Adaptation to Stable and Unstable Dynamics Achieved By Combined Impedance Control and Inverse Dynamics Model

David W. Franklin,[1,2] Rieko Osu,[1] Etienne Burdet,[3] Mitsuo Kawato,[1] and Theodore E. Milner[2]

[1]*ATR Computational Neuroscience Laboratories, Kyoto 619-0288, Japan;* [2]*School of Kinesiology, Simon Fraser University, Burnaby, British Columbia V5A 1S6, Canada; and* [3]*Department of Mechanical Engineering and Division of Bioengineering, National University of Singapore, 119260, Singapore*

Franklin, David W., Rieko Osu, Etienne Burdet, Mitsuo Kawato, and Theodore E. Milner. Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J Neurophysiol* 90: 3270–3282, 2003; 10.1152/jn.01112.2002. This study compared adaptation in novel force fields where trajectories were initially either stable or unstable to elucidate the processes of learning novel skills and adapting to new environments. Subjects learned to move in a null force field (NF), which was unexpectedly changed either to a velocity-dependent force field (VF), which resulted in perturbed but stable hand trajectories, or a position-dependent divergent force field (DF), which resulted in unstable trajectories. With practice, subjects learned to compensate for the perturbations produced by both force fields. Adaptation was characterized by an initial increase in the activation of all muscles followed by a gradual reduction. The time course of the increase in activation was correlated with a reduction in hand-path error for the DF but not for the VF. Adaptation to the VF could have been achieved solely by formation of an inverse dynamics model and adaptation to the DF solely by impedance control. However, indices of learning, such as hand-path error, joint torque, and electromyographic activation and deactivation suggest that the CNS combined these processes during adaptation to both force fields. Our results suggest that during the early phase of learning there is an increase in endpoint stiffness that serves to reduce hand-path error and provides additional stability, regardless of whether the dynamics are stable or unstable. We suggest that the motor control system utilizes an inverse dynamics model to learn the mean dynamics and an impedance controller to assist in the formation of the inverse dynamics model and to generate needed stability.

INTRODUCTION

Humans have exceptional abilities to move and interact with objects in the environment. When faced with novel tasks, they adapt to environmental disturbances in a way that indicates a fundamental knowledge of the mechanics of the external world (Conditt et al. 1997; Flanagan and Wing 1997; Flanagan et al. 2001; Krakauer et al. 1999; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999). Studies of individuals performing goal-directed movements in novel mechanical environments have shown that the CNS acquires internal models of the external world (Kawato 1999). However, most adaptation studies have employed paradigms that involve stable interactions with the environment, whereas many tasks that humans perform, particularly those involving tool use, are inherently unstable (Rancourt and Hogan 2001).

Adaptation to perturbations that do not compromise mechanical stability appears to involve the acquisition of an inverse dynamics model through feedback error learning (Kawato et al. 1987). However, conventional feedback error learning does not address the issue of modifying mechanical impedance to counteract mechanical instability, although this has been repeatedly observed (Akazawa et al. 1983; Burdet et al. 2001a; De Serres and Milner 1991; Milner 2002; Milner and Cloutier 1993, 1998; Milner et al. 1995).

When the dynamics do not induce mechanical instability, changes in muscle activation patterns closely follow adaptive changes in joint torques although there is excess activation, particularly in the early stages of learning (Milner and Cloutier 1993; Thoroughman and Shadmehr 1999). Agonist-antagonist muscle co-contraction was found to increase on exposure to the novel dynamics and then decrease as learning progressed. The ability to co-contract specific groups of muscles would permit selective changes to the geometry of the endpoint stiffness (Hogan 1985). Recently we were able to demonstrate that the magnitude, shape, and orientation of the endpoint stiffness of the arm can be controlled in a predictive way to compensate for environmental instabilities (Burdet et al. 2001a).

Previous research has suggested the existence of two separate motor control mechanisms: inverse dynamics models and impedance control. An inverse dynamics model is a controller that computes feedforward commands of the net joint torques for movement based on the estimated effects of internal and external dynamics. An impedance controller, in contrast, modifies the impedance of the limb by co-contraction of agonist and antagonist muscles without changing net joint torque. These two controllers can operate independently. While there is no direct evidence for (or against) separate brain mechanisms being implicated in these two types of control, they perform different functions and therefore should be considered as separate controllers. Different features of our visual experience, for example, color, contour orientation, motion, and retinal disparities, are not independently encoded at the sensory

Address for reprint requests: D. W. Franklin, ATR Computational Neuroscience Laboratories, 2-2-2 Hikaridai, Seika-cho, Soraku-gun, Kyoto 619-0288, Japan (E-mail: dfrank@atr.co.jp).

level but are represented separately in the visual cortex (Ts'o and Roe 1994). Functionally, these features of the visual world provide very different information about our environment, and so the brain processes each separately. It seems possible, therefore that functionally distinct features of motor output might be represented and processed in a similarly independent fashion by the motor system. Some evidence already exists that control of reciprocal activation and co-contraction occurs in separate areas of the cortex (Humphrey and Reed 1983). Recently both Takahashi et al. (2001) and Osu et al. (2002) have provided evidence that inverse dynamics models and impedance control operate as separate mechanisms for motor control.

The present study investigated adaptation to novel force fields, in which trajectories were initially stable or unstable, to compare features of inverse dynamics model formation and impedance control. Adaptation to the force field, in which hand trajectories were stable, could be achieved by simply modifying joint torque. An inverse dynamics model alone was sufficient to compensate for the perturbing effects of this force field (Franklin et al. 2003). Adaptation to the other force field, in which hand trajectories were unstable, required an increase in the endpoint impedance of the arm, but no change in net joint torque, i.e., only impedance control (Burdet et al. 2001a; Franklin et al. 2003). However, we hypothesized that both of these processes are generally active during learning. To test this hypothesis, we analyzed the time course of changes in muscle activation patterns as well as joint torques, during the learning of novel force fields. In particular, we were interested in comparing the extent to which the CNS used generalized muscle co-contraction in the early stages of learning to increase endpoint stiffness and how patterns of muscle activity were later refined. Two parallel processes were identified from the evolution profile of the muscle activation patterns. One was an activation process involved in increasing the endpoint stiffness of the arm by means of muscle co-contraction during the early stages of learning. The other was a deactivation process, which led to gradual reduction in muscle activity as learning progressed.

METHODS

Five healthy individuals participated in the entire study (20–34 yr of age; 1 female and 4 males; all right-handed). The institutional ethics committee approved the experiments and the subjects gave informed consent prior to participation.

Apparatus

Subjects sat in a chair and moved the parallel-link direct drive air-magnet floating manipulandum (PFM) (Fig. 1) in a series of forward reaching movements performed in the horizontal plane. Their shoulders were held against the back of the chair by means of a shoulder harness. The right forearm was securely coupled to the PFM
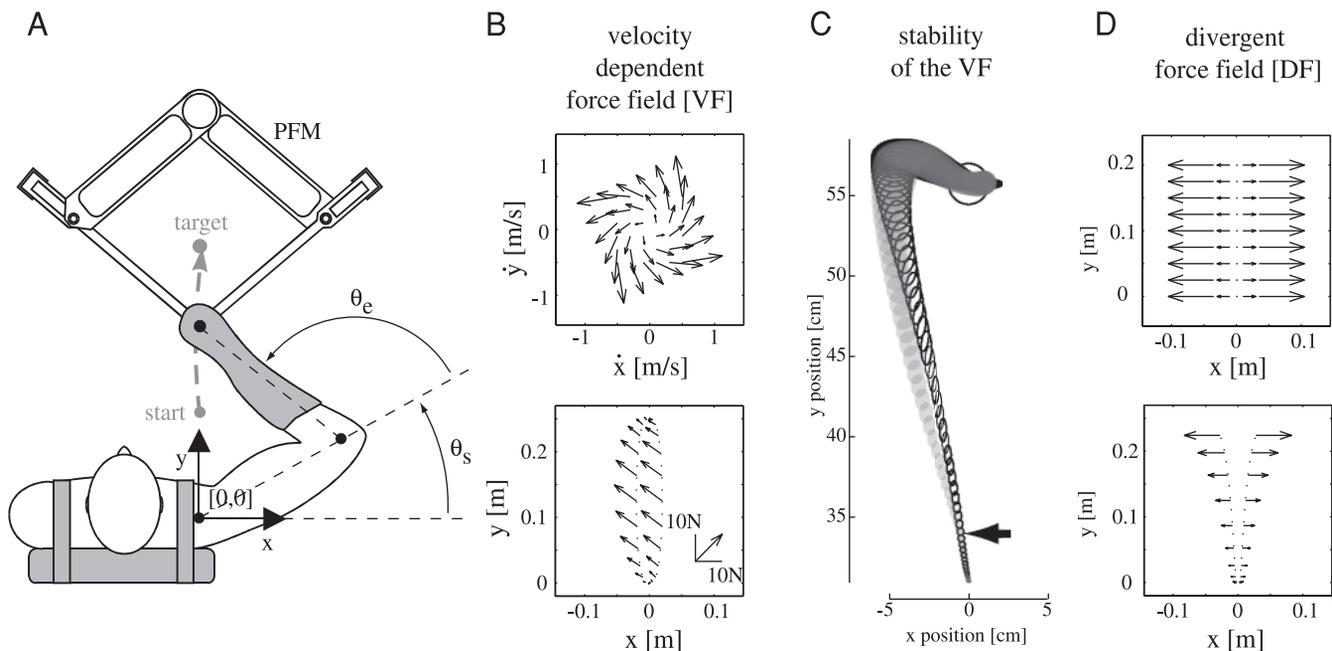

