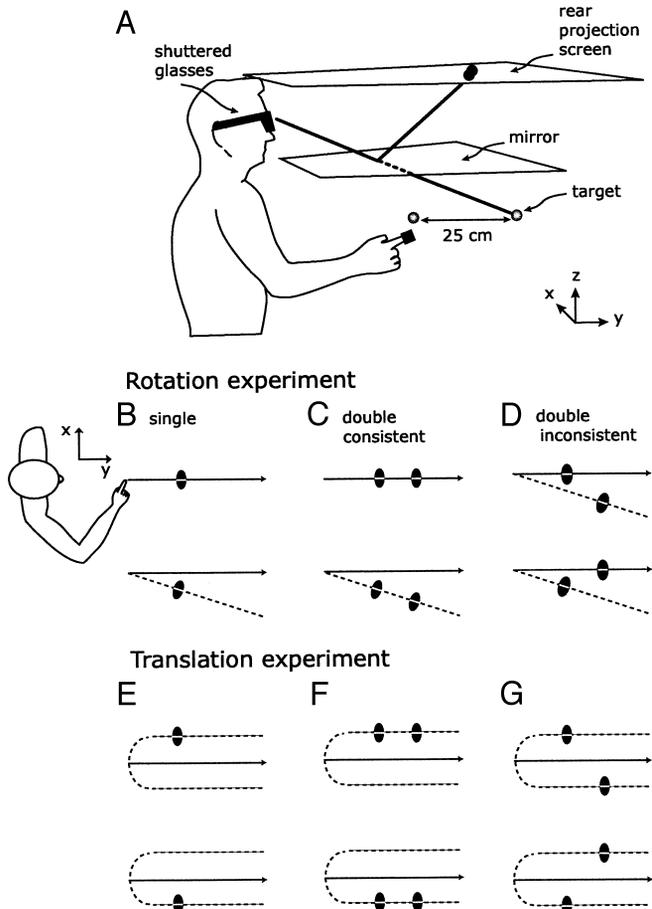
FIG. 1. *A*: apparatus for measuring unconstrained 3-dimensional arm movements under 3-dimensional virtual visual feedback. Looking down at the mirror, the subject sees virtual images of the finger and targets, which are projected onto the rear projection screen from the computer. Shuttered glasses alternately blanked the view from each eye in synchrony with the display, such that each eye was presented with the appropriate planar view—subjects therefore perceived a 3-dimensional scene. *B–D*: rotation experiment paradigm. →, a straight-line finger path; - - - , where the finger would be displayed under rotated feedback. ●, example locations where the finger feedback was briefly displayed. Single visual feedback condition: the finger is briefly displayed either at the veridical location (*top*) or rotated location (*bottom*). Double-consistent visual feedback condition: the finger is displayed twice either both veridically (*top*) or both rotated (*bottom*). Double-inconsistent visual feedback condition: the finger is displayed twice as an inconsistent pair, veridical:rotated (*top*) or rotated:veridical (*bottom*). Note that under right-rotated feedback, subjects have to point to the left to perceive their finger on target. *E–G*: translation experiment paradigm. *E–G* are analogous to *B–D* except that visual feedback is either left-translated or right-translated as opposed to veridical/rotated.

see Goodbody and Wolpert 1998). Subjects sat with their head in a chin rest and had infrared-emitting diodes (IREDs) mounted on their finger. Marker positions were recorded at 100 Hz using an Optotrak 3020 (Northern Digital, Waterloo, Ontario, Canada). The Optotrak was driven from a Silicon Graphics (SGi) Indigo 2 × Z workstation (Silicon Graphics, Mountain View, CA) where the position data were stored for later analysis. The markers allowed the position of the finger to be used on-line in the virtual visual feedback display on the SGi.

The targets and finger position were displayed as virtual three-dimensional images. This was achieved by projecting the SGi screen (1,280 × 500 pixels) with a cathode ray tube projector (Electrohome Marquee 8000 with P43 low-persistence phosphor green tube, Rancha Cucamonga, CA) onto a horizontal rear projection screen suspended

above the subject's head (Fig. 1*A*). A horizontal front-reflecting semi-silvered mirror was placed face up below the subject's chin (30 cm below the projection screen). The subject viewed the reflected image of the rear projection screen through field-sequential shuttered glasses (Crystal Eyes, Stereo-Graphic, San Rafael, CA) by looking down at the mirror. The SGi displayed left and right eye images of the scene to be viewed at 120 Hz. The shuttered glasses alternately blanked the view from each eye in synchrony with the display thereby allowing each eye to be presented with the appropriate planar view—subjects therefore perceived a three-dimensional scene.

Prior to each session the position of the IREDs relative to the projected image position was calibrated for each subject. By illuminating the semi-silvered mirror from below, the virtual image and the IRED could be lined up by eye. Each subject calibrated on 12 points on a three-dimensional grid covering the workspace. Image positions were linearly regressed onto IRED positions, and this fit was subsequently used on-line to position the target and finger images. Mean calibration errors in cross-validation sets were smaller than 1.0 cm.

During the experiments an opaque sheet was fixed beneath the semi-silvered mirror, thereby preventing any direct view of the arm. Finger feedback was provided as a 4-mm green wire cube in the virtual scene, and targets were presented as 4-mm-diam colored spheres. A static grid was displayed in the horizontal plane underneath the targets. By extinguishing the cube, which represented the finger position, movements in the absence of visual finger feedback could be examined. Computer-controlled discrepancies could also be introduced between the actual finger position and the displayed finger position.

Subjects

All subjects in the rotation (*n* = 9), translation (*n* = 8), and transition (*n* = 8) experiments were right-handed, gave their informed consent and were naive to the purpose of the experiment. Three subjects who had participated in the rotation experiment also participated in the translation experiment after an interval of at least 3 wk.

Rotation experiment

The aim of the experiment was to adapt subjects to two different contexts, that is visuomotor mappings, to examine how sensory feedback during a movement is used to estimate the current movement context. Trials were initiated by a movement to the start target (8 cm right of, 21 cm in front of and 28 cm below subjects' eyes) with visual feedback of finger position only given within 4 cm of the start. The actual trial was a pointing movement to a continuously displayed target 25 cm ahead (mean duration 1.15 s; Fig. 1*A*). Finger feedback was either withheld for the entire movement (no feedback), provided for the entire movement (full), or only provided for one (single) or two (double) brief instances of 50 ms. When provided, the visual feedback of the finger location was either veridical or rotated 7° to the right about the starting target in the horizontal plane (see Fig. 1*B*: note for clarity the angle is exaggerated). Note that the two contexts are indistinguishable at the start of the movement.

