

Temporal and Amplitude Generalization in Motor Learning

SUSAN J. GOODBODY AND DANIEL M. WOLPERT

Sobell Department of Neurophysiology Institute of Neurology, London WC1N 3BG, United Kingdom

Goodbody, Susan J. and Daniel M. Wolpert. Temporal and amplitude generalization in motor learning. *J. Neurophysiol.* 79: 1825–1838, 1998. A fundamental feature of human motor control is the ability to vary effortlessly over a substantial range, both the duration and amplitude of our movements. We used a three-dimensional robotic interface, which generated novel velocity dependent forces on the hand, to investigate how adaptation to these altered dynamics experienced only for movements at one temporal rate and amplitude generalizes to movements made at a different rate or amplitude. After subjects had learned to make a single point-to-point movement in a novel velocity-dependent force field, we examined the generalization of this learning to movements of both half the duration or twice the amplitude. Such movements explore a state-space not experienced during learning—any changes in behavior are due to generalization of the learning, the form of which was used to probe the intrinsic constraints on the motor control process. The generalization was assessed by determining the force field in which subjects produced kinematically normal movements. We found substantial generalization of the motor learning to the new movements supporting a nonlocal representation of the control process. Of the fields tested, the form of the generalization was best characterized by linear extrapolation in a state-space representation of the controller. Such an intrinsic constraint on the motor control process can facilitate the scaling of natural movements.

INTRODUCTION

When we walk, speak, reach, or dance we can vary the rate of the process without changing the spatial pattern of behavior. This rate modulation often is exploited so that new tasks, such as a tennis stroke, are practiced at a slow rate with the assumption that the acquisition at a faster rate will, in some way, be facilitated. Similarly, we can write and draw both on paper or on a blackboard while maintaining the same spatial properties of our impressions. When children are taught to write they start, and are encouraged, to form large letters with the assumption that the skill will someday transfer to the learning of the small characters exhibited in adult writing.

In this paper, we examine the generalization of motor learning to tasks of a different duration or amplitude. This can be considered a parallel process to the recently studied spatial generalization of motor learning in which the effects of learning in one part of the workspace is investigated in parts of the workspace that didn't form the training region. Several studies have investigated a learning paradigm in which subjects learn to make point-to-point movements in the presence of novel force fields that can be position dependent (Flash and Gurevich 1991, 1997; Gurevich 1993), velocity dependent (Gandolfo et al. 1996; Lackner and DiZio 1994; Shadmehr and Mussa-Ivaldi 1994), or acceleration dependent (Sainburg and Ghez 1995). The results show that

when very limited training is given, for example on only two movements, the effect of learning decays rapidly across space (Gandolfo et al. 1996), but that when a region of the workspace is learned the generalization is maintained over space and appears to generalize in intrinsic joint-based coordinates (Sainburg and Ghez 1995; Shadmehr and Mussa-Ivaldi 1994). In these experiments, the temporal components of the movements were maintained while the spatial location was altered systematically.

In the present work, we study motor learning in the temporal and amplitude domains by examining generalization when either the duration of a fixed amplitude movement is halved or the distance moved in a fixed time is doubled. Subjects learned to make movements of a particular duration and amplitude in a specific velocity dependent force field. After learning, subjects were tested on a new movement, either half the duration or twice the distance. To assess the generalization of learning for these faster movements, a variety of force fields were applied in an attempt to find the force field that made these movements kinematically normal.

Two experiments were performed to examine the generalization of the motor learning. In the first experiment, the test force fields were designed to discriminate between five specific hypotheses on the generalization of motor learning. The first possibility (*movement specific*) is that the control process is specific for each movement and as the movement is now different either in duration or amplitude, the controller will have no expectation of a force for any movement other than that of the exposure phase. A second hypothesis (*local*) is that learning is highly local and that no generalization will be seen to velocities not experienced. A third hypothesis (*r^2 rule*) follows from consideration of a possible strategy for scaling movements pointed out by Hollerbach (1982). By considering the dynamic equations of the arm, he noted that scaling the speed of movement produces a class of movement for which there are very simple computations involved. In the circumstance in which the velocity profile shape is maintained but simply scaled in time for movements of different speed, it is possible to avoid having to recompute the torque profile necessary for new movement speeds if the torque profile is already known for one speed. If the movement is made r times as fast, then scaling the time-dependent portion of the torque profile by a factor r^2 and playing it back r times as fast (then adding in the gravity component without any change in amplitude) will achieve the same path but at the new speed. The fourth hypothesis (*position*) is that, as there is a good correspondence between position and velocity for natural movements, the force is internalized in as a function of position. The fifth hypothesis (*linear*) is that the force-velocity relationship is internalized in a functional form and then linearly extrapolated to new speeds.

On the basis of the results of this first experiment, a second generalization experiment was performed in which new test fields were designed to examine the degree of linearity of the generalization. The fields were chosen to test whether the generalization to novel velocities was best characterized by a force-velocity relationship, which decayed smoothly to zero (*decay*), remained at the level experienced for the maximum velocity of the exposure phase (*level*), increased (*supra*), or linearly extrapolated (*linear*).