FIG. 1. *A*: experimental setup to study the adaptation to stable and unstable dynamics. Subjects were seated in a chair with their shoulders restrained by a harness, and their hand and forearm firmly attached to the parallel-link direct drive air-magnet floating manipulandum (PFM) with a thermoplastic splint. Reaching movements were performed from a start point ([x, y] = [0, 0.31] m) to a target located at ([0, 0.56] m) relative to the subjects shoulder ([0, 0]) for a total movement length of 0.25 m. The conventions for shoulder angle (θ_s) and elbow angle (θ_e) are shown. *B*: (*top*) force in the velocity-dependent force field (VF) plotted as a function of hand velocity. The vectors indicate the direction and magnitude of the forces. In the VF, the force vector rotates and increases in magnitude with velocity. *Bottom*: force vectors when hand trajectories (· · ·) are slightly deviated from a straight line along the *y* axis to the left and right. In the VF, the applied force vectors are similar despite trajectory variations. *C*: stability of the VF environment. Before effect trials [random VF trials in a majority of null force field (NF) trials] with (◎) and without (●) a force perturbation (15 N for 25 ms) were recorded. The onset of the force perturbation (100 ms after start) is shown by the black arrow. Both series of trajectories converge to the same path demonstrating the stability of the interaction between the arm and the VF. *D*: (*top*) force in the divergent force field (DF) plotted as a function of the hand position. The force vector increases in magnitude with the distance from the *y* axis. *Bottom*: force vectors when hand trajectories (· · ·) are slightly deviated from a straight line along the *y* axis to the left and right. The DF amplifies trajectory variations.

using a rigid custom-molded thermoplastic cuff. The cuff immobilized the wrist joint, permitting movement of only the shoulder and elbow joints. The subjects' right forearm rested on a support beam projecting from the handle of the PFM. Motion was, therefore limited to a single degree of freedom at the shoulder and at the elbow. The manipulandum and setup are described in detail elsewhere (Gomi and Kawato 1996, 1997).

Force fields

The experiment examined trajectory and electromyographic (EMG) adaptation in two force fields (Fig. 1, *B* and *D*): a velocity-dependent force field (VF), in which hand trajectories were stable from the outset, and a position-dependent (divergent) force field (DF), in which hand trajectories were initially unstable. Results were compared with those in a null field (NF). Details of the implementation and protocol have been described elsewhere (Burdet et al. 1999, 2001a). Only a general overview is given here. The force (F_x, F_y) (in N) exerted on the hand by the robotic interface in the VF was implemented as

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \chi \begin{bmatrix} 13 & -18 \\ 18 & 13 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \tag{1}$$

where ($\dot{x}$, $\dot{y}$) is the hand velocity (m/s) and the scaling factor, χ, was adjusted to the subject's strength ($2/3 \leq \chi \leq 1$). The stability of the initial trajectories was tested by applying small perturbations during before effect trials and comparing the resulting trajectories (Fig. 1*C*). While subjects were performing movements in the NF, the force field was switched to the VF on random trials (before effects). On some of these trials, a brief triangular force pulse (25-ms duration) with an amplitude of 15 N was applied to the hand 100 ms after movement onset. Both the perturbed and unperturbed before effect trajectories are perturbed by the force field. However, all trajectories converge illustrating that they are stable. This result is consistent with the defining characteristic of Lyapunov stability that the addition of a small perturbation does not produce divergent behavior.

The DF produced a negative elastic force perpendicular to the target direction with a value of zero along the *y* axis, i.e., no force was exerted when trajectories followed the *y* axis, but the hand was pushed away whenever it deviated from the *y* axis. The DF was implemented as

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \tag{2}$$

where the *x* component of the hand position was measured relative to the shoulder joint. $\beta = $ (300 to 500; N/m) was adjusted for each subject so that it was larger than the stiffness of the arm measured in NF movements so as to produce instability. For safety reasons, the DF force field was inactivated if the subjects' trajectory deviated more than 3 cm from the *y*-axis. Both force fields were inactivated once the subject reached the target position.

Learning

All subjects practiced in the NF on at least 1 day prior to the experiment. These training trials were used to accustom the subjects to the equipment and to the movement speed and accuracy requirements. Subjects were randomly assigned to one of two groups. Group 1 adapted to the DF on one day and adapted to the VF on another day, whereas group 2 adapted to the fields in reverse order.

Subjects first practiced in the NF until they had achieved 50 successful trials. Successful trials were those which ended inside a 2.5-cm-diameter target window within the prescribed time (0.6 ± 0.1 s). All movements were recorded whether successful or not. The movement distance was 0.25 m. Movements were self-paced so subjects were able to rest between movements if they wished. At the completion of 50 successful trials, the force field was activated. No

information was given to the subjects as to when the force field trials would begin. Subjects then practiced in the force field until achieving 75 successful trials. They took a short break and then performed 100 more movements, 20 of which were random trials in the NF. The NF trials were called after effects and were recorded to confirm that subjects had adapted to the force field.

Hand-path errors

The adaptation to the force fields was quantified by calculating the error relative to a straight line joining the centers of the start and target circles. The absolute hand-path error

$$S(|e_x|) = \int_{t=t_0}^{t_f} |x(t)||\dot{y}(t)| dt \tag{3}$$

represents the area between the actual movement path and the straight line. The signed hand-path error, defined as

$$S(e_x) = \int_{t=t_0}^{t_f} x(t)|\dot{y}(t)| dt \tag{4}$$

is a measure of the mean directional extent by which the path deviates from the straight line. Hand-path errors were calculated from the start time, t_0 (75 ms prior to crossing a hand-velocity threshold of 0.05 m/s), to the termination time, t_f (when curvature exceeded 0.07 mm^{-1}) (Pollick and Ishimura 1996).

The hand-path error for each subject was fit with an exponential curve using a least-square error method. This fitted error was expressed as

$$S(t) = Ae^{-t/\tau} + C \tag{5}$$

where A is the gain of the exponential process, τ is the time constant, C is the constant error, and t refers to the trial number.

Torque estimation

Time varying muscle torque at the shoulder and elbow was computed using the equations of motion for a two-link planar arm (cf. Hollerbach and Flash 1982). However, the equations used also include the contributions to joint torque from external forces applied at the hand. The joint torque was calculated as

$$\tau_s = \ddot{\theta}_s(2X \cos(\theta_e) + Y + Z) + \ddot{\theta}_e(X \cos(\theta_e) + Y) - \dot{\theta}_e^2 X \sin(\theta_e)$$
$$- 2\dot{\theta}_s\dot{\theta}_e X \sin(\theta_e) - (l_1 \sin(\theta_s) + l_2 \sin(\theta_s + \theta_e))F_x$$
$$+ (l_1 \cos(\theta_s) + l_2 \cos(\theta_s + \theta_e))F_y \tag{6}$$

$$\tau_e = \ddot{\theta}_e Y + \ddot{\theta}_s(X \cos(\theta_e) + Y) + \dot{\theta}_s^2 X \sin(\theta_e) - l_2 \sin(\theta_s + \theta_e)F_x$$
$$+ l_2 \cos(\theta_s + \theta_e)F_y$$

where

$$X = m_2 l_1 c_{m2} + m_c l_c c_{mc}$$

$$Y = I_2 + m_2 c_{m2}^2 + I_c + m_c c_{mc}^2$$

$$Z = I_1 + m_1 c_{m1}^2 + (m_2 + m_c)l_1^2$$

τ is joint torque, θ is joint angle (defined as in Fig. 1), I is moment of inertia about the center of mass (c_m) of the segment, l is segment length, and m is segment mass. The subscript 1 refers to the upper arm, 2 to the forearm, s to the shoulder, e to the elbow, and c to the wrist cuff. The mass and inertia of the subject's arm segments were estimated from the weight and segment lengths of each subject based on anthropometric scaling relations (Winter 1990).

Joint torque during learning was further quantified by estimating the variation in torque on a particular trial relative to the joint torque

after adaptation to the force field. The learned joint torque profile was calculated using the mean of the last 20 successful trials during the learning. The absolute torque error ($|\tau_{error}|$) represents the total difference between the joint torque on a particular trial and the mean final joint torque both in terms of extent and timing. It is calculated as

$$|\tau_{i,error}^{j}| = \int_{t=t_0}^{t_f} |\tau_i^j(t) - \tau_m^j(t)| dt \qquad (7)$$

where τ_i is the joint torque on a given trial, τ_m is the mean joint torque after adaptation, the superscript j refers to either the shoulder or elbow joint, t_0 is the time of movement onset, and t_f is the end of the time of interest. The torque error was estimated using a 1 s interval. The absolute torque error was fit with an exponential curve (*Eq. 5*) using least-squares similar to the hand-path errors.

To examine the development of the torque during the first few trials in the DF, we used two measures of torque change relative to the torque in the NF. The first measure, absolute torque development, was calculated as in *Eq. 7* but with τ_m referring to the mean joint torque in the NF. A second measure termed signed torque development ($\tau_{develop}$) was calculated as

$$\tau_{i,develop}^{j} = \int_{t=t_0}^{t_f} (\tau_i^j(t) - \tau_m^j(t)) dt \qquad (8)$$

where τ_m again refers to the mean joint torque in the NF. Both the absolute and signed torque development were calculated over the initial 400 ms to avoid the influence of corrective movements or the safety boundary.

Electromyography

Surface EMG activity of six arm muscles was recorded using pairs of silver-silver chloride surface electrodes during the learning sessions. The electrode locations were chosen to maximize the signal from a particular muscle while avoiding cross-talk from other muscles. The skin was cleansed with alcohol and prepared by rubbing in electrode paste. This was removed with a dry cloth, and pre-gelled electrodes were then attached to the skin with tape. The spacing between the electrodes of each pair was approximately 2 cm. The impedance of each electrode pair was tested to ensure that it was less than 10 kΩ.

