

Different Mechanisms Involved in Adaptation to Stable and Unstable Dynamics

Rieko Osu,^{1,*} Etienne Burdet,^{1,2,*} David W. Franklin,^{1,3} Theodore E. Milner,³ and Mitsuo Kawato¹

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

Submitted 27 January 2003; accepted in final form 31 July 2003

Osu, Rieko, Etienne Burdet, David W. Franklin, Theodore E. Milner, and Mitsuo Kawato. Different mechanisms involved in adaptation to stable and unstable dynamics. *J Neurophysiol* 90: 3255–3269, 2003; 10.1152/jn.00073.2003. Recently, we demonstrated that humans can learn to make accurate movements in an unstable environment by controlling magnitude, shape, and orientation of the endpoint impedance. Although previous studies of human motor learning suggest that the brain acquires an inverse dynamics model of the novel environment, it is not known whether this control mechanism is operative in unstable environments. We compared learning of multijoint arm movements in a “velocity-dependent force field” (VF), which interacted with the arm in a stable manner, and learning in a “divergent force field” (DF), where the interaction was unstable. The characteristics of error evolution were markedly different in the 2 fields. The direction of trajectory error in the DF alternated to the left and right during the early stage of learning; that is, signed error was inconsistent from movement to movement and could not have guided learning of an inverse dynamics model. This contrasted sharply with trajectory error in the VF, which was initially biased and decayed in a manner that was consistent with rapid feedback error learning. EMG recorded before and after learning in the DF and VF are also consistent with different learning and control mechanisms for adapting to stable and unstable dynamics, that is, inverse dynamics model formation and impedance control. We also investigated adaptation to a rotated DF to examine the interplay between inverse dynamics model formation and impedance control. Our results suggest that an inverse dynamics model can function in parallel with an impedance controller to compensate for consistent perturbing force in unstable environments.

INTRODUCTION

To manipulate objects it is necessary to interact with the environment. Skilled actions require that humans compensate for the resulting interaction forces. Previous studies that investigated learning in novel environments where the interaction with the arm was stable have shown that subjects learn the force necessary to compensate for the imposed dynamics in a feedforward manner (Brashers-Krug et al. 1996; Krakauer et al. 1999; Lackner and Dizio 1994; Shadmehr and Mussa-Ivaldi 1994). We recently observed learning in a divergent force field (DF), which produced an unstable interaction with the arm and showed that humans improved performance in a situation

where dynamics were unstable by learning optimal mechanical impedance of the arm (Burdet et al. 2001a). The CNS was able to predictively control the magnitude, shape, and orientation of the endpoint stiffness without varying endpoint force. These results suggest that humans can form internal models to compensate for both stable and unstable dynamics. However, the question arises whether such internal models are learned and employed in the same way for both stable and unstable dynamics.

The inverse dynamics model controls the net joint torque, which is produced by reciprocal activation of antagonistic muscle pairs (i.e., difference between agonist and antagonist muscle torque) (Fig. 1). Because of sensory motor delays, the inverse dynamics model must compute reciprocal activation commands from the desired trajectory instead of the current state. The term “inverse dynamics model” indicates a mapping from state to force, but we do not exclude the possibility that this computation is performed by a forward model (Bhushan and Shadmehr 1999; Miall et al. 1993). Impedance will also change because it is correlated with muscle activation. In addition to this obligatory change in impedance associated with inverse dynamics models, impedance can be independently controlled by generating specific central commands for co-contraction (i.e., the summation of agonist and antagonist stiffness without changing the net joint torque) (Smith 1996). Although there is relatively little direct evidence for separate brain mechanisms underlying these 2 types of control (Frysienger et al. 1984; Humphrey and Reed 1983; Milner et al. 2002), their interaction in motor adaptation has been a topic of interest in both behavioral and modeling studies (Feldman 1986; Franklin et al. 2003; Gribble and Ostry 2000; Levin et al. 1992; Osu et al. 2002; Takahashi et al. 2001; Thoroughman and Shadmehr 1999; Wang et al. 2001).

Although an inverse dynamics model should be effective in stable environments because the imposed dynamics is rather consistent over trials, it is known that humans, in some cases, adapt to novel stable dynamics by increasing impedance (Takahashi et al. 2001). Similarly, it may sometimes be possible to adapt to unstable dynamics by selecting a trajectory that evokes a consistent predictable perturbing force, so that the impedance derived from the compensating force could provide stabiliza-

*R. Osu and E. Burdet contributed equally to this work.

Address for reprint requests and other correspondence: R. Osu, ATR Computational Neuroscience Laboratories, 2-2-2, Hikaridai, Seika-cho, Sorakugun, Kyoto 619-0288, Japan (E-mail: osu@atr.co.jp).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked “advertisement” in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

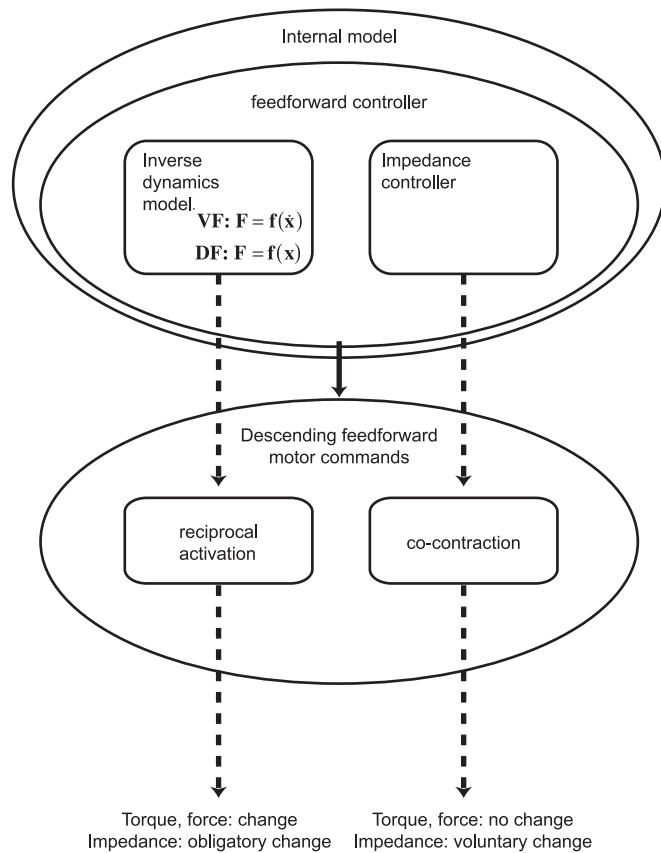


FIG. 1. Schema of an internal model. Internal model for motor system computes feedforward motor commands to compensate for external perturbation. This is done either by inverse dynamics model or impedance controller, or both. Inverse dynamics model for VF would be a mapping between hand velocity and external force, and that for DF would be a mapping between hand position and external force.

tion. However, relying on an inverse dynamics model would result in failure under unstable dynamics if the impedance associated with the compensating force along the selected trajectory was smaller than that necessary for stabilization. Therefore it is of interest to investigate when the 2 mechanisms are used selectively and when they are combined according to the environmental dynamics. In addition, it is not known whether impedance is learned in the same way as an inverse dynamics model. The “feedback error learning” strategy (Kawato 1990; Kawato et al. 1987) has been demonstrated to be effective for learning an inverse dynamics model both in computational studies and robotic implementations (Burdet et al. 1998; Kawato et al. 1988; Niemeyer and Slotine 1991) and is supported by physiological studies of the oculomotor system (Kawato 1999; Kobayashi et al. 1998; Shidara et al. 1993; Takemura et al. 2001; Yamamoto et al. 2002). However, no computational or behavioral investigation of impedance learning has been undertaken.

To address these questions, we observed learning during multijoint arm movements in the DF and in a velocity-dependent force field (VF), where interaction with the arm was stable. We also investigated the combination of an inverse dynamics model and impedance control by comparing adaptation to unstable and stable environments that incorporated a varying load bias, using a rotated divergent force field (rDF) and a rotated convergent force field (rCF). Our results provide

an accurate description of the distinct control behaviors used to adapt to stable versus unstable dynamics.

METHODS

Subjects

Nine subjects performed the learning experiment (24–34 yr of age; 5 males and 4 females). The institutional ethics committee approved the experiments and the subjects gave informed consent before participation.

Experimental setup

The movements studied were horizontal point-to-point movements away from the body (Fig. 2A). This movement direction corresponded to the y -axis in our coordinate system, with the x -axis oriented from left to right and the origin at the center of rotation of the shoulder. The hand was linked by means of a stiff brace to a handle at the end of a powerful robot that exerted computer-controlled forces during movement [Parallel-link Direct-Drive Air and Magnet Floating Manipulandum (PFM)] (Gomi and Kawato 1996). The forearm was supported against gravity by a beam in the horizontal plane at the level of the shoulder. Subjects performed reaching movements from a start circle to a target circle, both 2.5 cm in diameter. The center of the start circle was located 31 cm in front of the shoulder [i.e., at (0, 0.31) m relative to the shoulder], whereas the center of the target circle was 56 cm in front of the shoulder [i.e., at (0, 0.56) m]. The start circle, the target circle, and the instantaneous hand position, represented by a 0.5-cm-diameter cursor, were projected onto an opaque horizontal surface covering the arm. Positioning the hand in the start circle initiated a sequence of 3 beeps at 500-ms intervals. The subject was instructed to begin movement on the 3rd beep and complete it on a 4th beep, 600 ms later. Two additional beeps followed at 500-ms intervals to indicate the target hold time. Feedback of the performance indicating final hand position (OK, OUT) and movement duration (OK, LONG, SHORT) was given after each trial by displaying a message on a monitor in front of the subjects. Duration was considered OK if it was 600 ± 100 ms. The movement duration was the same as in Burdet et al. (2001a) for consistency because this duration had been used in previous measurements of null field (NF) stiffness. Position and force at the hand were sampled at 500 Hz.