To assess subjects' pointing behavior before any learning, subjects initially performed 18 trials with full veridical feedback followed by 18 trials with a brief instance of either veridical (9) or rotated (9) feedback 10 cm into the movement, in a pseudo-random order that was different for each subject. To adapt subjects to the two visuomotor contexts, subjects then performed full-feedback trials under veridical (18 trials) or rotated (18 trials) feedback in pseudo-random order.

To examine how sensory feedback during the movement affects the context estimate, we then examined (postlearning) subjects' pointing when they received either no, one, or two brief instances of visual feedback during the movement. Single instances of feedback (Fig. 1*B*) were given 4, 6, 8, 10, 12, 14, or 16 cm into the movement, and the feedback presented was either veridical or rotated. In double feedback

trials an additional instance of feedback was given at 10 cm. The two instances of visual feedback were either consistent, that is veridical: veridical or rotated:rotated (Fig. 1*C*), or inconsistent, that is veridical: rotated or rotated:veridical (Fig. 1*D*). Each of the 39 trial types—1 no feedback +7 (positions) × 2 (contexts) single feedback + 6 (positions excluding 10 cm) × 2 (contexts for 1st instance) × 2 (contexts for 2nd instance) for double feedback—was presented nine times in a pseudo-random order, interspersed with full feedback trials every five trials to prevent decay of learning. Movements with no or only brief visual feedback ended when finger velocity dropped below 3 cm/s, while movements with full visual feedback terminated when the finger was within 0.8 cm of the target.

Translation experiment

The Translation experiment differed from the rotation experiment in one main respect—the two contexts were a 0.75-cm leftward and a 0.75-cm rightward translation, which were introduced gradually over the first 4 cm of the movement (Fig. 1, *E–G*). Trials were initiated by moving the finger within 4 mm of the start target. Subjects were familiarized with 32 full veridical feedback trials followed by eight trials with no visual feedback. To assess prelearning pointing, subjects performed 16 trials with either left-translated (8) or right-translated (8) feedback 10 cm into the movement. During learning, subjects received either full left-translated (36 trials) or full right-translated (36 trials) feedback on each trial in a pseudo-random order. In the post-learning trials, single instances of feedback were given at 5, 8, 10, 12, or 15 cm, while in double feedback trials, as in the rotation experiment, an additional instance of feedback was given at 10 cm. Each of the 27 conditions was presented eight times, interspersed with full feedback trials every three trials.

Transition experiment

The aim of the transition experiment was to adapt subjects to two different visuomotor contexts to examine pointing behavior when feedback is withheld for several movements. Subjects started trials by moving their finger behind a notional fronto-parallel plane 18 cm in front of the eyes without visual feedback. A trial consisted of a variable number of pointing movements to successively neighboring targets in the corners and center of an 11 cm vertex horizontal hexagon (center, 6 cm right of, 28 cm in front of and 35 cm below subjects' eyes). Target order was pseudo-random, although the first movement was to one of the three most distal targets.

To factor out subjects' pointing errors in the absence of visual feedback, the first trial consisted of 5 movements with visual feedback followed by 30 movements without. The average height over the last 15 movements was taken as the subject's estimate of the correct target height, and all subsequent visual feedback was corrected by the discrepancy between this height and the veridical height. Before learning, we examined the effects on pointing of exposing subjects to a novel context, under which subjects had to point 4 cm higher than usual to be on target. Subjects performed seven trials consisting of three movements with feedback followed by four without. During learning, subjects made movements alternately with and without feedback. Each time feedback was turned on again, they had to point either 4 or 8 cm too high in a pseudo-random order with equal probability. Therefore when feedback re-appeared, there was a 0.5 chance that the context was the same and 0.5 that it had changed. The 12 learning trials consisted of eight blocks of three movements with feedback followed by four movements without.

For analysis, trials were split into movements on the basis of velocity profiles smoothed with a fifth-order, zero phase lag, low-pass filter with a 5 Hz cutoff. Movements started when the velocity exceeded 0.3 cm/s.
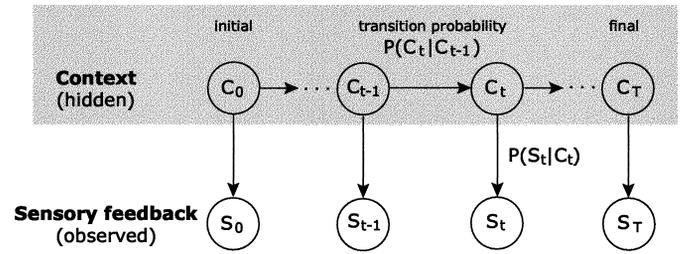

FIG. 2. Context evolution with sensory feedback as a hidden Markov process. *Top*: the evolution of the true context *C*, for example the weight of a bottle as we pour from it, which is hidden from the CNS but must be estimated for accurate control. *Bottom*: the sensory feedback *S* observed by the CNS enables it to infer the context. At each time *t*, the current context is determined by the previous context with a fixed transition probability given by $P(C_t|C_{t-1})$. As sensory transduction is a noisy process, the sensory feedback has a probability distribution that depends on the current context $P(S_t|C_t)$.

Context estimation model

The rotation and translation experiments were quantitatively fitted with the context estimation model (for a qualitative description of the model, see DISCUSSION) shown in Fig. 2. Here, the context at time *t*, C_t, is an *n* dimensional vector of indicator variables, which specifies which of *n* possible discrete contexts we are acting in, for example rotated or visual feedback. At each time step, a single element is one, reflecting the current context, while all others are zero. The context can change over time, and this can be represented by an $n \times n$ transition matrix whose elements determine the probability of transitioning from one context to another, $P(C_t|C_{t-1})$.