The activity of two monoarticular shoulder muscles, pectoralis major and posterior deltoid, two biarticular muscles, biceps brachii and long head of the triceps, and two monoarticular elbow muscles, brachioradialis and lateral head of the triceps, was recorded. The EMG signals were analog filtered at 25 Hz (high-pass) and 1.0 kHz (low-pass) using a Nihon Kohden amplifier (MME-3132) and then sampled at 2.0 kHz. All comparisons between force field EMG and NF EMG involved data recorded on the same day without removal of the electrodes. EMG was aligned on the movement onset and averaged over 20 trials to visually compare changes that had occurred during learning.

To quantify changes during learning, the root-mean-square (rms) value of the EMG was calculated from 100 ms before movement onset until 350 ms after movement onset to include all the early movement associated activity. The muscle activity during this time will include both the feedforward command and the reflex responses. We did not extend the time interval further to avoid the confounding effects of voluntary corrective actions that redirect the limb toward the target after being perturbed by the force field during the early phase of learning. In all cases, the recorded EMG data were characterized by an initial increase in activity followed by a more gradual decrease. Both processes appeared to be exponential. To compare the time course of the changes in EMG with the time course of other kinematic and dynamic parameters, we fit this data with a model. Our model expressed EMG as a function of trial number using a double exponential with four free parameters

$$EMG(t) = A(1 - e^{-(t/\tau_1)}) - B(1 - e^{-(t/\tau_2)}) + C \qquad (9)$$

A and B represent the gains of the two exponential processes, τ_1 and τ_2 represent the time constants, C is EMG at *trial 0* (fixed as the mean of the EMG in the NF), and t is the trial number. One process (the 1st term) represents an exponential increase in EMG in response to exposure to the force field. It is bounded by the maximal level of activation that is possible ($A + C$) that has been shown to be significantly lower in co-contraction than reciprocal activation (Milner et al. 1995). The other process (the 2nd term) represents an exponential decrease in EMG. For comparisons across subjects, the EMG was normalized such that the mean NF EMG (C) was equal to a value of one. This model corresponds well with the recent work of Osu et al. (2002) that describes EMG activity during motor learning as a gradually decreasing function with occasional increases or decreases that are related to the error in previous trials. To provide further support for the choice of this model, Akaike's Information Criterion (AIC) (Akaike 1974) was calculated for this and two simpler models (APPENDIX).

The time course of the EMG was compared with the time course of the hand-path error and absolute joint torque error. The hand-path error was calculated by one of two methods. In the DF, the absolute hand-path error $S(|e_x|)$, a measure of the area between the actual movement path and the straight line joining the start and end targets, was used. In the VF, the signed hand-path error $S(e_x)$, a measure of the mean directional extent by which the path deviates from the straight line, was used. Although different representations of the error were used for the DF and VF, we could have used absolute hand-path error for both DF and VF without changing any of the results or conclusions. The hand-path error and joint error for each subject were fit with an exponential curve using a least square error method. The time constants of these changes were then compared with the time constants of the change in EMG during the adaptation process using *t*-tests with a significance level of 0.05. These statistics were performed on the reciprocal of the time constants, i.e., the rates, to reduce the variance of processes that exhibited little change over time (long time constants). All time constants were significantly different from zero using a *t*-test at the 0.05 level ($P < 0.00001$).

Timing of reflex responses

To find an appropriate interval for EMG analysis, the onset of voluntary responses representing feedback correction in the VF and DF was examined using before effects. Movements were performed in the NF. On randomly selected trials this was changed to the force field to elicit reflex and voluntary reactions to the imposed force field. A total of 80 NF and 20 force field trials was recorded for each force field. A comparison of the before effect EMG to that of the NF trials allowed us to determine the onset of the corrective responses (Fig. 2). These generally occurred more than 200 ms after the onset of movement except in the case of the posterior deltoid muscle for which the responses were faster. In this case, a small early corrective response could be seen as early as 150 ms after the onset of movement.

RESULTS

The two force fields investigated in this study produced distinctive perturbations of the trajectories prior to adaptation. Initial trials in the VF consistently perturbed the trajectories to the left (Fig. 3*A*), but subjects quickly learned to compensate for the force field and soon began to make straighter movements. By the 25th trial, the trajectories were relatively straight and consistently reached the final target position. The trajectories of later trials were similar. The mean signed hand-path
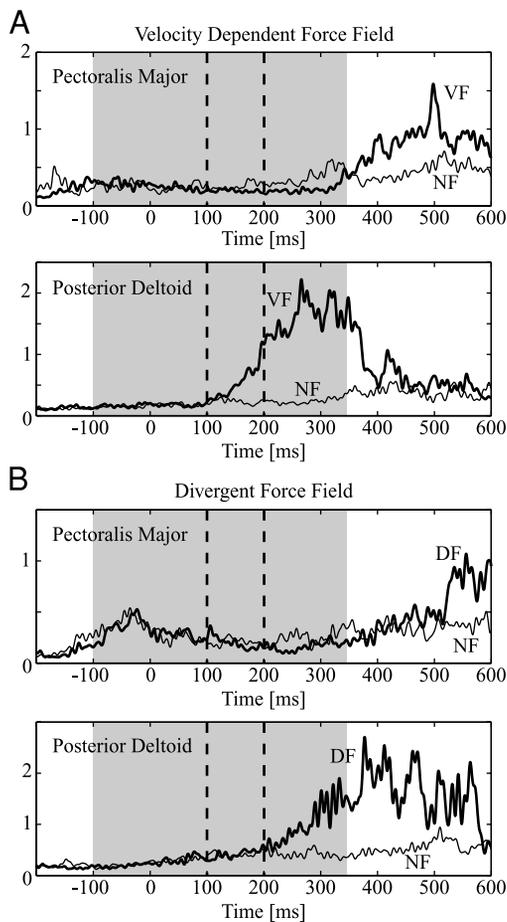
FIG. 2. Timing of corrective (reflexive and voluntary) electromyographs (EMG) in the VF (*A*) and DF (*B*). On random trials in a majority of NF movements, either the VF or DF (in separate experiments) was applied to elicit the initial corrective response to the field without learning (before effects). The timing of the corrective responses could then be seen by comparing the force field EMG (thick line) to the NF EMG (thin line). Each line represents the mean of 10 trials that have been rectified and low-pass (zero phase) filtered with a 5th-order 100-Hz cutoff Butterworth filter. Traces are shown for the muscle with the fastest response (posterior deltoid) and its antagonist muscle (pectoralis major). Except for the posterior deltoid muscle in the VF, no corrective EMG was seen prior to 200 ms, and in most muscles no activity was seen prior to 350 ms.

error illustrates the large reduction in trajectory error over the first 10 movements (Fig. 3*B*). The signed hand-path error was reduced almost to zero, with relatively little variability, by about trial 30. The absolute hand-path error changed in a similar fashion. Movements in the DF were at first perturbed either to the right or the left (Fig. 3*C*), depending on the initial deviation in the path. However, again subjects were able to adapt to the force field. After the 25th trial, subjects were able to successfully complete the task on most trials, exhibiting straight trajectories to the final target. After learning, trajectories in both the VF and DF were similar to NF trajectories. The mean signed error across subjects and its SD achieved their minimum value around trial 30 in the DF (Fig. 3*D*), similar to the VF. The trajectory of the initial trials in the DF was not perturbed as much as in the VF because of the safety zone. Nevertheless, the mean absolute error decreased only gradually over the first 75 trials. In both the VF and the DF, the signed error tended to be negative, as overall, subjects made move-

ments which were slightly biased to the left. This is consistent with NF movements that were also slightly biased to the left.

During learning, subjects adapted to the novel forces applied by the VF by modifying the joint torques (Fig. 4, *A–D*). After adaptation, the shoulder joint torque became an extensor torque throughout the entire movement and increased to several times the value in the NF trials. At the elbow, only a small change in torque in the extensor direction, during the second half of the movement, was required. To adapt to the VF, subjects gradually changed both the amplitude and shape of their joint torque profile until they had compensated for the force field's effects. Early in learning, the torque varied both above and below the final adaptation profile. In particular, early in the trial the torque tended to overshoot the final torque profile, whereas later in the movement, the opposite occurred. The early effect is due to the force applied by the force field and the corresponding changes in torque produced by changes in hand trajectory. The later effect is likely the result of reflexive feedback and voluntary correction in response to the disturbance. As subjects gained more experience performing movements in this force field, the amplitude and shape of the torque profile were adjusted until trajectories were no longer disturbed by the force field. The variability in the joint torque was large among the early trials (Fig. 4, *A* and *B*). As subjects modified the torque profile, they also reduced the trial-to-trial variability in the joint torque. The absolute torque error of all subjects decreased quickly during learning with most of the adaptation occurring in the first 30–40 trials (Fig. 4, *C* and *D*). However, subjects continued to decrease joint torque errors by small amounts as their performance became more skilled and less variable. While subjects were able to reduce their hand-path trajectory errors close to the final level by approximately the 20th trial during learning, changes in the joint torque profiles indicate that further updating and refining of the feedforward commands was still taking place throughout the learning process.

Early trials in the DF were characterized by unstable trajectories that varied either to the left or the right of the straight line joining the initial and final targets. Adaptation to this force field did not require a change in the net joint torques. The joint torque after adaptation to the force field was similar to that in the NF (Fig. 4, *E* and *F*). The moving average of the joint torque (over 5 consecutive trials) during the learning process was also quite similar to that in the NF, although the SD remained high throughout the learning process. The computed joint torque took into account the measured hand force so the effect of any force applied to the hand by the DF due to small fluctuations in the hand-path was included. Unlike the adaptation to the VF, the absolute torque error remained fairly constant (Fig. 4, *G* and *H*).