In the first experiment, we investigated the adaptation to 2 force fields activated during movement. The VF produced a stable interaction. It was selected to elicit large modification of impedance (Franklin et al. 2003a,b). It was realized as

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \alpha \begin{bmatrix} 13 & -18 \\ 18 & 13 \end{bmatrix} \begin{bmatrix} \dot{x} \\ \dot{y} \end{bmatrix} \quad (1)$$

where (\dot{x}, \dot{y}) is the hand velocity (m/s) and $\frac{2}{3} \leq \alpha \leq 1$ depended on the subject's stiffness (Fig. 2B, left). 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 force (F_x, F_y) exerted on the hand by the robotic interface was computed as

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \begin{bmatrix} \beta x \\ 0 \end{bmatrix} \quad (2)$$

where the x -component of the hand position was measured relative to the y -axis (Fig. 2B, right). $\beta > 0$ (N/m) was chosen to be larger than the stiffness of the arm measured in NF movements so as to produce an unstable interaction. Specifically, β was 125 N/m larger than the K_{xx} component of endpoint stiffness measured in the NF. Although hand paths were essentially straight in the NF, they varied slightly from trial to trial because of motor output variability or environmental

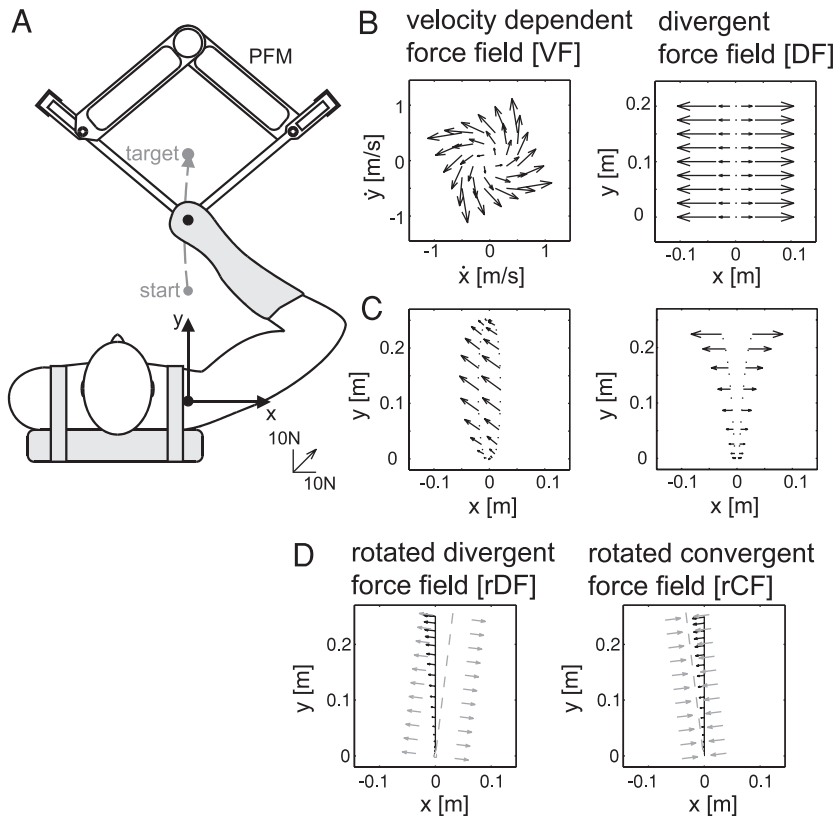


FIG. 2. Experimental setup. A: robotic interface exerts forces on hand during horizontal point-to-point arm movements away from body. B: force in velocity-dependent force field (VF) plotted as function of hand velocity (*left*) and force in divergent force field (DF) plotted as function of hand position (*right*). Vectors indicate direction and magnitude of forces. In VF force vector rotates and increases in magnitude with velocity. In DF force vector increases in magnitude with distance from y -axis. C: force vectors when hand trajectories (dotted curves) are slightly deviated from a straight line along y -axis to *left* and *right*. In VF (*left*) applied force vectors are similar despite trajectory variations. DF (*right*) amplifies trajectory variations. D: force in rotated divergent force field (rDF; *left*) and rotated convergent force field (rCF; *right*). Gray dashed lines and vectors indicate zero force line and direction and magnitude of force at origin of each vector. Black vector denotes applied force when hand trajectories are on straight line connecting start and end targets (black lines). Magnitude and direction of x -force along black lines are identical in rDF and rCF.

disturbance caused by the robot arm. The DF amplified such variations by pushing the hand with a force proportional to deviation from the y -axis. For safety reasons the DF was turned off when the trajectory deviated more than 3 cm from the y -axis. α and β were set for each subject so that the maximum magnitudes of the applied force in the VF and DF were similar. That is, assuming a minimum-jerk velocity profile in the y -direction with a duration of 600 ms, the maximum applied force in the DF would be 0.5, 1, and 1.5 times the maximum applied force in the VF if, in the DF, the deviation of the hand in the x -direction was 1, 2, and 3 cm, respectively. Therefore the perturbing force amplitude was roughly similar in both force fields within the safety zone of the DF.

Figure 2C shows perturbing force (arrows) applied to the hand when hand trajectories (dotted curves) deviate slightly from a straight line in the VF (*left*) and DF (*right*). In the VF, a slight difference in the trajectory to the left or to the right results in similar development of perturbing force both in magnitude and direction, whereas in the DF, the force develops in diametrically opposite directions. Although the magnitude of the force is similar in both fields, the variability in force direction resulting from small variations in the trajectory is quite different.

The second experiment was designed to investigate how impedance control and inverse dynamics model acquisition could be combined. The DF, as described above, requires that the same force be exerted as in the NF when moving along the y -axis (i.e., there is no perturbing force despite the instability). Everywhere else in the workspace there is a perturbing force as well as instability. The instability exists because overcompensating for the perturbing force would cause the hand to move from its equilibrium position toward the center of the DF with increasing acceleration, whereas undercompensating for the perturbing force would cause the hand to move away from the center of the DF with increasing acceleration. Everywhere except along the y -axis, adaptation to both the perturbing force and the instability is required. To compare the contribution of an inverse dynamics model and independently controlled impedance after adaptation to stable and

unstable dynamics we had subjects adapt to a stable force field (rCF) and an unstable force field (rDF), both of which incorporated the same varying bias force along the target trajectory. We then examined aftereffects of adaptation to the rDF and the rCF. The force exerted on the hand by the robotic interface was described as

$$\begin{bmatrix} F_x \\ F_y \end{bmatrix} = \beta [x \cos \gamma + (y - 0.31) \sin \gamma] \begin{bmatrix} \cos \gamma \\ \sin \gamma \end{bmatrix} \quad (3)$$

where $\gamma = -7^\circ$ for the rDF and $\gamma = 7^\circ$ for the rCF. For the rDF, β was 125 N/m larger than the K_{xx} component of endpoint stiffness measured in the NF. The x -force in the rDF was zero along the line connecting (0, 0.31) m and (0.03, 0.56) m (broken gray line in Fig. 2D, *left*) and increased perpendicular to this line, in proportion to distance, in the fashion of negative stiffness. Thus when moving along the target trajectory, (0, 0.31) m to (0, 0.56) m, the subject had to compensate for instability and apply a bias force, which increased in proportion to the distance moved toward the target. There was also a small assistive force in the y -direction, equal to about 12% of the x -force on the target trajectory. The rCF differed from the rDF only in the sign of γ (noted above) and β , which was negative, to create a stable interaction with the arm. As a result, the x -force was zero along the line connecting (0, 0.31) m and (-0.03, 0.56) m (broken gray line in Fig. 2D, *right*) and acted like positive stiffness perpendicular to this line. The x -force along the target trajectory was identical to that of the rDF (Fig. 2D, *right*), but the small y -force was resistive rather than assistive. The movement studied was the same as in the DF and the VF, except that the target movement duration was decreased to 500 ± 100 ms to reduce the possibility of voluntary correction.

Stability of the combined system

To verify that the initial interaction with the VF was stable we confirmed that the trajectories did not diverge when subjected to small force perturbations. A subject performed movements to the target with the NF present on 70% of the trials and the VF present on 30% of the

trials selected at random. During some of the VF trials, a brief triangular force pulse (25-ms duration) with an amplitude of 15 N in either the $+x$ or $-x$ direction was applied to the hand 100 ms after leaving the start circle. Trajectories in the VF, in the absence of the force pulse, were displaced in the direction of the force field, but terminated in the target circle. The force pulse produced an additional displacement, but the trajectory quickly returned to its original path, terminating in the target circle (Franklin et al. 2003). The result is consistent with the defining characteristic of Lyapunov stability that a small perturbation does not produce divergence.

Because the environmental stiffness was positive in the rCF, the combined endpoint stiffness of the arm and environment was also positive (i.e., stable). Instability in the DF and rDF was ensured by choosing the environmental stiffness so that the combined endpoint stiffness of the arm and environment would be negative (Burdet et al. 2001a). We measured endpoint stiffness during the movements in the NF by applying positional perturbations with the PFM (Burdet et al. 2000). The environmental stiffness was chosen for each subject so that the combined endpoint stiffness was considerably less than zero. In summary, before learning, the DF and rDF produced an unstable interaction with the arm, whereas the interaction was already stable in the VF and in the rCF. Consequently, we had a 2×2 design with respect to stability and load bias along the target trajectory (Table 1).

Aftereffects

Behaviorally, adaptations are characterized by “after-effect” trajectories when the imposed force is unexpectedly attenuated. Assuming that an inverse dynamics model compensates for imposed dynamics, after-effect trajectories with mirror image curvature compared with those produced by the imposed dynamics have been taken as evidence of an acquired inverse dynamics model. In reality, after-effect trajectories are determined both by the inverse dynamics model and the limb impedance (Takahashi et al. 2001). Theoretically, the amount by which the after-effect trajectories deviate from the target trajectory increases as the force produced by the imposed dynamics increases, whereas it decreases as the impedance increases. The amount by which after-effect trajectories deviate from the target trajectory depends on the magnitude of the subject’s impedance. If the imposed force in 2 environments is the same and aftereffects are compared, then greater impedance will result in a smaller deviation and less impedance in a greater deviation. Because the x -force along the target trajectory was identical in the rDF and rCF, the amount of deviation in the after-effect trajectories compared with the target trajectory indicated the relative contribution of impedance control and inverse dynamics model formation to adaptation.

From the theoretical perspective, the trajectory variance decreases as the impedance increases, whereas it increases as motor noise increases (Harris and Wolpert 1998). Unfortunately, motor noise cannot be measured directly, but it can be assumed to increase with the strength of the motor commands (i.e., as muscle activation) increases. By selectively increasing impedance, that is, achieving a balance between impedance and noise, trajectory variance can possibly be reduced. To determine whether the impedance controller functions to reduce trajectory variance, we compared trajectory variance in NF and DF aftereffects where mean trajectories were similar.

TABLE 1. *Experimental conditions*

Condition	Unstable	Stable
Bias force	rDF	VF, rCF
No bias force	DF	NF

DF, divergent force field; NF, null force field; rCF, rotated convergent force field; rDF, rotated divergent force field; VF, velocity-dependent force field.

Before effects

Because some adaptation was already evident on the second repeated trial in a force field, it was difficult to infer the effect of the force field before adaptation. To estimate the effect, we activated the force field on random trials during a series of movements in the NF (before effects), and examined trajectories generated in response to the force field.

Protocol for the experiments

Five naïve subjects participated in the first experiment with the VF and DF. They started with movements in the NF (the 1st day). Subjects were randomly placed into one of 2 groups. One group learned the VF first (the 2nd day), followed another day by the DF (the 3rd day). The second group started with the DF. Learning occurred over 100–300 trials. Another 100 trials, 20 of which had the force field removed unexpectedly, were recorded to check for aftereffects. On a subsequent day, subjects performed an additional 100 trials, of which 80 were in the NF, whereas 20 were in the force field, which was activated unexpectedly (before effects) to infer the effect of the force field before learning.