We assume that the CNS starts with estimates of the context at *time 0* given by the prior probability of each context being present $\hat{C}_0$ (we use ^ to denote an estimate). If at time $t - 1$, our context estimate is $\hat{C}_{t-1}$, then at time *t* our best estimate before considering sensory feedback is determined by the transition probability, $\tilde{C}_t = \hat{P}(C_t|C_{t-1})\hat{C}_{t-1}$. If no sensory feedback is available, then modeling how the context evolves is the best the CNS can do, that is $\hat{C}_t = \tilde{C}_t$. However, sensory feedback provides additional information. As different contexts lead to differing predictions of what the sensory feedback should be, the context estimate can be improved by determining how likely the observed sensory feedback is for each possible context. Using Bayes' rule $\hat{C}_t = [\hat{P}(S_t|C_t) \odot \tilde{C}_t]/[\hat{P}(S_t|C_t)^T\tilde{C}_t]$, where $\odot$ is element-wise multiplication and ^{T} denotes transpose.

We took transverse end points on individual trials, *x*, as an index of the CNS's final context estimate (Ghahramani and Wolpert 1997). These were considered a mixture of the learned end points under the two contexts, for example $e = (x_{\text{rotated}}, x_{\text{veridical}})$ weighted by the CNS's recursively derived estimate of the context, $x = C_t^T e$. To reduce the number of free parameters, the learned end points *e* were parametrized as deviations d_{rot} and d_{tr} from the observed no-visual feedback pointing positions in the rotation and translation experiments, reflecting the degree of adaptation to the contexts, which in general is not complete (Welch 1986). We assumed that visual sensory feedback, when present, represented the spatial location of the finger corrupted by Gaussian noise with fixed variance σ^2. Therefore the likelihood of a particular finger position given context 1, is $(2\pi\sigma^2)^{-1/2} \exp\{[-(S_t - \hat{S}_t)^2]/2\sigma^2\}$, where $\hat{S}$ is the predicted finger position given context 1. The transition probability matrix $\hat{P}(C_t|C_{t-1})$ was parametrized by a single parameter *q* of the context changing at each time step (and thus a probability of $1 - q$ of staying the same).

The four free parameters of the model, the transition probability *q*, the variance of the noise in visual feedback σ^2, and the maximum end-point deviations from the no-feedback condition in each experiment d_{rot} and d_{tr}, were fit to the 64 data points from both experiments simultaneously. Each trial was simulated by starting with $\hat{C}_0 = (0.5, 0.5)$ and using the parameters to simulate context estimation throughout a trial incorporating any instances of feedback, if present. The
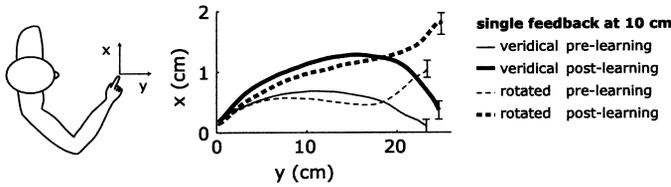
FIG. 3. Rotation experiment: finger paths before (- - -) and after learning (—), averaged over all subjects' paths. End points are shown with 95% confidence limits.

final context estimate at the end of the movement was used to predict the position the subject would point to and the four parameters were adjusted using the Simplex algorithm to minimize the squared error between the predicted and actual pointing location.

RESULTS

Rotation experiment

Before learning, when only a single instance of veridical feedback was given 10 cm into the movement, the finger path was roughly straight (Fig. 3, thin solid line). When instead a single instance of right-rotated feedback was given 10 cm into

the movement, this resulted in subjects compensating by pointing more to the left of the target (thin dotted line). As visual feedback was only given 10 cm into the movement, the paths for both the veridical and rotated feedback were, as expected, not significantly different before 10 cm. After learning, a single brief instance of finger position feedback affected ongoing pointing movement differently than before learning (Fig. 3, compare dotted and solid lines). Specifically, there was an increased difference in end points between trials with single veridical (thick solid line) or rotated feedback (thick dotted line) at 10 cm. Also, movements started with an initial bias slightly to the left of the true target.

We examined how the distance at which feedback was given affected pointing and found systematic changes in the movement paths (Fig. 4). When no visual feedback was given during the movement (black line Fig. 4), subjects pointed to a place intermediate between the end points under either of the two contexts. When a single instance of visual feedback was given, this biased pointing in a manner that depended both on the type of feedback and when it was given. A single instance of veridical feedback given during the movement biased pointing away from the no-feedback condition toward the true veridical
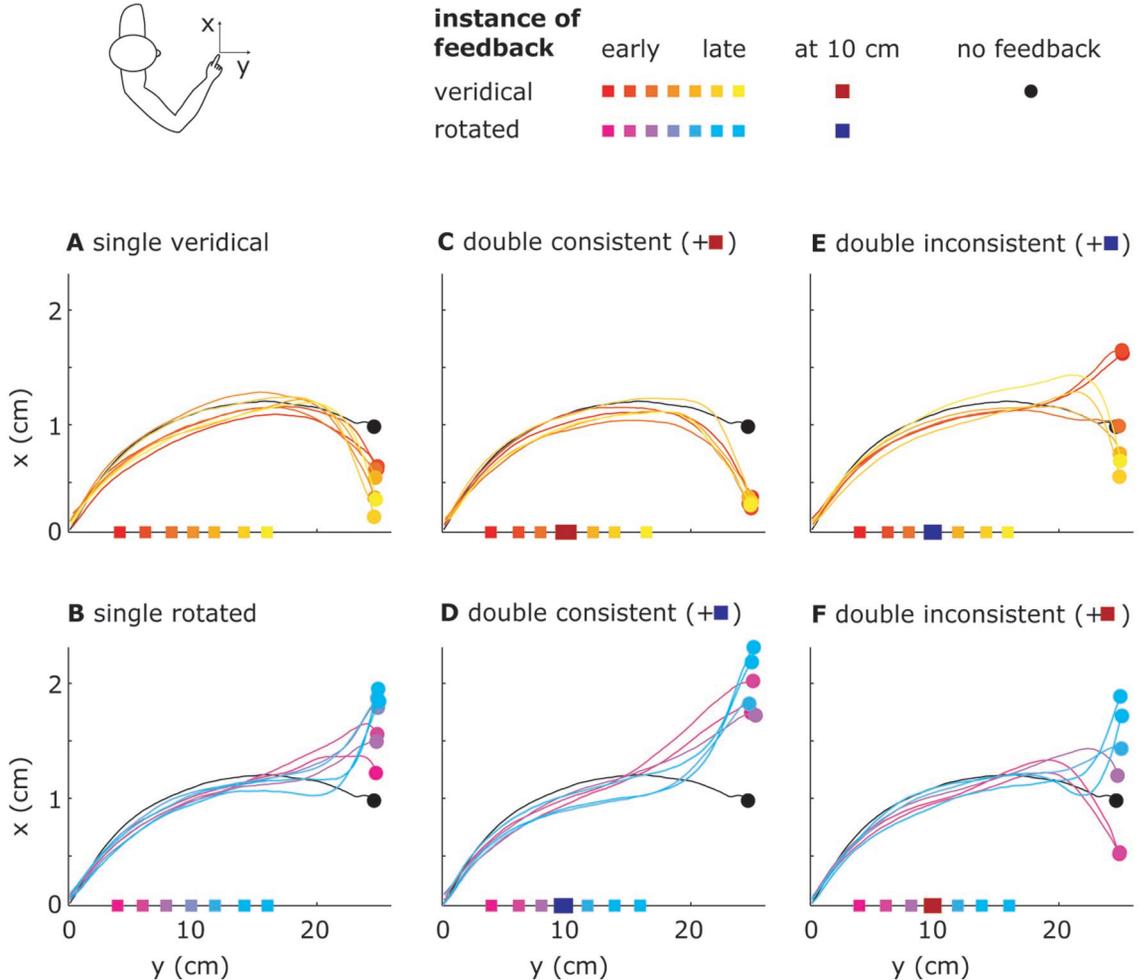