When subjects were initially presented with the DF, their trajectories deviated to the left or the right of the target, and they rarely completed the movement successfully to the target. After learning, however, subjects were able to produce successful movements to the target. Movement along the target trajectory in the DF did not require any change in the joint torque compared with the NF. However, during the first few movements the joint torque varied greatly from trial to trial (Fig. 5). In particular, it tended to vary from an extensor torque to a flexor torque and back again. This alternating pattern, which can also be seen in the signed torque development plots
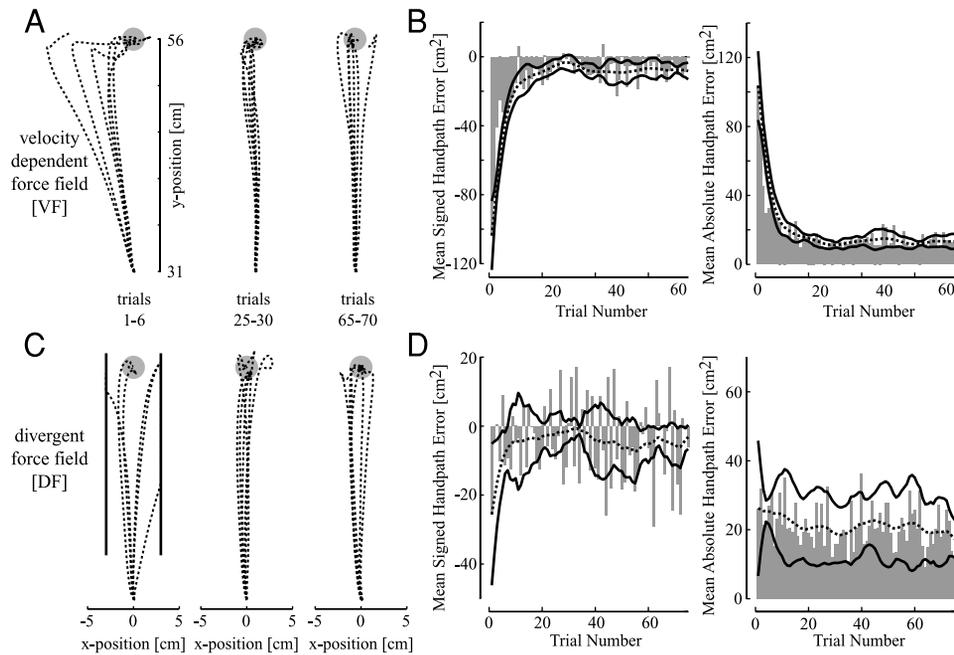
FIG. 3. Movements in the VF (*A*) and DF (*C*). Movements are shown for the initial movements in the force field (trials 1–6), the early portion of learning (trials 25–30), and the late portion of learning (trials 65–70). The black lines on either side of trials 1–6 for the DF indicate the safety boundary, outside of which the field was turned off for safety reasons. *B* and *D*: the bar graphs represent signed hand-path error (*left*) and absolute hand-path error (*right*) averaged over all subjects during the first 75 trials of learning in the VF (*B*) and DF (*D*). The middle dotted line indicates the mean values smoothed with an 8 point moving average. The solid lines (—) represent the SD about this mean (5 point moving average).
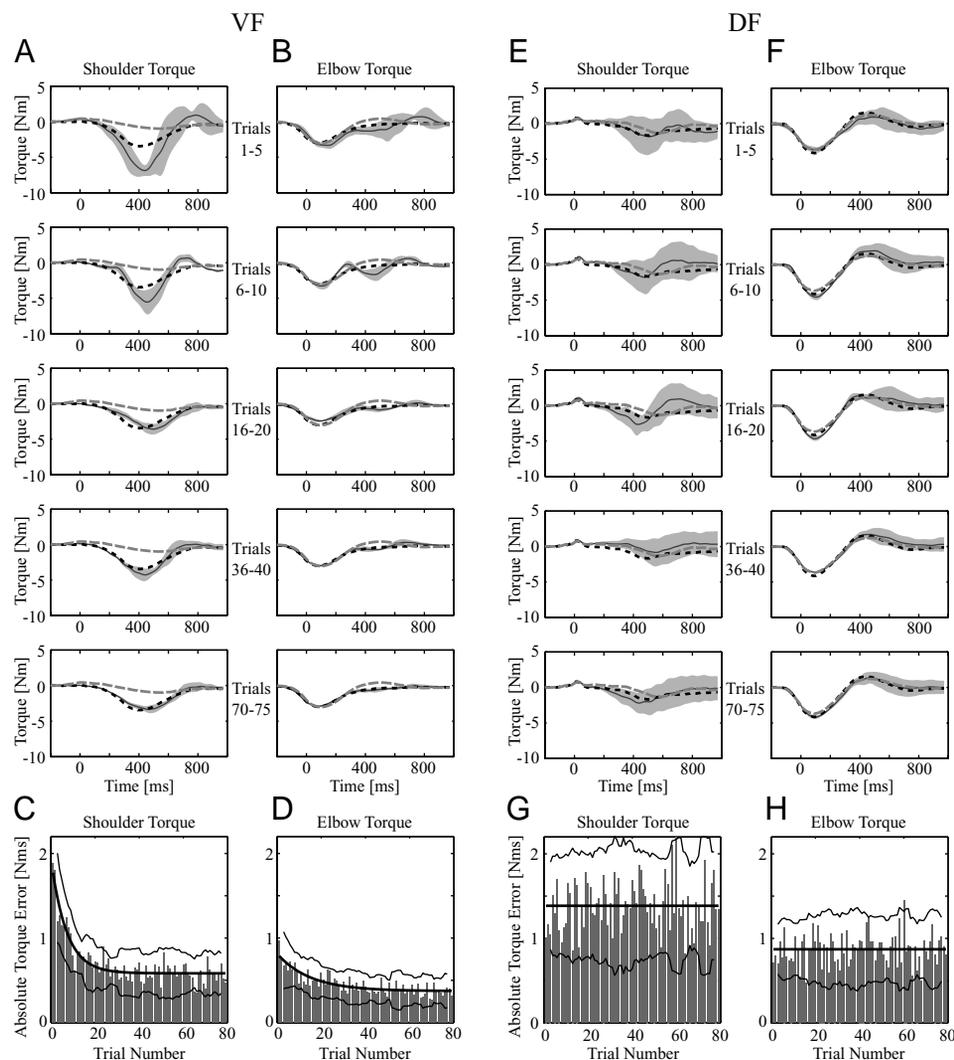

FIG. 4. Change in shoulder and elbow torque during learning in the VF (*left*) and DF (*right*). VF (*left*), *A*: mean shoulder torque during 20 NF movements (gray dotted line), during 20 movements after adaptation to the VF (black dotted line) and during the adaptation process (5 trials; solid gray line) along with SDs (light gray area). Individual plots are shown for 5 separate sets of trials during the learning process. Data are shown for 1 subject. *B*: mean elbow torque during learning in the VF. Data are plotted as in *A*. *C*: absolute shoulder torque error averaged over all subjects during learning. Absolute shoulder torque is the summation of the difference between the torque on a given trial and the mean of 20 successful trials after full adaptation to the force field. *D*: absolute elbow torque error during learning averaged over all subjects. DF (*right*). *E*: mean shoulder torque during 20 NF movements (gray dotted line), during 20 movements after adaptation to the DF (black dotted line) and during the adaptation process (5 trials; solid gray line) along with SDs (light gray area). Individual plots are shown for 5 separate sets of trials during the learning process. F: Mean elbow torque during learning in the DF. Data are plotted as in *E*. G: absolute shoulder torque error averaged over all subjects during learning. *H*: absolute elbow torque error during learning averaged over all subjects.
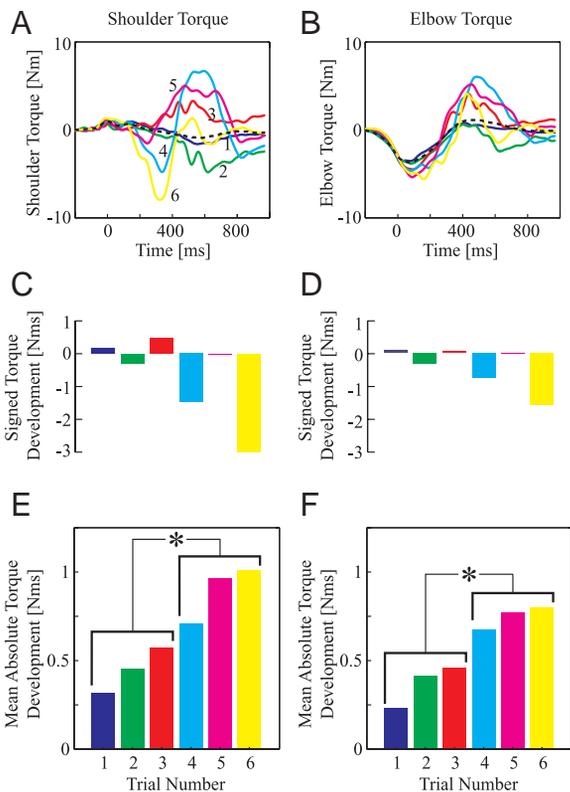
FIG. 5. Early changes in shoulder and elbow torque during learning in the DF. *A* and *B*: estimated shoulder torque (*A*) and elbow torque (*B*) during the 1st 6 trials of learning in the DF for 1 subject (colored lines). The level of joint torque during movements in the NF is shown with a black dotted line (mean of 20 successful trials). *C* and *D*: the signed torque error for the trials shown in *A* and *B,* respectively. Signed torque development is the sum of the difference between the torque on a given trial and the mean torque in the NF over the first 400 ms after movement onset. *E* and *F*: the absolute torque development for the 1st 6 trials averaged across all subjects. Absolute torque development is the sum of the absolute difference between the torque on a given trial and the mean torque in the NF (1st 400 ms after movement onset).