Five subjects also participated in the second experiment with the rDF and rCF. Only one of the 5 had participated in the first experiment. Subjects were randomly placed into one of 2 groups. One group first learned the rCF for 50 trials and then performed another 100 trials, 10 of which had the force field removed unexpectedly to record aftereffects. The same procedure was followed for the rDF later the same day. The second group started with the rDF, followed by the rCF. All movements were recorded during these sessions, including those not reaching the target.

Alternative strategies

Differences in behavior in the VF and the DF might simply be the result of differences in the mechanical constraints of the tasks. To demonstrate that the observed strategy was not the only strategy that could be used to successfully perform the task, we asked a subject to perform the task using an alternative strategy. We instructed an experienced subject to generate curved trajectories in the DF by moving through a 2-cm-diameter target placed at the midpoint of the movement path, 5 cm perpendicular to the straight line between the start and end targets, which would require inverse dynamics model learning. One hundred trials were performed for learning and of 100 additional trials, 20 were chosen randomly to examine aftereffects. The force field was inactivated on the after-effect trials. To test the possibility of using an alternative strategy in the VF, the subject was instructed to co-contract to resist the external force while reaching for the target. Fifty trials were performed for learning and of 100 additional trials 20 were after-effect trials, randomly chosen. As a demonstration that an inverse dynamics model would be ineffective in the DF even if it were possible to learn, we compared generalization in the 2 force fields. Two experienced subjects performed 50 trials in each force field, during which the target was occasionally repositioned 4 cm to the right or to the left of the original target before the start of the trial (5 trials for each direction). The DF safety zone was expanded to more than 7 cm in these experiments.

Hand path errors

Because the NF movements exhibited roughly straight trajectories, 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

$$S(|e_x|) = \int_{t_0}^{t_f} |x(t)| |y'(t)| dt \quad (4)$$

represents the area between the actual movement path and the straight line (Fig. 3A). The signed hand path error, defined as

$$S(e_s) = \int_{t_0}^{t_f} x(t)|y'(t)|dt \quad (5)$$

is a measure of the mean directional extent by which the path deviates from the straight line (Fig. 3B).

Hand path errors were calculated from the start time, t_0 (75 ms before crossing a hand velocity threshold of 0.05 m/s), to the termination time, t_f (when curvature exceeded 0.07 mm^{-1}) (Pollock and Ishimura 1996). Exponential curves were least-square fitted to the hand path error, after smoothing with a 10-trial moving average, to model the learning process as a function of trial number. Hand path error in rDF and rCF aftereffects was normalized by the total distance moved in the y -direction $y(t_f) - y(t_0)$ to account for any effect of the small y -force (assistive in the rDF vs. resistive in the rCF) on final position. In general, a repeated-measures ANOVA, with subject as a random variable, was performed using all the data from each subject.

EMG measurement

To examine the muscle activity before and after learning, surface EMG was recorded during before-effect and after-effect trials for 4 subjects after learning. These data were then compared with EMG recorded during NF movements and EMG recorded after adapting to the novel force fields. All EMGs were recorded on the same day with the same electrode placements. Activity was recorded from 6 muscles producing torque at the shoulder and elbow joints. The muscles included 2 monoarticular shoulder muscles: the pectoralis major and the posterior deltoid; 2 biarticular muscles: the biceps brachii and the long head of the triceps; and 2 monoarticular elbow muscles: the brachioradialis and the lateral head of the triceps. The EMG was recorded by using pairs of disposable silver–silver chloride surface electrodes in a bipolar configuration with a separation distance of about 2 cm. The skin was thoroughly cleaned with alcohol and prepared by rubbing in electrode paste. Excess paste was wiped from the skin before attaching the electrodes. The resistance of each electrode pair was tested to ensure that it was $<10 \text{ k}\Omega$. EMG signals were filtered at 25 Hz (high-pass) and 1 kHz (low-pass) and sampled at 2 kHz.

EMG during the DF after-effect trials was compared with EMG during NF movements and EMG after complete adaptation to the DF. The rectified EMG was integrated over the entire movement, from 100 ms before movement onset until 800 ms after movement onset. Twenty trials were used from each condition and the data for all subjects were used in an ANOVA with subjects as a random variable.

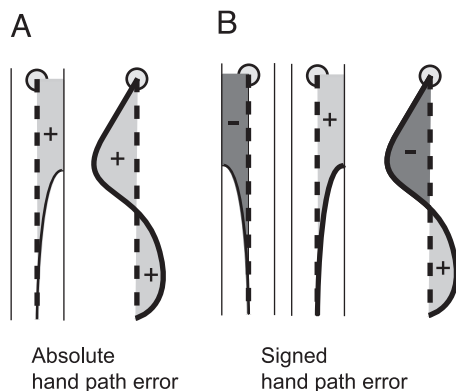


FIG. 3. Definition of absolute hand path error and signed hand path error. In DF area between y -axis and 3-cm safety zone was used after hand left safety zone. A: absolute hand path error is area between actual movement path and straight line. B: signed hand path error is net area with positive sign assigned to rightward displacement and negative sign to leftward displacement.

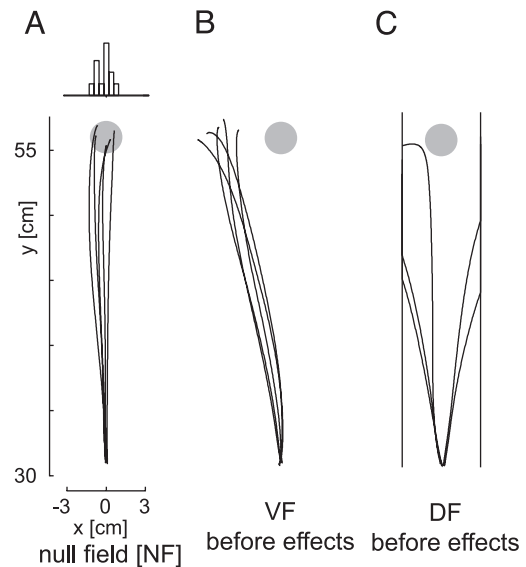


FIG. 4. Trajectories before learning force fields. Five trajectories of *subject 1* are shown for movements under different conditions. A: five consecutive trajectories for movements in null force field (NF). At top of trajectory plot, histogram shows distribution of endpoint positions (based on curvature $>0.03 \text{ mm}^{-1}$) for 5 movements being compared for each condition (all subjects). B: five trajectories during before effects in VF. C: five trajectories during before effects in DF. For safety reasons DF was turned off when trajectory deviated more than 3 cm from y -axis. Black lines indicate this safety zone.

A Scheffé post hoc comparison was then performed on the 3 conditions with a significance level of 0.05.

In the case of before effects, the 20 trials were sorted into 2 groups, based on whether the initial movement direction was to one side or the other of the mean movement direction. The averaged, rectified EMG was smoothed using a 125-point moving average and the equivalent NF EMG was subtracted, leaving only the change in EMG activity produced by reflex and voluntary responses. For comparisons across subjects and conditions this change in EMG was scaled so that the maximum value (positive or negative) was equal to one. The similarity of the pattern of EMG from these 2 groups of trials was examined by plotting one against the other every 10 ms for 500 ms from movement onset. A linear regression was then performed using the data from all subjects. If the EMG was similar in the 2 groups of trials then the result of the regression should have a high R^2 value with a slope close to one. These before-effect trials were also compared with the EMG of the final adaptation in both the VF and the DF conditions (mean of 20 trials) using the same method. In this case, the EMG of the final adaptation was plotted separately against trials to one side or the other of the mean trajectory in the case of the VF or the y -axis in the case of the DF. Again a linear regression was performed to examine the relation.

RESULTS

Trajectories before learning

To estimate the effect of the force field before adaptation, we examined before-effect trajectories generated in response to the force field activated on random trials during a series of movements in the NF (see METHODS). The behavior in the DF, which generated an unstable interaction, was distinctly different from that in the VF, which produced a stable interaction. Figure 4 shows trajectories of *subject 1* in the NF and when the VF and DF were activated unexpectedly. The movements performed in the NF had trajectories that were approximately straight. Almost all trials ended on target. The distribution of endpoints for

the 5 NF movements of all subjects, shown at the *top* of the *trajectory plot*, was bell-shaped. The movements performed in the VF were biased to the left. The movements performed in the DF diverged widely to either side as the DF amplified the initial deviation. Most trials crossed either the left or the right safety limit. The magnitude of the signed hand path error was significantly larger in the DF than in the VF (Siegel and Tukey method; $P < 0.01$), suggesting larger deviation in response to the DF than to the VF. The unstable interaction created by the DF caused trajectories to diverge, whereas the stable interaction with the VF produced biased, but not divergent trajectories.

Trajectories during and after learning

Although the initial movements were disturbed by the force field, subjects gradually adapted to the disturbances. Figure 5 shows the initial trajectories of *subject 1* and the trajectories

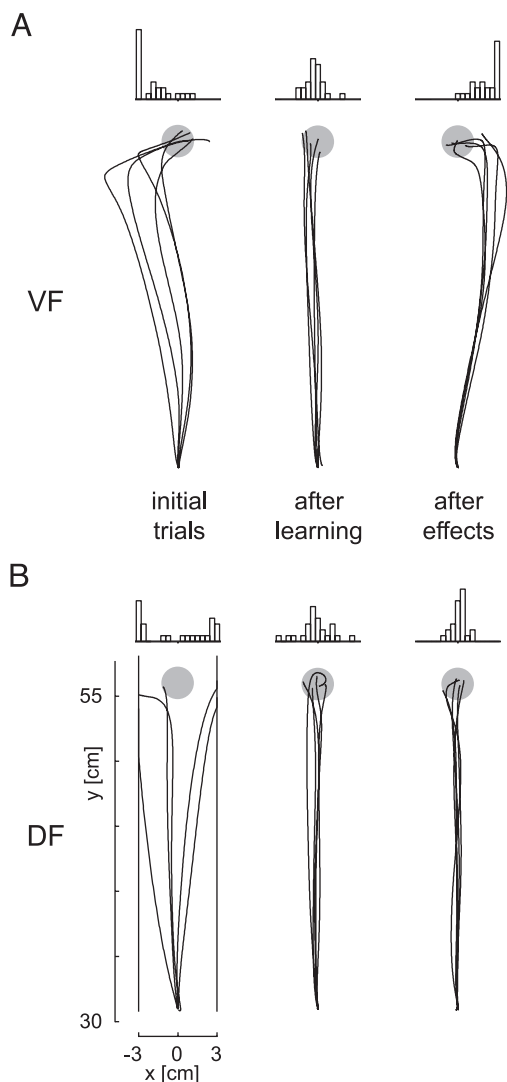


FIG. 5. Trajectories and endpoint distribution during and after learning in force fields. At *top* of each trajectory plot, histogram shows distribution of endpoint positions (based on curvature $>0.03 \text{ mm}^{-1}$) for 5 movements being compared for each condition (all subjects). A: learning in VF. Five trajectories of *subject 1* are shown, which include first 5 trials (*left*), last 5 trials (*middle*), and 5 consecutive after-effect trials (*right*). B: adaptation in DF. Format is identical to that in A.

after learning in the VF and DF. At the *top* of each *trajectory plot*, the distribution of endpoint positions for the 5 movements of all subjects is displayed. Initially, movements in the VF were systematically perturbed to the left as shown by the asymmetrical endpoint distribution. With practice the movements became straighter, similar to movements in the NF. The initial trials in the DF exhibited unstable behavior, generally diverging either to the right or to the left depending on the direction of the initial deviation. With practice, however, subjects gradually became proficient at producing straight trajectories along the y-axis, accompanied by a bell-shaped endpoint distribution, similar to that in the NF. A repeated-measures ANOVA was used to compare the last 5 trials in the NF to the last 5 trials of the learning phase in the VF and DF. At the end of the learning trials, no significant difference in the absolute error was found between movements in the NF, VF, and DF ($P = 0.906$), indicating that the subjects performed similar movements in all 3 conditions after adaptation.