FIG. 4. Finger paths under single (*A* and *B*), double-consistent (*C* and *D*), and double-inconsistent (*E* and *F*) visual feedback after learning, averaged over subjects for each condition. *A* and *B*: for each path, a single instance of feedback was provided at the distance indicated by the same-colored square on the abscissa. Finger end points are marked with circles for clarity. In *C* and *D*, an additional instance of consistent feedback was given after the finger had moved 10 cm (marked by a blue or red square), while in *E* and *F*, an additional instance of inconsistent feedback was provided at 10 cm.
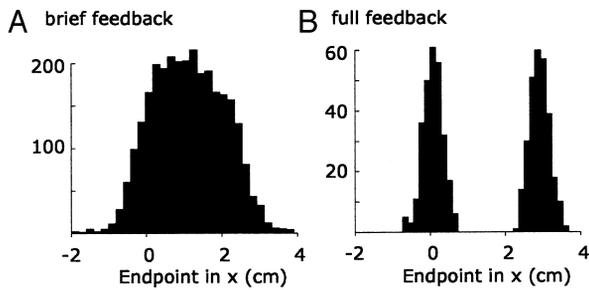
FIG. 5. End-point histograms of all postlearning trials from all subjects with brief instances of feedback (*A*) and full visual feedback (*B*).

end point (Fig. 4*A*). The amount of bias was smaller for feedback given early in the movement (red lines) than for feedback given later (yellow lines). Similarly, when a single instance of right-rotated feedback was given (Fig. 4*B*), subjects compensated by pointing further to the left than in the absence of feedback. Again this effect was greater for feedback given later in the movement (blue lines) than for earlier feedback (purple lines).

When in addition to this single instance of visual feedback a second consistent instance of feedback was given at 10 cm, the bias seen was enhanced for both the veridical:veridical (Fig. 4*C*) and rotated:rotated condition (Fig. 4*D*). However, when two inconsistent instances of feedback were given, a different pattern was observed. Figure 4*E* shows the pattern of trajectories seen when rotated visual feedback is given at 10 cm and a single instance of veridical feedback is given at a different location during the movement. When the single instance of veridical feedback is given before 10 cm and the rotated feedback is experienced second (orange lines), pointing is biased toward the left, appropriate for rotated feedback. When this single instance of veridical feedback is given after 10 cm and the rotated feedback is therefore experienced first (yellow lines), the pointing is biased toward the veridical location of the target. Therefore when inconsistent instances of visual feedback are given, the context of the second instance of visual feedback dominated pointing behavior. An analogous pattern was observed when veridical feedback was given at 10 cm and a single instance of rotated feedback at other locations (Fig. 4*F*).

When limited information as to the prevailing context is available (only brief feedback), two different strategies could be used. First, subjects could estimate how likely each context is and then mix the behaviors appropriate for each context weighted by the probability of being in that context. This strategy mathematically minimizes the expected mean squared pointing error (Abidi and Gonzalez 1992). Alternatively, subjects could simply estimate which context was more likely and then switch to the behavior appropriate for that context. While both strategies lead to the same mean end points, individual end points should cover the spectrum for mixing, while being bimodally distributed for switching. The histogram of transverse finger end points in brief feedback trials (Fig. 5*A*) is clearly not bimodal, even though the full feedback trials during learning show that training was on a bimodal data set (Fig. 5*B*). Subjects show a range of pointing locations, suggesting that they are not simply choosing one of the contexts.

While the full trajectories contain information about the context, they can only be interpreted if a model of the control-

ler and time delays are known. However, as subjects could choose to end the movement anywhere, the finger end point represents an estimate of the context subjects estimated they were moving in at the end of the movement (Ghahramani and Wolpert 1997). Figure 6 shows how the transverse end point varied as a function of the type of feedback and where it was given. The end point in the no feedback condition is shown as ▷. When a single instance of visual feedback was given (●/○), the end point deviated from the no-feedback condition. The deviation was in the direction appropriate for the context presented, that is to the right for veridical feedback (●) and in the opposite direction for rotated feedback (○). The deviation from the no feedback condition increased the further into the movement feedback was given. When an additional instance of consistent feedback was given 10 cm into the movement (□), this led to an equal or greater deviation than a single instance alone.

When two inconsistent instances of visual feedback were given, the context of the second instance had a greater influence on the pointing end points (Fig. 6 double inconsistent). However, the first instance of visual feedback still influenced the pointing, especially when the two instances were close together (compare the 3rd point in Fig. 6 double inconsistent with the other points). This ordering effect could arise from two possible processes. First, more recent visual feedback may be regarded as more "up to date" and thus more salient when estimating the current context. Second, since the difference between rotated and veridical feedback increases linearly with distance, distinguishing the contexts from the second instance of feedback is easier and thus more informative. To examine these two possibilities, we performed an additional experiment in which the difference between the two contexts was constant over the movement—the contexts were a fixed rightward and a fixed leftward translation.