(Fig. 5, *C* and *D*), did not necessarily occur on successive trials but was consistently seen on some time scale for every subject. The mean absolute torque development across all subjects increased progressively for the first six trials (Fig. 5, *E* and *F*). This occurred both for shoulder and elbow torque. An ANOVA was performed comparing the first three trials (trials 1–3) to the next three trials (trials 4–6) with subjects as a random variable. The mean absolute torque development for the second three trials was found to be significantly larger at the 0.05 level than for the first three trials for both the shoulder torque ($P = 0.033$) and the elbow torque (0.001). This indicates that subjects initially responded to the disturbing effects of the force field by trying to modify the net joint torques to reduce the error. However, changing the net joint torques modifies the hand-path trajectory, which generally causes larger perturbing forces to be applied by the DF.

Adaptation in the VF required a modification of the net joint torques at the shoulder and elbow. Specifically, adaptation produced a change in net extensor joint torque at the shoulder throughout the movement and a small extensor moment at the elbow late in the movement. Therefore we would expect to see increased activity in muscles contributing to these joint torques, particularly the posterior deltoid and long head of the

triceps. After adaptation to the DF, the joint torques were not different from those in the NF. Instead, as we have shown previously, subjects modified the endpoint impedance of the limb (Burdet et al. 2001a). We therefore expected to see increased activity in one or more muscle pairs contributing to increased co-contraction. The expected changes in muscle activation patterns were confirmed from the EMG after learning (Fig. 6). In the VF, the EMG increased predominantly in the posterior deltoid and long head of the triceps muscles, the two muscles contributing the extensor torque at the shoulder needed to compensate for the force field. In contrast, in the DF the muscle activity increased in both muscles of all antagonist pairs. Because the endpoint forces and joint torques were the same in the DF and the NF, this increase in muscle activation represented balanced co-contraction.

After subjects experienced the VF for the first time, the muscle activity increased dramatically over the next few trials particularly for the posterior deltoid and long head of the triceps. This initial increase in EMG activity gradually diminished as learning proceeded. The change in EMG activity as learning progressed was fit with the double exponential function of *Eq. 9,* which represents EMG as a function of trial
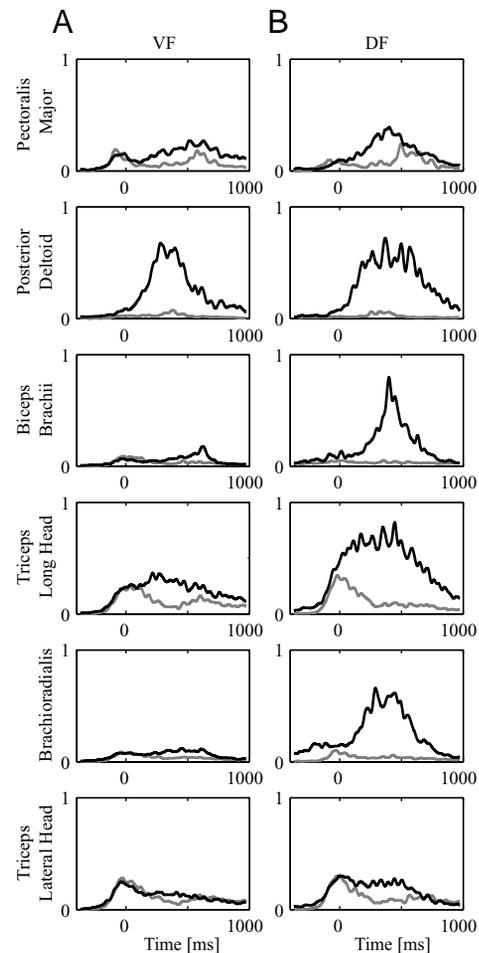

FIG. 6. Muscle activity after adaptation to the VF (*A*) and DF (*B*) shown for 1 subject. The EMG activity of 6 arm muscles is shown for the NF (gray), compared with either the VF (*A*) or DF (*B*) (solid black). The EMG was smoothed using a 75-point (37 ms) smoothing routine and averaged over 20 successful trials. The EMG is expressed in arbitrary units where the mean NF EMG level for each muscle is equal in *A* and *B*.

number. In the following sections, we present detailed analysis of the EMG from the interval of −100 to 350 ms from the start of movement. However, other intervals were also examined [(−100–200 ms), (250–350 ms), (−100–600 ms)] and all gave similar results. To illustrate the effect of the interval chosen for EMG analysis, the fitted curves for EMG activity using eight different time intervals were calculated for the long head of the triceps (Fig. 7). Although the rms EMG varies depending on the interval chosen, the overall trend for the way in which the activity changed on successive trials and the rate constants for the change in activity varied little. From this figure, we can see that the choice of interval is not critical. Our analysis using the interval −100–350 ms included both pre-programmed feedforward motor commands and feedback responses to the perturbations. However, based on before effect EMG (Fig. 2), we determined that the earliest response occurred at about 150 ms after movement start for the long head of the triceps and posterior deltoid, but for most muscles, no significant change in activity was seen prior to 350 ms. Analysis over this shorter interval (−100–200 ms) for the muscles that exhibited early responses gave similar results to the longer interval.

The evolution of muscle activity with practice in the VF is shown for one subject in Fig. 8. There was an increase in the EMG for all six muscles during the first few trials. A *t*-test was performed comparing the rms EMG of trials 2–21 to 20 trials in the NF for each muscle for each subject. Twenty-seven of the 30 comparisons showed significant increases in the rms EMG at the 0.01 significance level, indicating that the activity of all muscles increased early in learning. The peak EMG activity appeared to occur between the 3rd and 10th trials. With further practice, the EMG gradually decreased in all muscles to an asymptotic level. Similar to Fig. 6, the final rms EMG after adaptation was approximately the same as that of the NF for all muscles except those producing extensor torque at the shoulder. These muscles, the posterior deltoid and the long head of the triceps, compensated for the force applied by the VF. The actual difference after learning was tested using the same technique as for early learning. A *t*-test was used to compare the rms EMG of the final 20 trials in the VF to 20 trials in the NF for each muscle for each subject with the significance level set at 0.01. For the posterior deltoid and long head of the

FIG. 8. Evolution of the EMG for 6 arm muscles during learning in the VF for 1 subject. Subjects initially performed movements in the NF. ■, the EMG of the last 20 trials of the NF. A baseline was determined from the mean NF root-mean-square (rms) EMG (· · ·) and was set to a value of 1 for each muscle. ▨, all trials recorded in the VF. The best-fit double exponential curve (using least-squares error) is shown (—). All values are rms EMG, expressed relative to NF values, from 100 ms prior to movement initiation until 350 ms after movement initiation.

triceps, all comparisons (10/10) were significantly larger after learning compared with the NF. In the case of the lateral triceps, four of five comparisons were significantly larger after learning compared with in the NF. However, in the other three muscles, only 8 of the 15 comparisons were significantly larger in the force field.

The variation in EMG activity over trials was best fit by a double exponential process (*Eq. 9*). The results for all subjects are presented in Table 1. The activation and deactivation rates for each muscle were compared with the rate of reduction of the signed hand-path error. In the VF, subjects quickly adapted to the disturbing effects of the force field and reduced their hand-path error by two orders of magnitude within about 10 trials. Activation either led or paralleled hand-path error reduction and was characterized by a relatively short time constant, τ_1. The deactivation time constant, τ_2, was much longer. For three muscles (posterior deltoid, long head of the triceps, and brachioradialis), activation occurred more rapidly than hand-path error reduction ($P = 0.012$). Both the posterior deltoid and long head of the triceps contribute to the necessary increase in shoulder extensor torque. The activity of the brachioradialis acts to counteract the extensor torque at the elbow resulting from the increase in activity of the long head of triceps. The activity of the remaining muscles increased more slowly, at a rate not significantly different from the rate of hand-path error reduction ($P = 0.45$). The activation of both sets of muscle groups was also faster than the reduction in absolute torque error at either the shoulder ($P = 0.008$; $P = 0.026$) or elbow joint ($P = 0.008$; $P = 0.016$). Deactivation occurred at a similar rate in all muscles and proceeded much more slowly than either activation ($P = 0.007$), hand-path error reduction ($P < 0.0001$), reduction in absolute shoulder torque error ($P =$
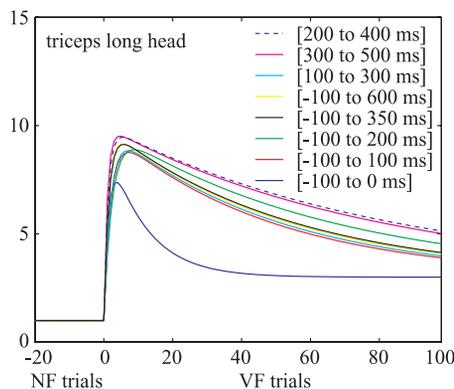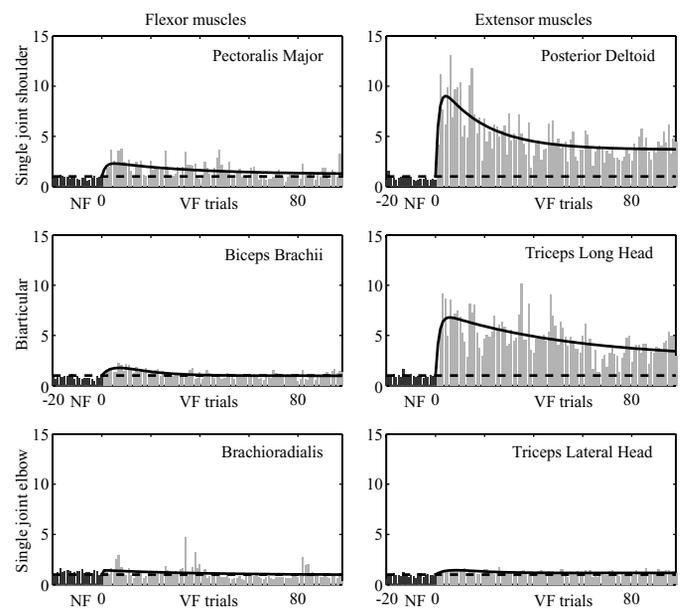

FIG. 7. Model fits to activity of the triceps long head during learning in the VF for various intervals. Fits for 8 different time intervals are shown (*Eq. 9*) to illustrate that the parameters of the model are relatively insensitive to the time interval chosen. The parameters *A, B* and the NF level have been normalized across all intervals so that the time constants can be examined.