To determine how compensation for the dynamics was achieved, we recorded aftereffects (i.e., trajectory deviations that occurred when the force field was unexpectedly removed on selected trials after learning). A repeated-measures ANOVA was used to compare the position error of the last 20 movements in the NF to that of 20 after-effects trials in either the VF, using signed error, or the DF, using absolute error. The after-effect movements in the VF were systematically displaced to the right relative to movements in the NF to compensate for the dynamics (signed hand path error of Eq. 5, $P = 0.001$; Fig. 5A, *right column*). In contrast, the after-effect movements in the DF were characterized by trajectories that deviated very little from the y-axis (Fig. 5B, *right column*). In fact, the after-effect trajectories were even closer to the y-axis than NF trajectories (absolute hand path error of Eq. 4, $P = 0.001$). In the VF after-effect trials, the subjects noticed that the force field had been removed soon after the movement start. In the DF, they were not aware that the force field had been removed, even after the movement had terminated. Thus learning in the DF exhibited significant aftereffects, which were fundamentally different from aftereffects following learning in the VF; that is, in the VF, subjects learned to produce the force necessary to predictively compensate for the force field, probably by using an inverse dynamics model. On the other hand, in the DF, subjects produced essentially the same net joint torques with less variance, and hence the same movement trajectories, as during NF movements.

Although endpoint forces, and therefore net joint torques, were not significantly different between the DF and NF at the midpoint of the movements as shown in Burdet et al. (2001a), the pattern of muscle activation was indeed very different (Franklin et al. 2003). The EMG was much higher during DF trials compared with that during NF trials ($P < 0.001$ for all 6 muscles), corresponding to the adapted impedance. Similarly, the after-effect EMG was much higher than that in NF trials ($P < 0.001$ for all 6 muscles), although the force field was identical.

Although the DF was off during the after-effect trials, the subjects were not aware for which trials the DF was on or for which it was off, and assumed that it was always on. If the impedance had been reactive to the activation of force field, EMG in aftereffects would have been similar to EMG in NF after learning because the force field was attenuated. However, the EMG profile during after-effect trials was more like that of

DF trials than of NF trials (Franklin et al. 2003). This reveals a preprogrammed motor command to compensate for the DF. The impedance was predictively and not reactively controlled.

Evolution of absolute hand path error

Because the trajectories after learning were roughly straight, similar to the NF movements, absolute hand path error relative to the straight line from the start to the target was used as an index of adaptation (see METHODS). Figure 6 shows the absolute error of *subject 1* during the different stages of learning in the

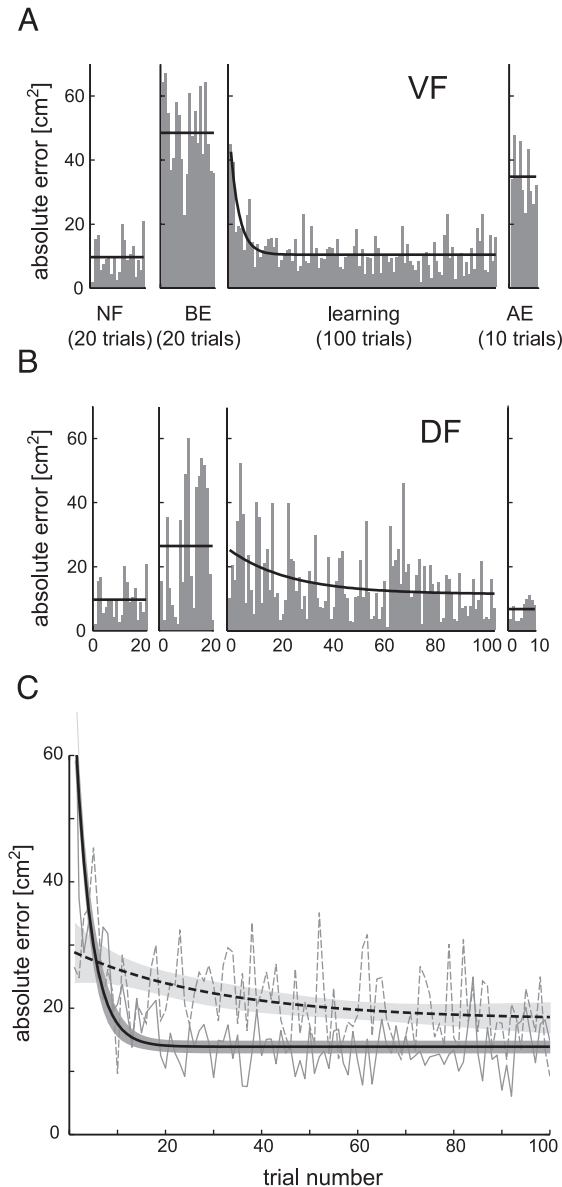


FIG. 6. Evolution of absolute hand path error during learning in force fields, for *subject 1*. *A*: absolute hand path error in VF is shown as function of trial number, which includes 20 trials in NF (*extreme left*), 20 before-effect trials (BE; *second from left*), 100 trials during learning (learning; *second from right*), and 10 after-effect trials (AE; *extreme right*). Superimposed lines for NF, BE, and AE indicate mean of data, whereas superimposed curve for learning indicates exponential fit to data. *B*: absolute hand path error in DF. Format is identical to that in *A*. *C*: evolution of absolute hand path error averaged across all subjects as function of trial number during learning in VF (solid) and DF (dashed). Exponential fit (solid curve for VF and dashed curve for DF) and its 95% confidence intervals (dark gray area for VF and light gray area for DF) are superimposed.

TABLE 2. Comparison of decay rate of the absolute error during adaptation to the VF (r_{VF}) and DF (r_{DF})

Subject	r_{VF} (CI)	r_{DF} (CI)	r_{DF}/r_{VF}
S1	0.1286 (0.0091)	0.0068 (0.0253)*	0.0529
S2	0.1613 (0.0080)	0.0045 (0.0103)*	0.0281
S3	0.0923 (0.0057)	0.0789 (0.0138)	0.8557
S4	0.1273 (0.0136)	0.0240 (0.0085)*	0.1889
S5	0.2038 (0.0217)	0.0010 (0.0161)*	0.0050
Mean	0.1349 (0.0061)	0.0238 (0.0115)*	0.1767

Values are exponential fit for each subject. CI, 95% confidence interval; * cases in which 95% confidence intervals of VF and DF do not overlap.

VF and DF compared with the NF. As already indicated in Figs. 4 and 5, absolute hand path error was large in before effects, and decreased as learning progressed in both fields. The after-effect trials of the VF exhibited larger absolute errors than those of NF trials. The absolute error during after-effect trials in the DF was significantly smaller than that in NF trials, suggesting that the after-effect trajectories were even closer to the straight line than NF trajectories. To test whether learning occurred, we compared absolute error in the first 5 trials and the last 5 trials using a repeated-measures ANOVA with random factor subjects. The absolute error decreased significantly between the first 5 trials and the last 5 trials of the learning, both in the VF ($P = 0.004$) and in the DF ($P = 0.002$), indicating that the subjects learned to adapt to the dynamics of both fields.

We also compared the speed of learning in the force fields, although caution must be exercised in interpretation of the results because of the existence of a safety zone in the case of the DF but not the VF. The absolute error in the DF decayed more slowly than that in the VF. Table 2 lists the decay rate r of the exponential fit for each subject and the mean for all subjects, along with the 95% confidence interval. The R^2 value of the fit, averaged across subjects, was 0.95 for the VF and 0.37 for the DF. The exponential decay rate for 4 of 5 subjects, as well as for the mean of the 5 subjects, was significantly slower in DF than in VF. Because decay rates varied considerably across subjects we computed the ratio of DF to VF decay rate for each subject (*right column* of Table 2). The ratio was significantly <1 (t -test, $P < 0.005$), suggesting that learning in the DF required more trials than in the VF. In an attempt to take into consideration the effect of the safety zone, we computed the decay rate using the absolute hand path error up to the time that the hand crossed the safety zone and performed the same statistical analysis. This analysis produced a similar result, with significantly slower learning in the DF than in the VF. We also computed the VF decay rate using the absolute hand path error, limited to the region of the safety zone of the DF (3 cm), and performed the same statistical analysis. This analysis again produced a similar result, with significantly slower learning in the DF than in the VF.