Translation experiment

Transverse end points for the translation experiment are shown in Fig. 7 in the same format as for the rotation experiment. When no visual feedback was given during the move-
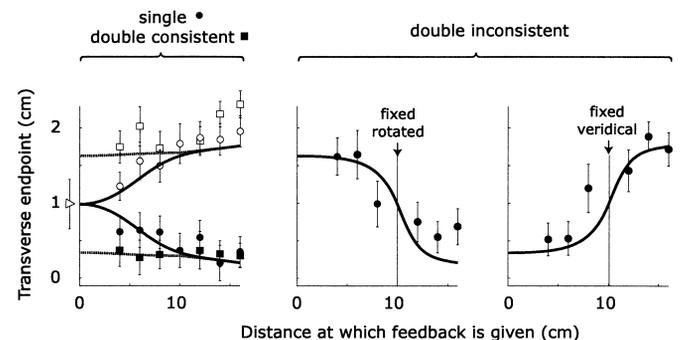

FIG. 6. Rotation experiment: transverse end points ±3 SE as a function of where and what kind of visual feedback was given, averaged over subjects for each condition. *A*: the end point in the absence of visual feedback is shown as ▷, while trials with a single instance of feedback are shown as ○ and ●. ●, trials in which feedback was veridical; ○, rotated feedback trials. The effect of an additional consistent instance of feedback at 10 cm is indicated by □. *B*: pointing end points when feedback is veridical, but at 10 cm an instance of inconsistent rotated feedback is given. Conversely, in *C* feedback at the distance indicated was rotated, while an instance of veridical feedback was given at 10 cm. —, the fit of the context estimation model.
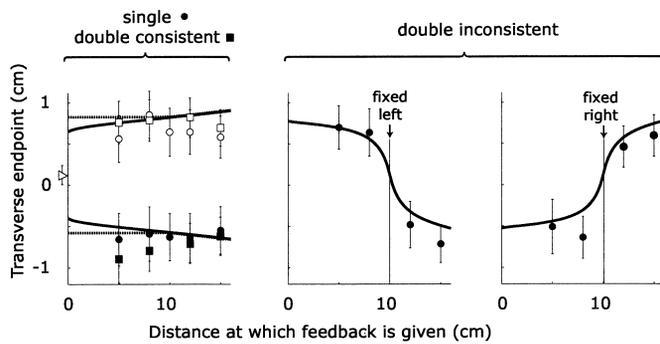
FIG. 7.   Translation experiment: transverse end points ±3 SE as a function of where and what kind of visual feedback was given, averaged over subjects for each condition. Trials with left-translated feedback are drawn as ○, trials with right-translated feedback as ●. Otherwise, the panels are analogous to the Fig. 6.

ment (▷), pointing was approximately straight. Single instances of feedback produced constant deviations from the no-feedback condition, regardless of when visual feedback was given (Fig. 7). To examine the time course of the change in pointing behavior induced by a single instance of feedback, we calculated the feedback-triggered average of the transverse finger position for both left-translated (Fig. 8 dark gray) and right-translated (light gray) feedback. This shows the average position aligned to the time feedback was given. The paths under left- and right-translated feedback diverged significantly 300 ms after a single brief instance of visual feedback ($P < 0.01$).

The constant pointing deviations contrast with the rotation experiment where later single instances had a greater effect on pointing. This suggests that the difference between the two contexts, which is greater for the translations than the rotations early in the movement, affects the way the context is estimated. If the difference between contexts was the only determinant of pointing behavior, then one would predict that in the inconsistent trials pointing end points should be identical to the no-feedback condition since averaging a leftward and a rightward translation would result in pointing straight ahead. However, the second instance of feedback again had a greater influence on the pointing end point (Fig. 6 double inconsistent), indicating that there is also a temporal component influencing context
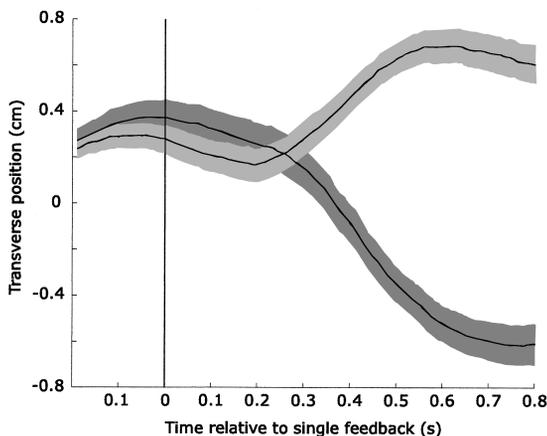

FIG. 8.   Translation experiment: transverse pointing positions relative to the time of feedback on single feedback trials. Pooled data from all subjects are shown with 95% confidence intervals. Right-translated feedback is shown in light gray and left-translated in dark gray.
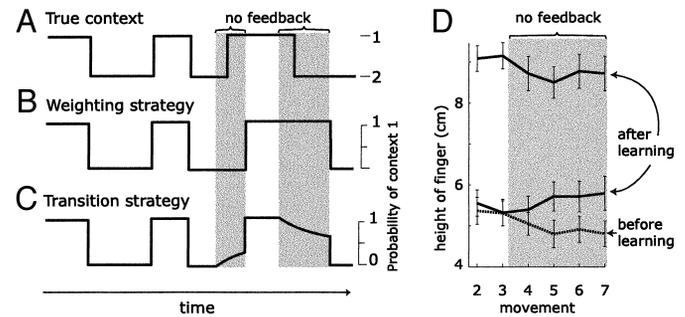

FIG. 9.   A–C: schematic model prediction of the relationship between the true evolving context and 2 strategies for estimating it. A: the true context switches between 2 states. In the presence of visual feedback, this evolution is closely tracked by 2 different context estimation strategies (B and C). Time delays are ignored for simplicity. However, when visual feedback is not available (shading) a weighting strategy (B) will leave the context estimate unchanged, while a transition strategy (C), which takes into account the likely evolution of the context, will gradually make the context estimate converge to a value between 0 and 1. D: transition experiment: finger height, 1st averaged over trials and then subjects (with inter-subject ± SE). Movements 2 and 3 are made with visual feedback and movements 4–7 are made without. - - -, pointing prior to learning with visual feedback 4 cm too low. —, pointing after training on contexts of 4 cm and 8 cm too low.

estimation. This suggests that both temporal and spatial attributes of the feedback influence context estimation.