TABLE 1. *Summary of the least-square fit to the adaptation of hand-path error and EMG during learning in the VF and DF*

	A	τ_1	B	τ_2	A-B
			A. Adaptation to VF		
EMG					
Pectoralis major	1.08 ± 1.00	2.89 ± 3.91	0.97 ± 0.95	22.27 ± 11.85	0.11 ± 0.22
Posterior deltoid	8.17 ± 6.14	1.49 ± 2.10	5.90 ± 3.44	17.69 ± 5.92	2.27 ± 2.71
Biceps brachii	1.84 ± 1.59	5.42 ± 3.55	1.74 ± 1.31	21.20 ± 17.50	0.11 ± 0.35
Triceps long head	3.75 ± 2.24	1.74 ± 1.66	3.31 ± 1.88	27.16 ± 19.55	0.44 ± 0.70
Brachioradialis	1.51 ± 1.03	0.96 ± 0.87	1.32 ± 0.86	29.16 ± 25.39	0.20 ± 0.20
Triceps lateral head	1.39 ± 1.01	3.20 ± 3.00	1.30 ± 0.92	22.57 ± 19.95	0.09 ± 0.12
Handpath error		3.32 ± 2.21			
Torque error					
Shoulder		6.85 ± 14.51			
Elbow		9.78 ± 11.39			
			B. Adaptation to DF		
EMG					
Pectoralis major	2.06 ± 3.08	6.51 ± 5.69	1.46 ± 2.15	26.68 ± 18.32	0.60 ± 0.93
Posterior deltoid	4.14 ± 2.17	8.09 ± 5.06	3.39 ± 1.72	58.40 ± 54.57	0.72 ± 0.70
Biceps brachii	7.40 ± 7.93	13.75 ± 13.53	6.35 ± 6.15	73.60 ± 66.55	1.05 ± 1.95
Triceps long head	5.92 ± 7.36	17.53 ± 22.15	4.53 ± 5.20	73.45 ± 93.11	1.39 ± 2.31
Brachioradialis	5.21 ± 6.34	6.43 ± 1.72	4.61 ± 6.67	65.02 ± 67.98	0.60 ± 0.67
Triceps lateral head	2.90 ± 1.07	8.92 ± 5.83	1.98 ± 0.71	21.06 ± 13.08	0.92 ± 0.49
Hand-path error		13.65 ± 15.55			

Values for the gain of the two exponential functions (A, B, and A-B) are expressed relative to the null force field (NF) value (normalized to 1). The mean ± SD over all subjects is shown for each muscle and for the hand-path error. EMG, electromyograph; VF and DF, velocity- and position-dependent (divergent) force field.

0.012) or reduction in absolute elbow torque error ($P <$ 0.0001). The fact that the reduction in hand-path error occurred faster than muscle deactivation indicates that subjects only gradually reduced their reliance on increased impedance to resist the force field even after hand-path error had been reduced to NF values. The largest increase in the activation process (*A* in *Eq. 9*) was seen for the posterior deltoid, followed by the long head of the triceps. However, all muscles showed a significant increase in muscle activation. The final level of adaptation, expressed as an increase from NF values (*A-B* in *Eq. 9*), indicates that the activation of most muscles, posterior deltoid and long head of the triceps being the exceptions, was reduced to near NF levels. The rates of the decrease in absolute shoulder and elbow torque errors were not found to be significantly different from each other ($P = 0.4062$). The rate of decrease in the absolute elbow torque error was found to be significantly slower than that of hand-path error ($P = 0.041$). However, the shoulder torque error was not found to be significantly different from the hand-path error at a level of 0.05 ($P = 0.063$).

During adaptation to the DF, there was also an initial increase in EMG activity of all muscles, which then gradually declined. The initial increase in EMG activity reached its peak after 20–40 trials. A *t*-test was performed comparing the rms EMG of trials 11–30 to 20 trials in the NF for each muscle for each subject. Twenty-seven of the 30 comparisons showed significant increases in the rms EMG at the 0.01 significance level indicating that all muscles increased activity early in learning. The EMG then gradually declined to an asymptotic level. Again, a *t*-test was performed comparing the rms EMG of the last 20 trials in the DF to 20 trials in the NF for each muscle for each subject. Twenty-six of the 30 comparisons showed a significant increase in the rms EMG at the 0.01 significance level after adaptation. The initial increase in activation, which represented co-contraction, occurred at a similar

rate in all muscles. The ensuing deactivation proceeded much more slowly. This suggests that the selective control of endpoint impedance was a slow process, which likely involved tuning of the relative activation of muscle pairs. The double exponential model of *Eq. 9* was again able to accurately capture the variation in EMG activity over the period of adaptation to the DF (Fig. 9).
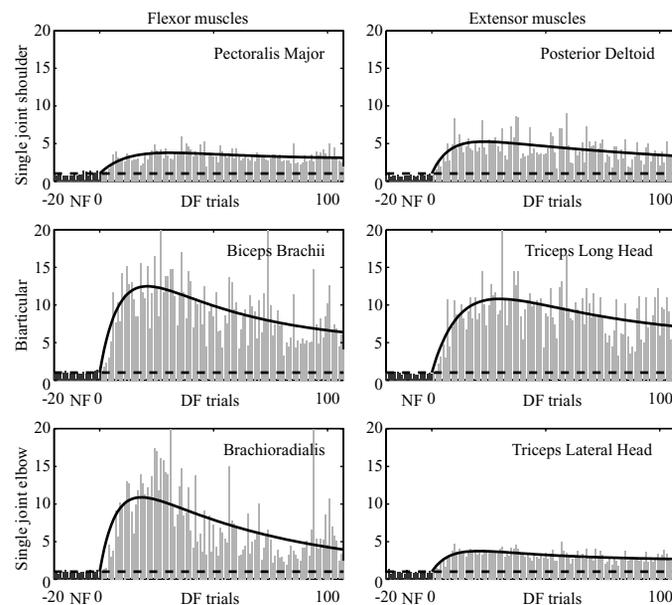

FIG. 9. Evolution of the EMG for 6 arm muscles during learning in the DF for 1 subject. Subjects initially performed movements in the NF. ■, the EMG of the last 20 trials. A baseline was determined from the mean NF rms EMG (···) and was set to a value of 1 for each muscle. ▫, all trials recorded in the DF. The best-fit double exponential curve (using least-squares error) is shown (—). All values are rms EMG, expressed relative to NF values, from 100 ms prior to movement initiation until 350 ms after movement initiation.

The activation and deactivation time courses in the DF for each muscle were compared with the time course of the hand-path error reduction. Activation proceeded at a similar rate to hand-path error reduction ($P = 0.59$), whereas deactivation proceeded much more slowly ($P < 0.0001$). The time courses were similar for the six muscles. As in the case of the VF, activation was significantly faster than deactivation ($P < 0.0001$). Values for all parameters are shown in Table 1. The gain of the activation process (A) is larger for the biarticular muscles than the single joint elbow or shoulder muscles. Interestingly, the deactivation process (B) also is largest for the biarticular muscles, even though the final level after adaptation (expressed as an increase from NF values; A-B) shows the largest increase for the biarticular muscles.

DISCUSSION

This study compared the learning processes during adaptation to a force field in which hand trajectories were stable and one in which they were initially unstable. The force fields were designed so that adaptation could be achieved solely by a change in the net joint torques in the stable case, in contrast to the unstable case, which required only a change in the limb impedance. Initial trajectories were consistently displaced in the same direction in the stable case (VF), but subjects quickly reduced their hand-path error and learned to produce straight movements. The modification of joint torques occurred more slowly. The adaptation in muscle activity as learning progressed was well described by concurrent activation and deactivation processes. Learning was characterized by a rapid increase in activation of the muscles needed to compensate for the environmental force together with generalized co-contraction, which lagged slightly behind. This generalized co-contraction was later reduced, but much more slowly than the reduction of hand-path error. In the unstable case (DF), initial

trajectories were displaced in both directions, usually ending outside a safety boundary. With practice, subjects were able to reduce the hand-path error and learned to produce straight movements to the final target location. The mean joint torques changed relatively little during the adaptation process. However, the early trials were characterized by alternating errors in joint torque from trial to trial and increasing absolute joint torque. Again, the modification of muscle activity was well described by a concurrent activation and deactivation process. Differential co-contraction occurred rapidly and matched the rate of hand-path error reduction. The later reduction in muscle activity occurred slowly, gradually falling to a level where endpoint stiffness would guarantee a normal safety margin for stability (Burdet et al. 2001a). These processes are summarized in Fig. 10.