Evolution of signed hand path error

Whereas the absolute error decreased monotonically during learning in both the VF and DF, the evolution of signed error (see METHODS) was markedly different in the 2 fields. Figure 7 shows the signed error of *subject 1* during the different stages of learning in the VF and DF compared with the NF. During

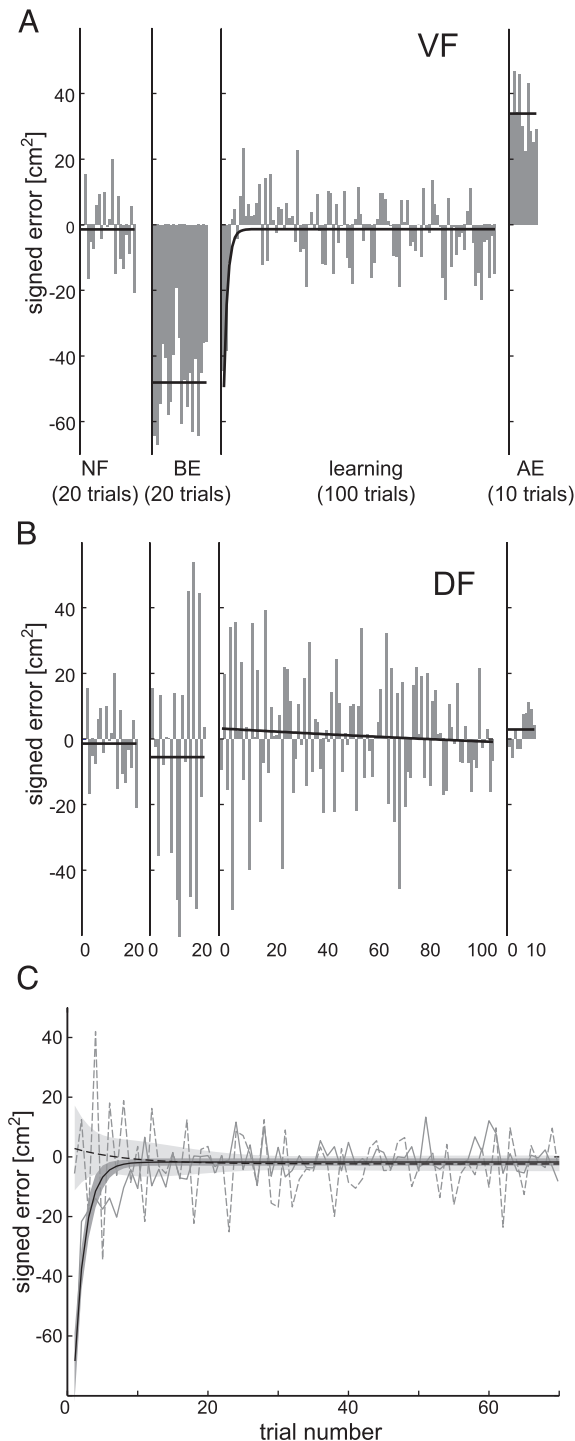


FIG. 7. Evolution of signed hand path error during learning in VF and DF, for subject 1 in same format as in Fig. 6. A: signed hand path error in VF. B: signed hand path error in DF. C: signed hand path error averaged across all subjects.

before-effect trials and during the initial stage of learning in the VF, the signed error was consistently biased to the left. On the other hand, during DF before-effect trials, the signed error tended to alternate to the left and right. In the VF the signed error exponentially approached the mean signed error in the NF. In contrast, the mean signed error in the DF never varied much from that in the NF throughout the entire learning period. The 2 learning curves were clearly different at the onset of

learning, both in mean value and deviation. In the VF the error decreased monotonically and approached zero; that is, it approached the mean of NF movements. In the DF the mean was closer to the NF mean, but the deviation was much larger. The mean changed very little but the deviation gradually decreased. At the end of learning, the means were close to the NF mean for both fields, although the SD was slightly larger in the DF. The mean signed error during the after-effect trials in the VF was biased to the right (positive), whereas in the DF it was similar to the NF error but with less variability (Figs. 4 and 5).

To determine whether the signed error changed during learning in the VF and DF, the first and last 5 trials were compared using a Wilcoxon signed-ranks test. In the VF the magnitude of the signed error decreased during learning ($P = 0.002$), indicating that the CNS could use this error information for learning of the dynamics. In contrast, the signed error in the DF was not modified between the first and last 5 movements ($P = 0.600$), whereas SD significantly decreased (t -test, $P < .005$). This indicates that in the DF, even if the magnitude of the signed error of each trial is large, the mean over successive trials does not provide information that could be used for the conventional feedback error learning (Burdet et al. 2001b).

As seen in Fig. 7, the evolution of signed error was more apparent during the initial stage of learning. To examine the initial learning in the DF, trial by trial, we plotted the mean signed hand path error of the 5 subjects and its SD for the first 20 trials (Fig. 8). The sign convention of the error for each subject was chosen to make the error on the first trial negative. In the DF the error was alternately negative and positive; that is, the trajectory alternated to the left and right of the straight line joining the center of the start and target circles. This alternating movement pattern disappeared after about the 10th trial.

The alternating behavior can be explained if the CNS attempts to acquire an inverse dynamics model of the DF by means of feedback error learning (Kawato et al. 1987). Suppose that the trajectory deviated to the left on the first trial. The CNS would detect that a perturbing force had pushed the hand to the left. On the following trial it would generate motor commands to compensate for that perturbing force, resulting in a trajectory biased to the right. The DF would exaggerate this bias and push the hand further to the right. The CNS would then suppose that the force field was actually pushing in the opposite direction and would try a leftward force on the next trial. Such alteration in force direction would be expected when the CNS tries to improve performance by acquiring an inverse dynamics model on a short time scale, which estimates the compensating force based on error information about displace-

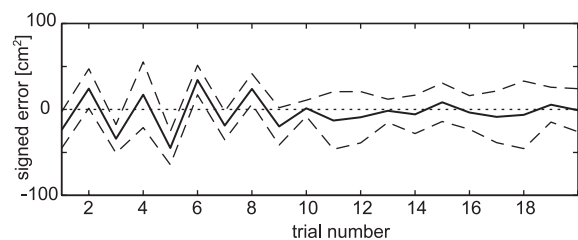


FIG. 8. Evolution of signed hand path error during initial trials in DF. Solid and dashed lines correspond to mean absolute hand path error of 5 subjects with SD. Sign of error for each subject was adjusted so that it would be negative on first trial.

ment direction on the previous trial. After several trials, it would become apparent that this strategy did not work, and the CNS would switch its strategy to increase impedance. This would be consistent with the observed attenuation of the oscillation in signed error.

Differences in muscle activation patterns in the VF and DF

The activation patterns of biarticular muscles in before effects and after learning illustrate differences in adaptation to the VF and DF. EMG patterns in the VF during before effects were very similar to those after learning and were consistent from trial to trial, whereas those in the DF were not. Figure 9 shows the EMG of the biarticular muscles of *subject 1* during before-effect trials and after learning, compared with NF trials. In the before effects, subjects assumed that the force field was off (i.e., they issued feedforward motor commands for the expected NF). Because the interaction with the VF is stable, before-effect trials, where the VF was activated in only 20% of the trials, have very similar trajectories. In contrast, before effects in the DF have trajectories deviating either to the right or to the left as a result of the unstable interaction. The difference between the adaptation to the VF versus DF dynamics can be best illustrated by separately analyzing movements according to their relative position with respect to the mean trajectory (Fig. 9, A and D). Note that in some cases there are different numbers of trajectories to the left and right of the mean.

In the VF the before-effect EMG was almost identical for movements to the left and right of the mean (Fig. 9B). Furthermore, the EMG after learning in the VF, which represents the acquired feedforward motor commands (EMG after learning minus NF EMG) was similar to the before-effect EMG, which represents reflex responses and voluntary error correction (before-effect EMG minus NF EMG) (Fig. 9C). The ratio of muscle activity in the final adaptation trials compared with that in the before-effect trials was 0.71 ± 0.03 (95% confidence interval on mean) in the posterior deltoid and 0.80 ± 0.08 in the long head of the triceps. In contrast, in the DF the before-effect EMG was distinctly different, depending on the direction in which the trajectory deviated from the mean (Fig. 9E). The EMG after learning in the DF was also very different from the before-effect EMG (Fig. 9F). Reflex responses contributed substantially to the before-effect EMG. If we assume that reflex responses are indicative of the sensory feedback information received by the CNS then the information during the initial stage of learning was consistent from trial to trial in the VF, but inconsistent in the DF. Further, the final feedforward motor commands were similar to initial feedback responses in the VF, but different in the DF.

We can examine these results in greater detail. The long head of the triceps serves as the best example to illustrate differences in the adaptation process because large changes in its activity were observed after adaptation to the VF and DF. The pattern of reflex and voluntary changes in EMG during the before-effect movements to the left and to the right of the mean trajectory in the VF for the long head of triceps were similar ($R^2 = 0.78$) and the resulting final adaptation in the EMG was also highly correlated with both sets of before-effect EMG ($R^2 = 0.71$ and $R^2 = 0.72$). In contrast, the biceps brachii shows little correlation between the 2 sets of before-effect

EMG ($R^2 = 0.22$) or between the final adaptation EMG and the before-effect EMG ($R^2 = 0.19$, $R^2 = 0.05$), in that no reflex activation was observed in this muscle during the before-effect trials. Similar results were observed in the shoulder muscle pair (i.e., high correlations in the posterior deltoid), but low correlations in the pectoralis major. Overall, high correlation between before-effect EMG and final EMG was observed only in the muscles that contributed to compensate for VF dynamics. The correlation was low in the muscles where little change was observed after the VF adaptation. The trial-to-trial variability in the EMG was much larger in the DF than in the VF. This was consistent across subjects, with no relation between the EMG during trials deviating left and right in either the long head of the triceps ($R^2 = 0.01$) or the biceps brachii ($R^2 = 0.02$). Clearly, the reflex response in these 2 antagonistic muscles varied from trial to trial depending on the direction of deviation produced by this force field. The EMG after learning in the DF was also very different from the before-effect EMG. The EMG after adaptation was not correlated with the EMG in before effects for the long head of the triceps ($R^2 = 0.09$, $R^2 = 0.05$) or for the biceps brachii [$R^2 = 0.43$ (slope $\neq 1$), $R^2 = 0.08$]. Similar results were observed in the shoulder muscle pair (i.e., low correlations in both posterior deltoid and pectoralis major).

Thus muscle activation patterns indicate 2 principal differences between adaptation to the VF and DF. First, the stability of the VF resulted in before-effect EMG that was consistent from trial to trial, whereas the instability of the DF resulted in EMG that was quite variable from trial to trial. Second, the EMG generated by feedback pathways to correct for perturbations in the VF during before effects was very similar to the EMG generated by feedforward pathways to compensate for the VF after learning. In contrast, EMG generated by feedback pathways in the DF during before effects was very different from the selective co-contraction generated by feedforward pathways to compensate for the instability of the DF after learning.