There are two contrasting possibilities that can account for this temporal effect. First, in a *weighting* strategy context estimates are only updated when sensory information is available. So when a new piece of visual information is received, it is deemed more important than older information and therefore contributes more to the estimate. However, such a weighting strategy ignores information about how the context evolves over time, which could be used to update the context estimate even in the absence of visual feedback. Such a *transition* strategy uses knowledge of the way the context evolves to update the estimate at all times. For example, consider a situation in which two contexts are presented with equal probability and that transitions occur randomly between them (Fig. 9A). In the presence of continuous visual feedback, the context estimate is likely to track the true context. If visual feedback is removed, the weighting strategy maintains the current context estimate until new information is provided (Fig. 9B). However, in the transition strategy, the context estimate can change in the absence of visual feedback (Fig. 9C). If at some point in time we are certain that we are in context 1 and then visual feedback is turned off, then at the next point in time, the context may have transitioned to context 2, and therefore our estimate of context 1 should decrease. This process continues so that whatever the initial probability of each context, the probability of either contexts converges toward the overall probability of each context being present, that is, 0.5 in this case. This example is analogous to the rotation and translation experiments in which there are only two contexts, and the presentation of each is probabilistic and equally likely.

Transition experiment

To distinguish between weighting and transition strategies of context estimation over time, we conducted an experiment in which we could monitor explicitly the development of the context estimate in the absence of visual feedback. Before learning, when subjects had to point 4 cm too high, the finger

drifted downwards toward the veridical height on removal of visual finger feedback. Subjects were then trained on two contexts (unnoticed) in which subjects had to point either 4 or 8 cm too high in a pseudo-random order with equal probability, interspersed by movements without visual finger feedback. Having learned these two new contexts, the prediction is that in the absence of visual feedback the context estimate should converge to between 4 and 8 cm. Specifically, before learning the finger should drop from 4 cm, while after learning it should rise when feedback is eliminated. As predicted, the finger rose significantly from 4 cm after learning compared with before learning (1-tailed paired t-test, $P < 0.002$, *bottom solid* and *dotted lines,* Fig. 9*D*). This result lends strong support to the transition strategy for context estimation in which the estimates are updated even in the absence of visual feedback.

We fitted the data in the rotation and translation experiments assuming that that the CNS uses such a transition strategy to update its context estimate in the absence of sensory feedback and that sensory feedback, when available, is noisy (see METH-ODS and DISCUSSION for details). The former ensures that recent feedback is more salient than older feedback, while the latter ensures that contexts are easier to distinguish the larger the discrepancy between them. The fit of this model captures the features of the single, double-consistent, and inconsistent feedback in both the rotation and translation experiments (Figs. 6 and 7, solid lines, $r^2 = 0.95$ $P < 0.001$, parameters are $q = 0.45$ s^{-1}, $\sigma^2 = 0.22$ cm^2, $d_{rot} = 1.3$ cm and $d_{tr} = 1.3$ cm; see METHODS for details). Fitting each subjects' data individually gave a mean $r^2 = 0.86$ (range 0.74–0.95) Comparison of the full model to reduced models which either had no intra-movement context transitions ($q = 0$) or no noise in sensory feedback ($\sigma^2 = 0$) showed that these fits were significantly worse than the full model, even when taking into account the reduced number of parameters ($P < 0.001$).

DISCUSSION

We have investigated how subjects use sensory information to determine the context of their movement, that is, the properties of objects in the world and the prevailing environmental conditions. Subjects learned to point to a single target under two different contexts, for example rotated and veridical visual feedback, which were chosen to be indistinguishable at the start of the movement. After learning, the majority of the movement was made without visual feedback. When no feedback was given, subjects pointed to a location intermediate to those appropriate for each context. A single instance of feedback biased the end point toward the location appropriate for the context presented. When the difference between the finger's visual location under the two contexts increased with distance into the movement (rotated vs. veridical feedback), later feedback caused a greater deviation from the no feedback condition but not when it remained constant (leftward vs. rightward translation). Two instances of consistent visual feedback had a greater effect than one alone. However, when the two instances of inconsistent visual feedback were inconsistent, the second had a far greater influence. To examine whether this ordering effect was due to weighting of information by its recency or due to the CNS modeling how the context evolves over time, we explored the way in which the context estimate evolved in the absence of visual feedback. Subjects

pointed in a situation where the context varied with equal likelihood back and forth between two possibilities. When feedback had been turned off, subjects gradually adopted a behavior intermediate to that for either context. This behavior is consistent with the CNS modeling how the context is likely to evolve, but cannot be explained by just weighting visual feedback according to its recency (see also Fig. 9, *A–C*).

All three experiments can be parsimoniously captured in a general probabilistic framework of how the context evolves over time. Figure 2 shows a model of how the context evolves (time increasing to the right), with the true context at time t given by C_t. The context can change over time, which is represented by a transition matrix $P(C_t|C_{t-1})$ whose elements determine the probability of transitioning from one context to another. This transition matrix can capture both deterministic and probabilistic transitions reflecting uncertainty in the way the world evolves. At each time step, the context will influence the sensory feedback. As our sensors are noisy, sensory feedback is generated from a distribution that is related probabilistically to the current context.

Given this structure of the world, the task of the CNS is to estimate the current context. Two special cases of estimation in the framework of Fig. 2 can be considered. If the CNS assumes the context to remain constant over time, estimation reduces to a discrimination task often studied in sensory psychophysics in which some underlying parameter to be estimated is corrupted by Gaussian noise (Boff et al. 1986). However, such a discrimination model would be unable to capture the time-dependency of the context estimation in our experiment when feedback is inconsistent (Fig. 6, *E* and *F*). Alternatively, if sensory feedback is effectively noise free, then the CNS should only use the most recent instance of sensory feedback to estimate the context and would predict a switching type of behavior between contexts—a pattern inconsistent with the smooth end-point variations in our data (Figs. 6 and 7).