Learning in the VF

The VF is similar to previously studied dynamical environments in which it has been suggested that inverse dynamics models are formed (Conditt et al.1997; Flanagan and Wing 1997; Flanagan et al. 2001; Krakauer et al. 1999; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994; Thoroughman and Shadmehr 1999). Early in learning, the activation of three muscles (posterior deltoid, long head of the triceps, and brachioradialis) increased rapidly. The increased activity may have comprised reflex activity arising from muscle stretch, voluntarily activation during the movement to correct for the initial displacement caused by the force field and predictive activation to counteract the expected force. The posterior deltoid and long head of the triceps contributed to the increase in shoulder extensor torque necessary to counteract the VF while the brachioradialis would have counteracted the extensor torque at the elbow resulting from the increase in activity of the long head of triceps. This adaptation is consistent with the
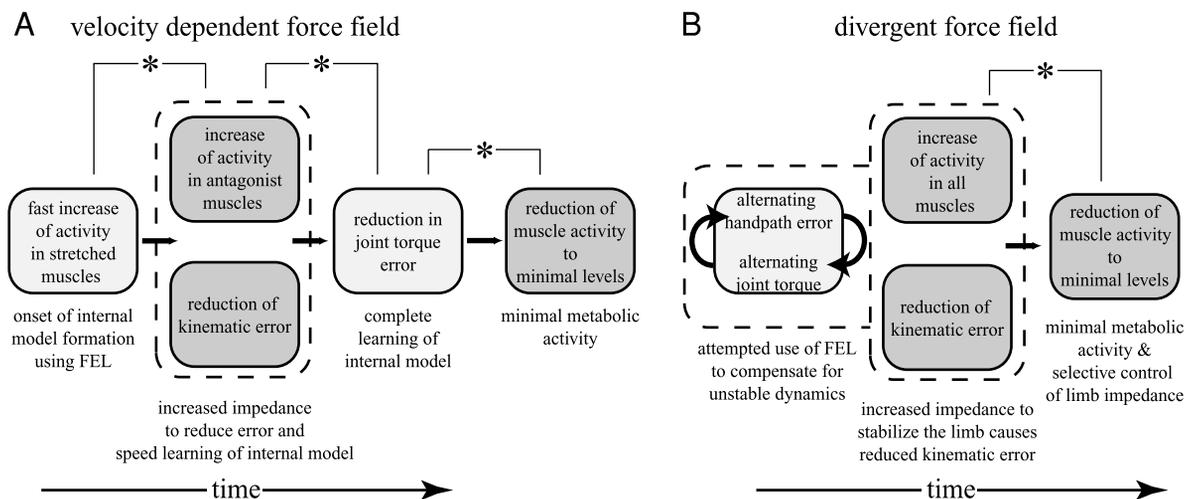

FIG. 10. Summary of the time course of events taking place during learning in the VF (*A*) and DF (*B*). Asterisks indicate a significance difference in time constants at the 0.05 level. *A*: VF. Initially, fast increases are seen in the stretched muscles suggesting quick onset of internal dynamics model formation. This is followed by the increase in activity of antagonist muscles (onset of impedance controller) and reduction in kinematic error. Later, the joint torque error is reduced, which would indicate the complete learning of the internal dynamics model (IDM). Finally muscle co-contraction is reduced to minimal levels by the impedance controller. *B*: DF. Early trials in the DF are characterized by alternation in the direction of hand-path error and joint torque which indicates that development of an IDM using feedback error learning is being attempted. Also, during the early trials, increased co-contraction occurs in all muscles (impedance control) causing a reduction in kinematic error. Finally, the muscle activity is reduced to minimize metabolic activity while ensuring the stability of the movements. This results in a selective control of the endpoint impedance.

development of an inverse dynamics model by means of feedback error learning (Kawato et al. 1987).

In addition, muscle activity increased due to agonist-antagonist co-contraction, which can be inferred from the increase in the activity of the remaining muscles. This co-contraction acts to increase the endpoint stiffness of the arm. Increased co-contraction during the learning of novel force fields is certainly not a new finding. Thoroughman and Shadmehr (1999) have also shown increased co-contraction during adaptation to stable dynamics, which they termed wasted contraction. However, we suggest that this increased contraction of antagonist muscles is not wasted but performs an essential role in the learning of novel dynamics by stabilizing the limb, resisting the disturbing effects of the force field, and allowing for the refinement of the feedforward inverse dynamics model (Osu et al. 2002). With improvement in the inverse dynamics model, co-contraction (limb stiffness) was reduced. This was associated with a slow deactivation process, which resulted in the activity of most muscles returning close to NF levels.

Learning in the DF

When subjects initially performed movements in the DF, their movements varied to either side of the mid-line from one trial to the next. This is evident from the joint torques, which alternated between extensor and flexor moments. The torque, relative to the NF, increased on each of the first six trials. This suggests that the subjects may have been attempting to incorporate the error information from the previous trial into the feedforward command for the next movement as would occur during feedback error learning of an inverse dynamics model. In the case of stable dynamics, this method works well and allows for quick adaptation. In contrast, this mechanism alone will not succeed when dynamics are unstable (Burdet et al. 2001b). If error information from the first trial was used to update the feedforward command for the next movement, that movement would be made to the opposite side of the mid-line of the force field because subjects would produce a force opposite to direction of the previous disturbance. The resulting error would be larger and opposite in sign to the previous error. As this process continued, joint torque would tend to increase and alternate between extensor and flexor moments, although the degree to which this would occur would depend on parameters such as the learning factor (cf. Scheidt et al. 2001) and the magnitude of motor noise.

However, early in the learning period there was an increase in agonist-antagonist co-contraction, which can be interpreted as an increase in endpoint stiffness. The hand-path error was reduced at about the same rate as the increase in stiffness. This provides evidence that the increased stiffness directly contributed to the reduction in hand-path error. This is not unexpected because increased endpoint stiffness counteracts the instability of the DF.

As the subject became more successful in counteracting the instability of the DF, the EMG was gradually reduced. This reduction in superfluous co-contraction would reduce metabolic energy requirements and possibly also reduce variability in motor output that tends to increase with muscle activity (Clancy and Hogan 1995; Harris and Wolpert 1998; van Galen and van Huygevoort 2000). We suggest that an interplay between two competitive processes, increased muscle activation

in response to trajectory errors, and a drive to minimize muscle activation necessary to perform the task, results in a selective control of endpoint impedance as seen after learning in the DF (Burdet et al. 2001a; Franklin et al. 2003). The final level of muscle activity, relative to the NF, is expressed as *A-B* (Table 1). The largest values are found for the biarticular muscles (long head of the triceps and biceps brachii), indicating that these muscles may play the most prominent role in the selective control of the endpoint impedance in the DF.

Impedance controller

It is clear that the CNS is able to control the impedance of the limb. Increased muscle activity and stiffness has been seen during adaptation to many types of environments (De Serres and Milner 1991; Milner 2002; Milner and Cloutier 1993; Takahashi et al. 2001; Thoroughman and Shadmehr 1999). More recently, we have shown that this impedance can be directionally tuned to the environment (Burdet et al. 2001a). However, it is not yet clear how the CNS produces optimally oriented limb impedance. We propose the existence of an impedance controller that employs a fast activation process in response to error signals and a slower deactivation process. This controller will initially increase the impedance of the limb during adaptation to any novel dynamics. It will also attempt to minimize the activation of all muscles. By employing these two opposing processes, the initial high-impedance will eventually be reduced to a minimal level necessary for stability.

It appears that the CNS engages a similar process of impedance control during the initial period of adaptation whether trajectories are inherently stable or unstable. When an additional force must be applied to the environment, as in the case of the VF, there is a natural increase in limb impedance due to muscle activation (Franklin et al. 2003). Even though stability may be guaranteed by this naturally occurring impedance, the CNS chooses to augment it by co-contraction during the earliest phase of learning. The superfluous impedance is later eliminated once the inverse dynamics model had been acquired. However, when the naturally occurring impedance is insufficient to provide stability, as in the case of the DF, the impedance controller generates a global increase in stiffness, which is selectively reduced as the optimal stiffness geometry is determined (Franklin et al. 2003). Takahashi et al. (2001) have also suggested that impedance control can coexist with the formation of inverse dynamics models for control. Robotic implementations for learning novel tasks often use high impedance as a method of achieving faster learning (Katayama et al. 1998; Sanger 1994). The increased stiffness can reduce the disturbing effects of the novel dynamics to provide better tracking of a desired trajectory during the early stages of learning. The dynamics during learning are then closer to the desired final dynamics, which will increase the speed of learning.

Learning rates for impedance control and inverse dynamics model formation

The inverse dynamics model learning and impedance learning occur simultaneously, although the former appears to proceed at a faster rate than the latter. Several lines of evidence support this conclusion. Learning in the VF, ultimately realized

by inverse dynamics model formation, is faster than learning in the DF, achieved by impedance control. Deactivation of all muscles (late phase of impedance learning) is slower than acquisition of accurate torque profiles (late phase of inverse dynamics model learning) in the VF. Activation of all muscles in the DF (early phase of impedance learning) is slower than activation of the posterior deltoid and the triceps long head in the VF (early phase of inverse dynamics model learning). In summary, the early phase of inverse dynamics model learning is almost simultaneous or slightly faster than the early phase of impedance learning, which is much faster than the late phase of inverse dynamics model learning, which, in turn, is slightly faster than the late phase of impedance learning. The interrelation of these two learning processes suggests that they could be unified in a single model.

Unified model for motor learning

The following natural extension of feedback error learning could coherently unify the two learning processes and at least qualitatively reproduce our results as well as other recent data on motor learning (Burdet et al. 2001a; Osu et al. 2002). First, centrally generated feedforward motor commands would comprise both a reciprocal component for agonist and antagonist muscle pairs (difference in muscle activation similar to net joint torque) and a co-activation component (summation of agonist and antagonist muscle activation similar to joint stiffness). This concept originated with Feldman and has been elaborated by his collaborators (Feldman 1980 a,b; Gribble et al. 1998; Latash 1992; Levin et al. 1992). Second, the feedforward co-activation signal to antagonists should increase on trials following perturbation of the hand path during early learning even when only agonist muscles are stretched. This hypothesis is required to account for activation of pectoralis major up to the fifth trial in the VF (Fig. 8) despite the fact that it was not stretched, as all the trajectories deviated to the left (Fig. 3A, trials 1–6). This mechanism also contributes to the co-activation of all muscles in the DF because trajectories deviate to left on some trials and to the right on others. Third, the feedforward co-activation signal decays with a large time constant as manifested by the deactivation time constants of all muscles. We are currently developing methods to quantitatively evaluate this conceptual model both experimentally and by means of computer simulations.