Aftereffects in rotated force fields

The DF adaptation described above occurred in the region where force was close to zero. Therefore there was little need to compensate for a perturbing force. This represented the extreme where only impedance could be learned. Everywhere else in the workspace, adaptation to both a perturbing force and instability is required (see METHODS). To compare the contribution of an inverse dynamics model and impedance in the adaptation to stable and unstable environments when the perturbing force is nonzero, the hand path errors of aftereffects in the rDF were compared with those in the rCF. Because the rCF was stable, it was assumed that only an inverse dynamics model would be acquired in the rCF after learning. DeSerres and Milner (1991) previously showed that the relation between joint stiffness and joint torque is essentially identical for constant torque loads and elastic loads, like the rCF. This implies that in stable environments the impedance is determined only by the muscle activation necessary to produce the force required to match the perturbing force, which is the basis for our assumption that in the rCF there would be no impedance over and above that associated with an inverse dynamics model. Subjects were able to adapt to both the rDF and rCF. A

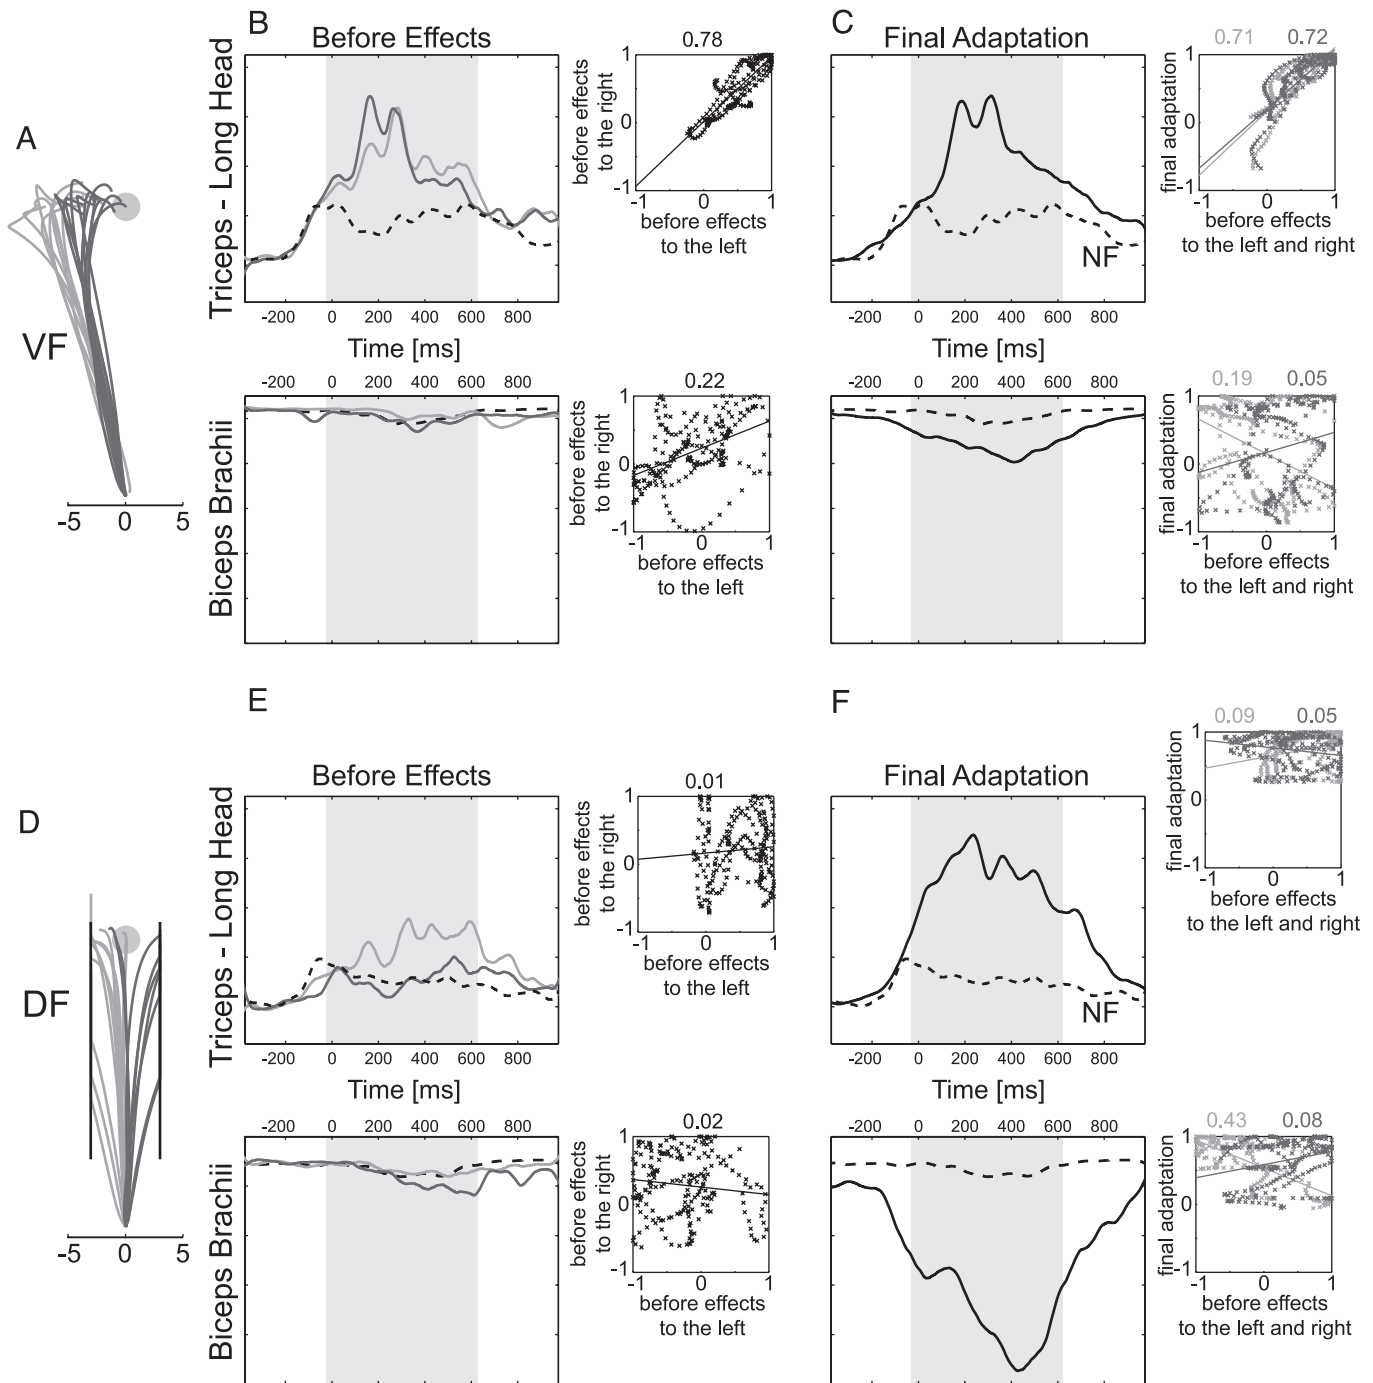


FIG. 9. Rectified, averaged, and smoothed surface EMG of long head of triceps (*top panels*) and biceps brachii (*bottom panels*) muscles during before-effect trials and after learning for *subject 1*. *A* and *D*: before-effect trajectories were divided into 2 groups (dark and light gray lines), according to direction in which trajectory diverged with respect to mean in VF (*A*) and with respect to *y*-axis in DF (*D*). *B*: EMG during before-effect trials in VF. Dark and light gray lines correspond to mean EMG for dark and light gray trajectories in *A*, respectively, whereas dashed line corresponds to intervening NF trials. To *right* of each time course plot is a scatter plot where EMG during movements to *right* is plotted against EMG during movements to *left* for all subjects. If change in EMG is similar during both of these movements, scattered data will exhibit a linear relationship with a slope close to one. Best-fit line of each relation is shown and corresponding R^2 values appear above plot. *C*: EMG after learning (Final Adaptation) in VF (solid line) compared with NF (dashed line). Again, to *right* of each time course plot, a scatter plot illustrates relation between final EMG adaptation and before-effect EMG in movements either to *left* (in light gray) or to *right* (in dark gray) with best-fit lines and R^2 values. *E*: EMG during before-effect trials in DF. Dark and light gray lines correspond to mean EMG for dark and light gray trajectories in *D*, respectively, whereas dashed line corresponds to intervening NF trials. Scatter plots are similar to those in *B*. *F*: EMG after learning (Final Adaptation) in DF (solid line) compared with NF (dashed line). Scatter plots are shown as in *C*.

repeated-measures ANOVA applied to the signed error of the final 10 trials in the rDF, rCF, and NF showed no significant effect of the type of force field ($P > 0.05$). Figure 10 shows rDF and rCF after-effect trajectories of *subject 1*. As expected, the after-effect trajectories were displaced opposite to the perturbing force in both the rDF and rCF. The signed error of the after-effect trajectories was significantly positive for both rDF ($P < 0.001$) and rCF ($P < 0.001$), suggesting that an inverse dynamics model was acquired. However, the signed error was significantly smaller for rDF aftereffects than rCF aftereffects ($P < 0.001$), indicating higher impedance in the rDF than in the rCF, given that the perturbing force to which subjects had adapted was essentially the same. The ratio of signed error in the rCF to signed error in the rDF for the 5 subjects was 2.07 (SD 0.39), suggesting that impedance was, in fact, much higher in the rDF than in the rCF. This suggests that in the rDF, an inverse dynamics model was acquired to compensate for the perturbing force and, at the same time, impedance was significantly increased to compensate for the instability.

Curved trajectories in the DF, co-contraction in the VF

Although we did not give explicit instructions, subjects tended to generate straight trajectories in the DF, which did not require acquisition of an inverse dynamics model of the force field. However, this is not the only possible strategy in the DF. Another possibility is that subjects could intentionally or unconsciously produce biased trajectories to the left or right, so as to evoke a consistent (i.e., predictable) perturbing force, so that impedance derived from the compensating muscle force could provide stabilization. Figure 11A shows the mean and SD of trajectories in the DF (solid) when explicitly instructed to move along a curved path, and the after-effect trajectories (dashed). After-effect trajectories were shifted opposite to the direction of the imposed force relative to the curved trajectories in the DF, suggesting an acquired inverse dynamics model. In the VF, subjects produced the necessary force (i.e., by reciprocal

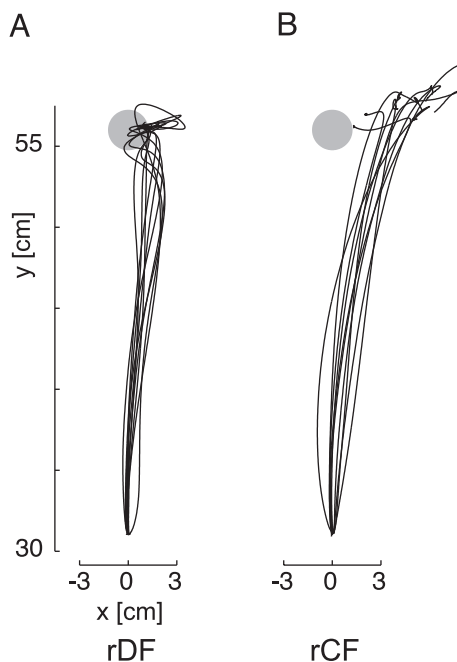


FIG. 10. After-effect trajectories in rDF (A) and in rCF (B) for *subject 1*.