By simulating context estimation that assumes both a time-varying context and noisy feedback, we were able to fit our data set with few free parameters (Fig. 6). This implies that the CNS internally models both how the context is likely to evolve over time and the expected sensory feedback given each possible context. During a movement, in the absence of visual information, the context estimate is updated solely using knowledge about how the context evolves over time. Any instances of visual feedback are compared with the sensory feedback predicted for each of the possible contexts. Contexts that lead to small prediction errors are judged more likely. This context estimation model captures the key features of all three experiments both qualitatively and quantitatively (Figs. 6 and 7). First, the model suggests that the probability of each context is estimated and that therefore the optimal behavior is to mix the behaviors appropriate for each of the visuomotor contexts in proportion to their probability. Second, the model assumes sensory feedback to be noisy, thereby making contexts that have large sensory differences easier to distinguish than those with similar sensory feedback, in agreement with our finding that the size of the difference between contexts influences the context estimate. Third, the context estimate is continuously updated, even in the absence of feedback. This is achieved by modeling the likely transitions between contexts and effectively makes newer feedback more salient than old, as observed. Finally, the model predicts that in the absence of visual
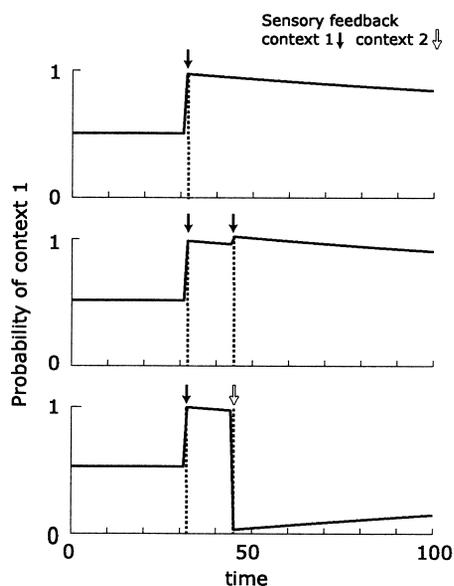
FIG. 10. Simulation of context estimation during a movement. Simulations are taken from the rotation experiment when feedback ($\downarrow$) was given after having moved 10 and 16 cm. Parameters as in Fig. 6. At movement onset, the prior probability of being in context 1 is estimated as 0.5. During the movement, sensory feedback becomes available ($\downarrow$) and the context estimate is updated accordingly. Over time, the estimate decays (*top*), unless it is modified by a 2nd instance of sensory feedback that is consistent (*middle*) or inconsistent (*bottom*) with the 1st instance.

feedback, the context estimate will gradually change toward the prior probability of the contexts. This is because after extended periods without visual feedback, knowing what the context was when you last had feedback provides little information. This prediction was tested and confirmed in the transition experiment.

To visualize the estimation process, the evolution of the context estimate during three simulated movements is shown in Fig. 10. Initially, the probability of context one (rotated feedback) is 0.5, reflecting the experimentally imposed probability of each context. On presentation of a single instance of rotated feedback (Fig. 10, *top*), the probability of this context increases and subsequently decays slowly. This decay is a result of the continual updating of the context estimate in the absence of feedback and ensures that more recent feedback is deemed more reliable than older feedback. A second consistent instance of rotated feedback (Fig. 10, *middle*) further increases the probability estimate, which is reflected in the greater bias for two consistent instances of feedback compared with a single instance in Fig. 6, *A* and *D* ($\square$ vs. $\bigcirc$). Conversely, a second instance of inconsistent feedback decreases the probability estimate (Fig. 10, *bottom*).

The estimation model proposed requires the CNS to construct and use two types of internal model—brain processes that simulate the behavior of both the body and the outside world (Ito 1984; Kawato et al. 1987; Merfeld et al. 1999; Miall and Wolpert 1995; Wolpert 1997; Wolpert et al. 1995). First it must model how the context evolves and second the likely sensory feedback given each possible context. This is consistent with recent neurophysiological studies in primates showing that the CNS represents the expected visual feedback for a particular context (Eskandar and Assad 1999) as well as the likelihood of the observed visual feedback given that context

(Kim and Shadlen 1999). Previous studies have shown that internal models play a role in maintaining accurate control in the presence of sensory feedback delays (Ito 1984; Miall et al. 1993), generating anticipatory responses (Forssberg et al. 1992; Gordon et al. 1993; Jenmalm and Johansson 1997), and in distinguishing our own actions from externally produced stimuli (Blakemore et al. 1998, 1999; Jeannerod 1997). The necessity for multiple internal models in context estimation comes from the need to predict the sensory consequences of actions for each possible context (Wolpert and Kawato 1998; Wolpert et al. 1998).

Various investigators have looked at how intermittent visual feedback affects accuracy during pointing movements (Spijkers and Spellerberg 1995) and ball-catching (Elliott et al. 1994). In general, movements with intermittent feedback were similar to movements with full visual feedback, with accuracy improving the more and the later visual feedback was given. Similarly a number of studies have addressed how rapidly visual feedback is available to correct ongoing movements, for instance, if the target position changes at movement onset (Castiello et al. 1991; Paillard 1996). While both types of study provide information about how visual feedback is used to refine movement, they were all performed under a single context and therefore do not address how the visual information is used to estimate the context.

Several studies have examined adaptation to multiple contexts but have focused on whether multiple contexts can be learned or cued prior to movement. Welch et al. (1993) found that alternating prism exposure causes dual adaptation. In terms of our framework, subjects learned two different visuomotor contexts and were able to switch between them rapidly when given full visual feedback. While Welch kept the context constant within trials, we explored, whether switching could also occur within trials. Ghahramani and Wolpert (1997) had subjects point to a single target under two different visuomotor contexts. In their experiment, each visuomotor context was linked to a particular starting position. When pointing from starting locations intermediate between the two learned, subjects' output was a mixture of that observed from either starting point. Within the model, the starting position determines the prior expectation of each mapping, and because no feedback is given during the movement, this estimate does not change.

We have shown that human pointing behavior is consistent with the CNS estimating the movement context using two types of internal model. The first captures the way in which contexts change over time and the second the likely sensory feedback given each context. Using either model in isolation is insufficient to reproduce the data. This provides a novel framework for sensorimotor control in which internal models must learn the probabilistic structure of both the world and their own sensors and combine these two sources of information to estimate the current context, thereby achieving accurate control.