APPENDIX

The change of EMG during the learning in the force fields was characterized by an initial increase and a gradual decrease. We therefore modeled this as a double-exponential process with two functions, one contributing to each of these changes. To assess the accuracy of this model, we compared the residuals of three related models using Akaike's information criterion (Akaike 1974). The EMG at *trial 0* was set equal to the mean EMG in the NF prior to presentation of the force field. In the first model, EMG was expressed as a linear function of trial number. In the second model, EMG was expressed as an exponential function of trial number as used previously to characterize hand-path error (Burdet et al. 2001a; Flanagan et al. 1999) and suggested by the results of Thoroughman and Shadmehr (1999). The third model expressed EMG as a double exponential function of trial number with four free parameters (*Eq. 9*).

The three models were fit to the rms EMG of consecutive trials during learning in both the VF and DF. Because increasing the number of parameters of a model can over fit the data without adding any information, the residuals were used in Akaike's information criterion (AIC) (Akaike 1974) to determine if the added parameters of the double-exponential function were justified by explaining more of the variation in the data. To calculate a single representative value of AIC for each model, the data from six muscles and five subjects were combined. Because the EMG magnitude varied among muscles and subjects, the data were first normalized by dividing the residual by the sum of the fitted EMG data for each muscle and subject. The residuals, now expressed as a function of how well they fit the data, were then combined to calculate a single AIC value

$$\text{AIC} = N\left(\log 2\pi\left(\frac{1}{\text{sub}}\sum_{i=1}^{\text{sub}}\frac{1}{\text{mus}}\sum_{j=1}^{\text{mus}}\hat{\sigma}_{i,j}^2\right)+1\right)+2(m+1) \qquad (A1)$$

where N is the total number of data points, m is the number of fitted parameters in the equation multiplied by the number of muscles and subjects, and $\hat{\sigma}_2^{i,j}$ is the integrated squared error in EMG modeling for the jth muscle of the ith subject. This was performed separately for VF and DF data to determine which model fit the data for learning in each force field most appropriately.

The AIC values were calculated for each model in both the VF and the DF. In both fields, the double-exponential model fit the data better. In the VF, the lowest AIC was obtained for the double-exponential model (42,696). This was 50 less than AIC for the single exponential model and 2,023 less than AIC for the linear model. A difference of more than 2 is usually considered statistically significant (Sakamoto et al. 1986). The results for the DF were similar, with the lowest AIC (45,374) for the double exponential model, which was 209 less than AIC of the single exponential model and 1,521 less than AIC for the linear model. This supports our choice of the double-exponential model, comprising an activation and a deactivation process, to describe adaptation to both the VF and DF dynamics.

REFERENCES

Akaike H. A new look at the statistical model identification. *IEEE Trans Automat Control AC* 19: 716–723, 1974.

Akazawa K, Milner TE, and Stein RB. Modulation of reflex EMG and stiffness in response to stretch of human finger muscle. *J Neurophysiol* 49: 16–27, 1983.

Burdet E, Osu R, Franklin DW, Milner TE, and Kawato M. Measuring stiffness during arm movements in various dynamic environments. In: *Proceedings of the 1999 ASME Annual Symposium on Haptic Interfaces and Virtual Environments for Teleoperator Systems,* Nashville, TN: ASME, 1999, p. 421–428.

Burdet E, Osu R, Franklin DW, Milner TE, and Kawato M. The central nervous system stabilizes unstable dynamics by learning optimal impedance. *Nature* 414: 446–449, 2001a.

Burdet E, Tee KP, Chew CM, Franklin DW, Osu R, Kawato M, and Milner TE. Stability and learning in human arm movements. *Proceedings of the 2001 International Conference on Computational Intelligence, Robotics and Autonomous Systems,* Singapore, p. 335–361, 2001b.

Clancy EA and Hogan N. Multiple site electromyograph amplitude estimation. *IEEE Trans Biomed Eng* 42: 203–211, 1995.

Conditt MA, Gandolfo F, and Mussa-Ivaldi FA. The motor system does not learn the dynamics of the arm by rote memorization of past experience. *J Neurophysiol* 78: 554–560, 1997.

De Serres SJ and Milner TE. Wrist muscle activation patterns and stiffness associated with stable and unstable mechanical loads. *Exp Brain Res* 86: 451–458, 1991.

Feldman AG. Superposition of motor programs. I. Rhythmic forearm movements in man. *Neuroscience* 5: 81–90, 1980a.

Feldman AG. Superposition of motor programs. II. Rapid forearm flexion in man. *Neuroscience* 5: 91–95, 1980b.

Flanagan JR, King S, Wolpert DM, and Johansson RS. Sensorimotor prediction and memory in object manipulation. *Can J Exp Psychol* 55: 87–95, 2001.

Flanagan JR, Nakano E, Imamizu H, Osu R, Yoshioka T, and Kawato M. Composition and decomposition of internal models in motor learning under altered kinematic and dynamic environments. *J Neurosci* 19: RC34, 1999.

Flanagan JR and Wing AM. The role of internal models in motion planning and control: evidence from grip force adjustments during movements of hand-held loads. *J Neurosci* 17: 1519–28, 1997.

Franklin DW, Burdet E, Osu R, Kawato M, and Milner TE. Functional significance of stiffness in adaptation of multijoint arm movements to stable and unstable environments. *Exp Brain Res* In press.

Gomi H. and Kawato. Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement. *Science* 272: 117–120, 1996.

Gomi H and Kawato M. Human arm stiffness and equilibrium-point trajectory during multi-joint movement. *Biol Cybern* 76: 163–171, 1997.

Gribble PL, Ostry DJ, Sanguineti V, and Laboissiere R. Are complex control signals required for human arm movement? *J Neurophysiol* 79: 1409–1424, 1998.

Harris CM and Wolpert DM. Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.

Hogan N. The mechanics of multi-joint posture and movement control. *Biol Cybern* 52: 315–331, 1985.

Hollerbach MJ and Flash T. Dynamic interactions between limb segments during planar arm movement. *Biol Cybern* 44: 67–77, 1982.

Humphrey DR and Reed DJ. Separate cortical systems for control of joint movement and joint stiffness: reciprocal activation and co-activation of antagonist muscles. *Adv Neurol* 39: 347–372, 1983.

Katayama M, Inoue S, and Kawato M. A strategy of motor learning using adjustable parameters for arm movement. *Proceedings of the 20th Annual International Conference of the IEEE Engineering in Medicine and Biology Society.* Hong Kong: IEEE, 1998, p. 2370–2373.

Kawato M. Internal models for motor control and trajectory planning. *Curr Opin Neurobiol* 9: 718–727, 1999.

Kawato M, Furukawa K, and Suzuki R. A hierarchical neural-network model for control and learning of voluntary movement. *Biol Cybern* 57: 169–185, 1987.

Krakauer JW, Ghilardi MF, and Ghez C. Independent learning of internal models for kinematic and dynamic control of reaching. *Nat Neurosci* 2: 1026–1031, 1999.

Lackner JR and Dizio P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J Neurophysiol* 72: 299–313, 1994.

Latash ML. Independent control of joint stiffness in the framework of the equilibrium-point hypothesis. *Biol Cybern* 67: 377–384, 1992.

Levin MF, Feldman AG, Milner TE, and Lamarre Y. Reciprocal and coactivation commands for fast wrist movements. *Exp Brain Res* 89: 669–677, 1992.

Milner TE. Adaptation to destabilizing dynamics by means of muscle co-contraction. *Exp Brain Res* 143: 406–416, 2002.

Milner TE and Cloutier C. Compensation for mechanically unstable loading in voluntary wrist movement. *Exp Brain Res* 94: 522–532, 1993.

Milner TE and Cloutier C. Damping of the wrist joint during voluntary movement. *Exp Brain Res* 122: 309–317, 1998.

Milner TE, Cloutier C, Leger AB, and Franklin DW. Inability to activate muscles maximally during cocontraction and the effect on joint stiffness. *Exp Brain Res* 107: 293–305, 1995.

Osu R, Franklin DW, Kato H, Gomi H, Domen K, Yoshioka T, and Kawato M. Short- and long-term changes in joint co-contraction associated with motor learning as revealed from surface EMG. *J Neurophysiol* 88: 991–1004, 2002.

Pollick FE and Ishimura G. The three-dimensional curvature of straight-ahead movements. *J Mot Behav* 28: 271–279, 1996.

Rancourt D and Hogan N. Stability in force-production tasks. *J Mot Behav* 33: 193–204, 2001.

Sakamoto Y, Ishiguro M, and Kitagawa G. *Akaike Information Criterion Statistics.* Dordrecht: Kluwer Academic, 1986.

Sanger TD. Neural network learning control of robotic manipulators using gradually increasing task difficulty. *IEEE Trans Robotics Automat* 10: 323–333, 1994.

Scheidt RA, Dingwell JB, and Mussa-Ivaldi FA. Learning to move amid uncertainty. *J Neurophysiol* 86: 971–985, 2001.

Shadmehr R and Mussa-Ivaldi FA. Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208–3224, 1994.

Takahashi CD, Scheidt RA, and Reinkensmeyer DJ. Impedance control and internal model formation when reaching in a randomly varying dynamical environment. *J Neurophysiol* 86: 1047–1051, 2001.

Thoroughman KA and Shadmehr R. Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19: 8573–8588, 1999.

Ts'o DY and Roe AW. Functional compartments in visual cortex: segregation and interaction. In: *The Cognitive Neurosciences,* edited by Gazzaniga MS. Cambridge, MA: MIT Press, 1994, p. 325–337.

van Galen GP and van Huygevoort M. Error, stress and the role of neuro-motor noise in space oriented behaviour. *Biol Psychol* 51: 151–171, 2000.

Winter DA. *Biomechanics and Motor Control of Human Movement.* New York: Wiley, 1990.