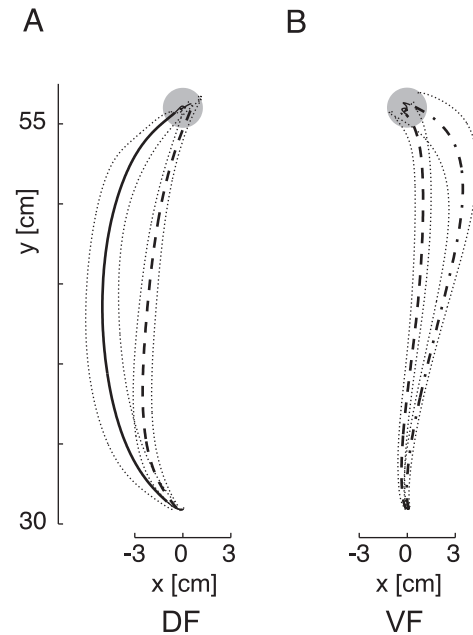


FIG. 11. Movements with alternative strategies for one subject. A: mean trajectory in DF when explicitly instructed to follow curved path (solid), and its mean after-effect trajectory (dashed). Dotted curves are SD of corresponding mean trajectories. B: mean after-effect trajectory in VF when explicitly instructed to co-contrast during adaptation (dashed), and mean after-effect trajectory in VF when no instruction was given (dashed-dotted). Dotted curves are SD of corresponding mean trajectories.

activation) to compensate for the perturbing force. However, subjects can also be instructed to compensate for the perturbing force by increasing impedance. Figure 11B shows the mean and SD of after-effect trajectories in the VF when subjects were explicitly asked to co-contrast (dashed), in comparison to those when given no explicit instructions (dashed-dotted). Smaller aftereffects when subjects were instructed to co-contrast in comparison to those without instruction suggest that increased impedance with little evidence of an inverse dynamics model is effective in the VF. These results demonstrate that subjects in the first experiment strongly preferred one strategy over another even though both strategies were possible.

Generalization tests of inverse dynamics models

Previous studies have demonstrated significant capability of inverse dynamics models for local generalization (Gandolfo et al. 1996; Goodbody and Wolpert 1998; Imamizu et al. 1995; Sainburg et al. 1999; Shadmehr and Mussa-Ivaldi 1994). The ability to generalize decayed smoothly as the movement direction deviated farther from that of training (Gandolfo et al. 1996; Sainburg et al. 1999). Therefore if an inverse dynamics model was acquired, we should expect similar local generalization of learning in our force fields. In particular, we would expect subjects to be able to move with similar accuracy to nearby targets as to the learned target. In the VF both subjects easily mastered the new targets (Fig. 12A), suggesting local generalization by means of an acquired inverse dynamics model. However, in the DF both subjects missed the new targets (Fig. 12B), revealing poor directional generalization. Note that in the DF, the forces in the vicinity of the new targets were no larger than those experienced during the curved movements of Fig. 11A. If an inverse dynamics model had been

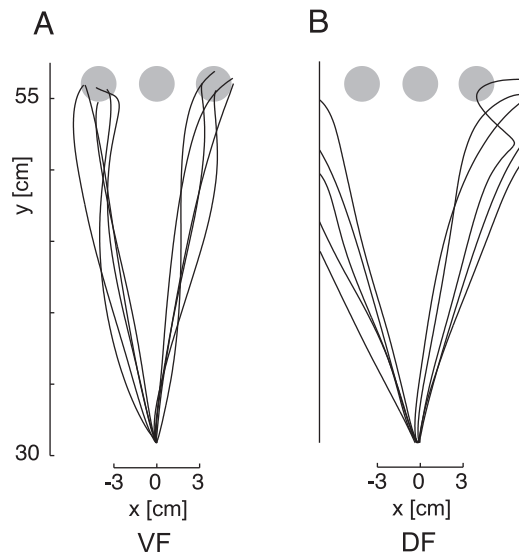


FIG. 12. Test of generalization for one subject. A: trajectories to displaced targets (*left* and *right*) in VF. B: trajectories to displaced targets (*left* and *right*) in DF.

learned in the DF, local generalization would have been expected and the subjects should have been able to move accurately to the new targets. The results support the conclusion that an inverse dynamics model is learned in the VF, but is not learned in the DF.

DISCUSSION

Our results demonstrate qualitative differences between adaptation to stable and unstable dynamics, which are summarized in Table 3. These differences may be indicative of different mechanisms for the acquisition of an inverse dynamics model and the control of mechanical impedance. We also demonstrated that the 2 kinds of adaptation are combined when compensating for a perturbing force in an unstable environment.

Inverse dynamics model versus impedance control

Our results suggest that the CNS uses distinctly different control mechanisms to adapt to stable and unstable dynamics. Because the subjects avoided perturbing forces in the DF, it does not appear that an inverse dynamics model was learned in compensating for DF dynamics. The lack of local generalization supports this view. The inverse dynamics model provides no stability when no perturbing force occurs along desired

trajectory (i.e., it produces no change in mechanical impedance). However, we have evidence that mechanical impedance *was* controlled independently of force in a predictive manner. The lack of aftereffects in the DF, even though error decreased during the adaptation, is proof of this. The persistence of increased muscle coactivation in DF after-effect trials, when the force field was inactive, also supports this (Franklin et al. 2003).

On the other hand, all evidence points to the formation of an inverse dynamics model in the VF, that is, the curvature of VF after-effect trajectories, local generalization, and the reciprocal patterns of muscle activation. The demonstration by counterexamples that curved trajectories (inverse dynamics model) in the DF, and co-contraction (impedance control) in the VF are viable strategies proves that the observed behavioral difference is not the result of the physical constraints imposed by the force fields. However, subjects strongly preferred impedance control in the DF as opposed to an inverse dynamics model in the VF.

Inverse dynamics model formation is realized through feedback error learning, whereas impedance learning is not

When the interaction between the arm and the novel dynamics is stable, the signed error between the planned and realized trajectories is systematically and reliably correlated with the effect of the imposed dynamics. Furthermore, the signed error is similar from trial to trial. The CNS can use this information as a teacher error signal to acquire an inverse dynamics model of the novel dynamics in the manner of supervised learning. As this inverse dynamics model becomes more accurate at predicting the dynamics and compensating for them in subsequent movements, the signed error will gradually decrease, as was indeed observed in the VF. In this “feedback error learning” strategy (Kawato 1990; Kawato et al. 1987), feedback information indicates the direction in which the feedforward motor commands should be modified. In the VF we can infer that the CNS receives consistent feedback about how to modify and shape the feedforward motor commands. The similarity between the feedforward EMG after adaptation and the before-effect EMG for the muscles used in compensating for the VF, which constituted both reflex feedback commands and voluntary correction, suggests that these feedback motor commands might be representative of the error signals used during learning of the feedforward command (i.e., the signals used to update the inverse dynamics model).

In contrast, feedback was inconsistent in the DF because the hand was pushed to the right or to the left, depending on the initial deviation. As a result, both during and after learning, the

TABLE 3. Comparison of the adaptation to stable and unstable dynamics

Force field	VF and rCF	DF and rDF
Stability		
Before learning	Stable	Unstable
After learning	Stable	Stable
Aftereffect	Large bias	Deviation attenuated
Learning		
Absolute error	Fast decrease	Slow decrease
Signed error	Exponential decrease (initially biased)	No change (initially alternating)
EMG		
Feedback	Consistent	Inconsistent
Feedforward after learning	Correlated with feedback in the first trial	Uncorrelated with feedback in the first trial

signed error was randomly positive or negative with a mean close to that in the NF. It is theoretically possible to learn an inverse dynamics model of the DF that maps actual trajectories to force by using the signed error. However, such an inverse dynamics model can be used only for straight line movement in the DF when on-line control is possible using the current state as the input to the inverse dynamics model or when there is no delay or noise, neither of which is the case in the biological system. Instead, the CNS learns to selectively increase impedance along the preferred straight trajectory to stabilize an unstable interaction, for which conventional feedback error learning predicts no improvement of performance (Burdet et al. 2001b). If learning occurs with a fast decay rate (i.e., based primarily on the previous trial) (Scheidt et al. 2000), the signed error will oscillate and learning will not converge, similar to initial trials in the DF. If learning occurs with a slow decay rate (i.e., over many trials), the zero mean of the signed error on the longer time scale will result in no change in motor commands. Therefore the conventional feedback error learning algorithm based on signed error could not have been used by the CNS to guide learning of the unstable dynamics. The signed error contained information about the dynamics experienced on the current trial, which would have produced an incorrect prediction of the dynamics on the next trial.

The alternating behavior observed during the first few trials in the DF suggests that the CNS at first attempted to acquire an inverse dynamics model of the DF by means of feedback error learning, which failed (Fig. 8). The before-effect EMG shown in Fig. 9 indicates that reflex responses attributed to a leftward deviation in the DF are markedly different from those attributed to a rightward deviation. If we assume that the reflex responses are representative of the error information received by the CNS then there would be a large variability in feedback information from trial to trial. Furthermore, the feedforward command after adaptation was quite different from the initial reflex activity for either leftward or rightward trajectory deviation. Therefore it does not appear that compensation for the DF was achieved by any form of supervised learning of the environmental dynamics analogous to conventional feedback error learning (Kawato 1990; Kawato et al. 1987). The precise mechanism by which the CNS learns impedance to compensate

for unstable dynamics remains to be elucidated. Alternative models may be assessed in future studies (Bhushan and Shadmehr 1999; Gribble and Ostry 2000; Mussa-Ivaldi and Bizzi 2000; Wang et al. 2001).

Interaction between inverse dynamics models and impedance

The CNS preferred to use an inverse dynamics model in the VF and impedance control in the DF. That is, under mechanically stable conditions, a predictable perturbing force is counteracted by an inverse dynamics model. On the other hand, impedance control is used to counteract instability when no perturbing force is present. If both a perturbing force and instability occur, and compensation for the perturbing force cannot provide sufficient impedance to stabilize the movement, the CNS might employ both an inverse dynamics model and impedance control. Comparison of aftereffects in the rDF and rCF suggests that an inverse dynamics model was formed in both cases. However, the attenuation of aftereffects in the rDF compared with the rCF demonstrates that additional impedance was generated to counteract its instability. Thus the 2 mechanisms are not mutually exclusive, but can function in parallel according to the environment or to different learning phases (Franklin et al. 2003; Osu et al. 2002; Takahashi et al. 2001).