Wellcome four-year PhD programme in Neuroscience at University College London.

REFERENCES

ABIDI MA AND GONZALEZ RC. *Data Fusion in Robotics and Machine Intelligence.* San Diego, CA: Academic, 1992.

BAKER JF, PERLMUTTER SI, PETERSON BW, RUDE SA, AND ROBINSON FR. Simultaneous opposing adaptive changes in cat vestibulo-ocular reflex directions for two body orientations. *Exp Brain Res* 69: 220–224, 1987.

BLAKEMORE SJ, FRITH CD, AND WOLPERT DM. Perceptual modulation of self-produced stimuli: the role of spatio-temporal prediction. *J Cognit Neurosci* 11: 551–559, 1999.

BLAKEMORE SJ, WOLPERT DM, AND FRITH CD. Central cancellation of self-produced tickle sensation. *Nat Neurosci* 1: 635–640, 1998.

BOFF KR, KAUFMAN L, AND THOMAS JP. *Handbook of Perception and Human Performance. Sensory Processes and Perception.* New York: Wiley, 1986, vol. 1.

CASTIELLO U, PAULIGNAN Y, AND JEANNEROD M. Temporal dissociation of motor responses and subjective awareness: a study in normal subjects. *Brain* 114: 2639–2655, 1991.

ELLIOTT D, POLLOCK BJ, LYONS J, AND CHUA R. Intermittent vision and discrete manual aiming. *Percept Mot Skills* 80: 1203–1213, 1995.

ELLIOTT D, ZUBEREC S, AND MILGRAM P. The effects of periodic visual occlusion on ball catching. *J Mot Behavior* 26: 113–122, 1994.

ESKANDAR EN AND ASSAD JA. Dissociation of visual, motor and predictive signals in parietal cortex during visual guidance. *Nat Neurosci* 2: 88–93, 1999.

FALMAGNE JC, COHEN SP, AND DWIVEDI A. Two-choice reactions as an ordered memory scanning process. In: *Attention and Performance,* edited by Rabbitt P and Dornic S. New York: Academic, 1975, vol. 5, p. 20.296–20.344.

FORSSBERG H, KINOSHITA H, ELIASSON AC, JOHANSSON RS, WESTLING G, AND GORDON AM. Development of human precision grip. II. Anticipatory control of isometric forces targeted for objects weight. *Exp Brain Res* 90: 393–398, 1992.

GANDOLFO F, MUSSA-IVALDI FA, AND BIZZI E. Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843–3846, 1996.

GHAHRAMANI Z AND WOLPERT DM. Modular decomposition in visuomotor learning. *Nature* 386: 392–395, 1997.

GOODBODY SJ AND WOLPERT DM. Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825–1838, 1998.

GORDON AM, WESTLING G, COLE KJ, AND JOHANSSON RS. Memory representations underlying motor commands used during manipulation of common and novel objects. *J Neurophysiol* 69: 1789–1796, 1993.

HAY JC AND PICK HL. Gaze-contingent prism adaptation: optical and motor factors. *J Exp Psychol Hum Percept Perform* 72: 640–648, 1966.

ITO M. *The Cerebellum and Neural Control.* New York: Raven, 1984.

JEANNEROD M. *The Cognitive Neuroscience of Action.* Oxford: Blackwell, 1997.

JENMALM P AND JOHANSSON RS. Visual and somatosensory information about object shape control manipulative fingertip forces. *J Neurosci* 17: 4486–4499, 1997.

KAWATO M, FURAWAKA K, AND SUZUKI R. A hierarchical neural network model for the control and learning of voluntary movements. *Biol Cybern* 56: 1–17, 1987.

KIM J, AND SHADLEN MN. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2: 176–185, 1999.

KOHLER I. Development and alterations of the perceptual world: conditioned sensations. *Proc Aust Acad Sci* 227: 1950.

KOWLER E, MARTINS AJ, AND PAVEL M. The effect of expectations on slow oculomotor control. IV. Anticipatory smooth eye movements depend on prior target motions. *Vision Res* 24: 197–210, 1984.

KRAVITZ JH AND YAFFE F. Conditioned adaptation to prismatic displacement with a tone as the conditional stimulus. *Percept Psychophys* 12: 305–308, 1972.

MERFELD DM, ZUPAN L, AND PETERKA RJ. Humans use internal model to estimate gravity and linear acceleration. *Nature* 398: 615–618, 1999.

MIALL RC, WEIR DJ, WOLPERT DM, AND STEIN JF. Is the cerebellum a Smith Predictor? *J Mot Behav* 25: 203–216, 1993.

MIALL RC AND WOLPERT DM. The cerebellum as a predictive model of the motor system: a Smith Predictor hypothesis. In: *Neural Control of Movement,* edited by Ferrell WR and Proske U. New York: Plenum, 1995, p. 215–223.

PAILLARD J. Fast and slow feedback loops for the visual correction of spatial errors in a pointing task: a reappraisal. *J Physiol Pharmacol* 74: 401–417, 1996.

SHELHAMER M, ROBINSON DA, AND TAN HS. Context-specific gain switching in the human vestibuloocular reflex. *Ann NY Acad Sci* 656: 889–891, 1991.

SPIJKERS W AND SPELLERBERG S. On-line visual control of aiming movements? *Acta Psychol* 90: 333–348, 1995.

WELCH RB. Adaptation to space perception. In: *Handbook of Perception and Performance,* edited by Boff KR, Kaufman L, and Thomas JP. New York: Wiley-Interscience, 1986, vol. 1, p. 24.1–24.45.

WELCH RB, BRIDGEMAN B, ANAND S, AND BROWMAN KE. Alternating prism exposure causes dual adaptation and generalization to a novel displacement. *Percept Psychophys* 54: 195–204, 1993.

WOLPERT DM. Computational approaches to motor control. *Trends Cognit Sci* 1: 6: 209–216, 1997.

WOLPERT DM, GHAHRAMANI Z, AND JORDAN MI. An internal model for sensorimotor integration. *Science* 269: 1880–1882, 1995.

WOLPERT DM AND KAWATO M. Multiple paired forward and inverse models for motor control. *Neural Networks* 11: 1317–1329, 1998.

WOLPERT DM, MIALL RC, AND KAWATO M. Internal models in the cerebellum. *Trends Cognit Sci* 2: 338–347, 1998.