Implication to neural mechanisms and integrated control schema

Although we identified 2 distinct control strategies, it does not necessarily mean that the neural processes underlying these 2 strategies are also distinct. The results of electrophysiological studies suggest that the neural substrate for inverse dynamics models of external dynamics resides in the cerebellum (Kobayashi et al. 1998; Shidara et al. 1993). Imaging studies report cerebellar activation related to acquisition of internal models (Imamizu et al. 2000; Shadmehr and Holcomb 1997). Recent studies also suggest that primary motor cortex is involved in compensating for external dynamics (Cabel et al. 2001; Gandolfo et al. 2000; Li et al. 2001). In contrast, few studies have focused on the neural substrate of impedance control. Humphrey and Reed (1983) reported separate neuronal systems

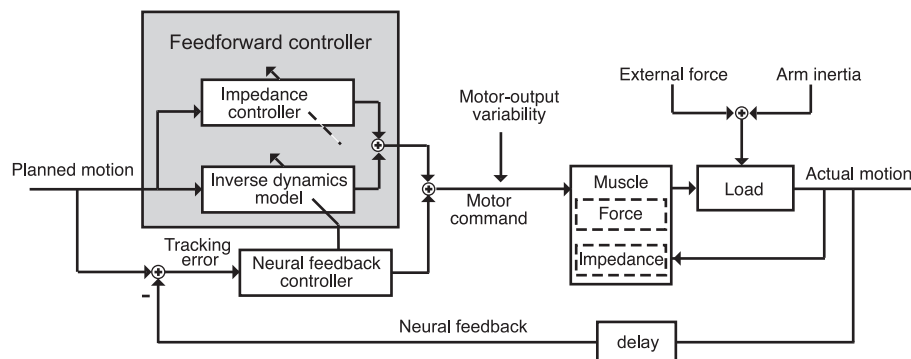


FIG. 13. Control scheme of neuromechanical control of human movements. Planned motion is input to internal dynamics model, as well as to impedance controller. Internal dynamics model computes necessary motor commands to compensate for external force and arm dynamics. Impedance controller computes necessary motor command to achieve mechanical stability. Sum of output of internal dynamics model, impedance controller, and neural feedback controller, with added noise (motor output variability), represents final motor command fed to muscle, which generates joint torque and impedance to compensate for environment dynamics. Actual motion is fed back to neural feedback controller. Both internal dynamics model and impedance controller are modified by learning. Signals from neural feedback controller can be used as learning signals for internal dynamics model (feedback error learning), but not for impedance controller.

existing within the monkey motor cortex for the generation of reciprocal activation and co-contraction of muscles. Smith (1996) discussed the possibility of cerebellar control of impedance based on clinical and experimental lesion studies. Our recent brain imaging experiments also suggest involvement of the cerebellum in dealing with instability (Milner et al. 2002). Behavioral experiments cannot definitively prove that the 2 control and learning mechanisms are neurally distinct, but the separate neuronal systems identified for the generation of reciprocal activation and co-contraction suggest distinct neural mechanisms for the inverse dynamics model and the impedance control.

In Fig. 13 we propose a computational scheme, which includes an impedance controller. The impedance controller operates in parallel with the inverse dynamics model and computes feedforward commands to generate impedance compensating for the environmental instability while receiving information about the environment and the planned trajectory. Impedance depends on intrinsic muscle properties preprogrammed by the impedance controller, as well as on neural feedback reactive to the environment. The sum of feedforward commands from the inverse dynamics model and the impedance controller as well as feedback motor commands constitute the final motor command. This command, contaminated by neural noise, activates muscles, which generate the necessary joint torques to move the limb, while simultaneously selectively controlling impedance to ensure stability.

We thank T. Yoshioka for technical support.

DISCLOSURES

This research was supported by the Telecommunications Advancement Organization of Japan; the Special Coordination Fund to Brain Science, Japan; the Natural Sciences and Engineering Research Council of Canada; and the Human Frontier Science Program.

REFERENCES

- Bhushan N and Shadmehr R.** Computational nature of human adaptive control during learning of reaching movements in force fields. *Biol Cybern* 81: 39–60, 1999.
- Brashers-Krug T, Shadmehr R, and Bizzi E.** Consolidation in human motor memory. *Nature* 382: 252–255, 1996.
- Burdet E, Codourey A, and Rey L.** Experimental evaluation of nonlinear adaptive controllers. *IEEE Control Systems Magazine* 18: 39–47, 1998.
- 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, Osu R, Franklin DW, Yoshioka T, Milner TE, and Kawato M.** A method for measuring endpoint stiffness during multi-joint arm movements. *J Biomech* 33: 1705–1709, 2000.
- Burdet E, Tee KP, Chew CM, Franklin DW, Osu R, Kawato M, and Milner TE.** Stability and learning in human arm movements. *Proc. Int. Conf. on Computational Intelligence, Robotics and Autonomous Systems*, 2001b, p. 355–360.
- Cabel DW, Cisek P, and Scott SH.** Neural activity in primary motor cortex related to mechanical loads applied to the shoulder and elbow during a postural task. *J Neurophysiol* 86: 2102–2108, 2001.
- 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.** Once more on the equilibrium-point hypothesis (lambda model) for motor control. *J Mot Behav* 18: 17–54, 1986.
- Franklin DW, Burdet E, Osu R, Kawato M, and Milner TE.** Functional significance of stiffness in adaptation of multi-joint arm movements to stable and unstable dynamics. *Exp Brain Res* 151: 145–157, 2003.
- Franklin DW, Osu R, Burdet E, Kawato M, and Milner TE.** Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. *J Neurophysiol* 90: 3270–3282, 2003.
- Frysinger RC, Bourbonnais D, Kalaska JF, and Smith AM.** Cerebellar cortical activity during antagonist cocontraction and reciprocal inhibition of forearm muscles. *J Neurophysiol* 51: 32–49, 1984.
- Gandolfo F, Li C, Benda BJ, Schioppa CP, and Bizzi E.** Cortical correlates of learning in monkeys adapting to a new dynamical environment. *Proc Natl Acad Sci USA* 97: 2259–2263, 2000.
- Gandolfo F, Mussa-Ivaldi FA, and Bizzi E.** Motor learning by field approximation. *Proc Natl Acad Sci USA* 93: 3843–3846, 1996.
- Gomi H and Kawato M.** Equilibrium-point control hypothesis examined by measured arm stiffness during multi-joint movement. *Science* 272: 117–120, 1996.
- Goodbody SJ and Wolpert DM.** Temporal and amplitude generalization in motor learning. *J Neurophysiol* 79: 1825–1838, 1998.
- Gribble PL and Ostry DJ.** Compensation for loads during arm movements using equilibrium-point control. *Exp Brain Res* 135: 474–482, 2000.
- Harris CM and Wolpert DM.** Signal-dependent noise determines motor planning. *Nature* 394: 780–784, 1998.
- Humphrey DR and Reed DJ.** Separate cortical systems for control of joint movement and joint stiffness: reciprocal activation and coactivation of antagonist muscles. *Adv Neurol* 39: 347–372, 1983.
- Imamizu H, Miyauchi S, Tamada T, Sasaki Y, Takino R, Putz B, Yoshioka T, and Kawato M.** Human cerebellar activity reflecting an acquired internal model of a new tool. *Nature* 403: 192–195, 2000.
- Imamizu H, Uno Y, and Kawato M.** Internal representations of the motor apparatus: implications from generalization in visuomotor learning. *J Exp Psychol Hum Percept Perform* 21: 1174–1198, 1995.
- Kawato M.** Feedback-error-learning neural network for supervised motor learning. In: *Advanced Neural Computers*, edited by Eckmiller R. Amsterdam: North-Holland/Elsevier, 1990, p. 365–372.
- 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.
- Kawato M, Uno Y, Isobe M, and Suzuki R.** A hierarchical neural network model for voluntary movement with application to robotics. *IEEE Control Systems Magazine* 8: 8–16, 1988.
- Kobayashi Y, Kawano K, Takemura A, Inoue Y, Kitama T, Gomi H, and Kawato M.** Temporal firing patterns of Purkinje cells in the cerebellar ventral paraflocculus during ocular following responses in monkeys. II. Complex spikes. *J Neurophysiol* 80: 832–848, 1998.
- 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.
- Levin MF, Feldman AG, Milner TE, and Lamarre Y.** Reciprocal and coactivation commands for fast wrist movements. *Exp Brain Res* 89: 669–677, 1992.
- Li CS, Padoa-Schioppa C, and Bizzi E.** Neuronal correlates of motor performance and motor learning in the primary motor cortex of monkeys adapting to an external force field. *Neuron* 30: 593–607, 2001.
- Miall RC, Weir DJ, Wolpert DM, and Stein JF.** Is the cerebellum a smith predictor? *J Mot Behav* 25: 203–216, 1993.
- Milner TE, Franklin DW, Imamizu H, and Kawato M.** Selective increase in cerebellar activity for unstable compared to stable dynamics. *Program No 360-3, Abstracts of the Society for Neuroscience 32nd Annual Meeting*, CD-ROM, 2002.
- Mussa-Ivaldi FA and Bizzi E.** Motor learning through the combination of primitives. *Philos Trans R Soc Lond B Biol Sci* 355: 1755–1769, 2000.
- Niemeyer G and Slotine J-JE.** Performance in adaptive manipulator control. *Int J Rob Res* 10: 149–161, 1991.
- 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.

- Sainburg RL, Ghez C, and Kalakanis D.** Intersegmental dynamics are controlled by sequential anticipatory, error correction, and postural mechanisms. *J Neurophysiol* 81: 1045–1056, 1999.
- Scheidt RA, Reinkensmeyer DJ, Conditt MA, Rymer WZ, and Mussa-Ivaldi FA.** Persistence of motor adaptation during constrained, multi-joint, arm movements. *J Neurophysiol* 84: 853–862, 2000.
- Shadmehr R and Holcomb HH.** Neural correlates of motor memory consolidation. *Science* 277: 821–825, 1997.
- Shadmehr R and Mussa-Ivaldi FA.** Adaptive representation of dynamics during learning of a motor task. *J Neurosci* 14: 3208–3224, 1994.
- Shidara M, Kawano K, Gomi H, and Kawato M.** Inverse-dynamics model eye movement control by Purkinje cells in the cerebellum. *Nature* 365: 50–52, 1993.
- Smith A.** Does the cerebellum learn strategies for the optimal time-varying control of joint stiffness? *Behav Brain Sci* 19: 399–410, 1996.
- 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.
- Takemura A, Inoue Y, Gomi H, Kawato M, and Kawano K.** Change in neuronal firing patterns in the process of motor command generation for the ocular following response. *J Neurophysiol* 86: 1750–1763, 2001.
- Thoroughman KA and Shadmehr R.** Electromyographic correlates of learning an internal model of reaching movements. *J Neurosci* 19: 8573–8588, 1999.
- Wang T, Dordevic GS, and Shadmehr R.** Learning the dynamics of reaching movements results in the modification of arm impedance and long-latency perturbation responses. *Biol Cybern* 85: 437–448, 2001.
- Yamamoto K, Kobayashi Y, Takemura A, Kawano K, and Kawato M.** Computational studies on acquisition and adaptation of ocular following responses based on cerebellar synaptic plasticity. *J Neurophysiol* 87: 1554–1571, 2002.