METHODS

In all sessions, subjects sat with their head in a chin rest and grasped, with their right hand, a handle attached to a lightweight, carbon-fiber robotic manipulator (Phantom haptic interface, Sensable Devices, Cambridge, MA). This robot, which is free to move in three dimensions, can exert forces of ≤ 20 N, in any direction, at its endpoint (backdrive friction 0.02 N, closed loop stiffness 1 N/mm, apparent mass at the tip < 150 g). The handle was free to rotate in all directions about its center, thereby transmitting only translational forces and preventing torques being applied to the hand. The position of the motors (and through the kinematic equations of the robot the position of the hand) were sampled on-line by three optical encoders (10,160 counts per revolution, sampling rate 3,000 Hz) mounted on the three motors. The velocity of the endpoint was obtained by differencing this position signal over a 10-ms window and applying a low-pass digital filter with a 90-ms time constant. This velocity was used in the calculation of the velocity-dependent forces exerted by the Phantom on the arm during movement. The robot was controlled through a Pentium PC. Two infrared emitting diodes (IREDs) were mounted on the robot's distal link. An Optotrak 3020 (Northern Digital, Waterloo, Ontario) also was used to record the position of the markers at 400 Hz. The optotrak was driven from a Silicon Graphics (SGi) Indigo 2 XZ workstation (Silicon Graphics, Mountain View, CA) where the position data were stored for later analysis. Based on these two markers, the position of the center of the hand could be reconstructed for use in the virtual visual feedback display on the SGi.

Visual feedback

The targets and feedback of hand position (as defined by the center point of the handle) were presented as virtual three-dimensional images. This was achieved by projecting the scene from the SGi with a cathode ray tube (CRT) projector (Electrohome Marquee 8000 with P43 low-persistence phosphor green tube, Rancho Cucamonga, CA) onto a horizontal rear projection screen suspended above the subject's head (Fig. 1). A horizontal front-reflecting semisilvered mirror was placed face up below the subject's chin (30 cm below the projection screen). The subject viewed the reflected image of the rear projection screen through field-sequential shuttered glasses (Crystal Eyes, Stereo-graphic, CA) by looking down at the mirror. The SGi workstation displayed left and right eye images ($1,280 \times 500$ pixels) of the scene to be viewed at 120 Hz. The shuttered glasses alternately blanked the view from each eye in synchrony with the display thereby allowing each eye to be presented with the appropriate planar view—subjects therefore perceived a three-dimensional scene. To maintain a high quality force field, the PC was dedicated to controlling the robot, whereas the SGi was used to generate the virtual images and for data capture of the hand position through the Optotrak markers.

Before each session, the position of the IREDs relative to the projected image position was calibrated for each subject. By illuminating the semisilvered mirror from below, the virtual image and the IRED could be lined up by eye. Each subject calibrated on 24




FIG. 1. Experimental apparatus for measuring unconstrained 3-dimensional arm movements under 3-dimensional virtual visual and force feedback. Looking down at the mirror through field sequential glasses, the subject sees the virtual image of the hand and targets. The Phantom haptic interface can generate state-dependent force fields.

points on a three-dimensional grid covering the workspace. A linear regression fit of image position to IRED position was performed, and this then was used on-line to position the targets and hand feedback images. Cross-validation sets gave a mean calibration error of < 0.8 cm.

During the experiments, an opaque sheet was fixed beneath the semisilvered mirror thereby preventing any direct view of the arm. Hand feedback was provided by a 1 cm green wire cube in the virtual scene, and the targets were presented as 1-cm diam colored spheres. By extinguishing the cube, which represented the hand position, movements in the absence of visual feedback could be examined.

Experimental design

EXPERIMENT 1. Six naive, normal, right-handed students (age range 20–27), who gave their informed consent before their inclusion, participated in *experiment 1*. The subjects were familiarized with the equipment and performed two sessions, *Amp* and *Dur*, of arm movements on separate days with the order balanced across subjects.

The subjects were asked to reach “naturally” between the targets—no instructions were given as to the movement path. In all the sessions, the subject's task was to move his arm so as to place the hand cursor at the illuminated target. When the hand was at the target and stationary, the target was extinguished and a tone signaled that the subject should move to another target that became illuminated. Subjects were considered to be on target when they were within 1 cm of the target and their speed was < 3.0 cm s^{-1} .

In session *Dur*, movements were made diagonally between two targets 15 cm apart in the horizontal plane. The subjects were required to make movements of either 500 or 1,000 ms duration cued by two different tones. The target positions were at $(-3, 37.2, -37.6)$ and $(7.6, 26.6, -37.6)$ cm relative to a point midway between the subject's eyes (Fig. 2 shows the target positions and coordinate system). Movements between these targets required subjects to make movements away or toward the body at an angle of 45° to transverse. For the *Amp* session, the targets were either 12.5 or 25 cm apart and the movement duration was fixed at 700 ms for both movement distances. The targets for this session, (Fig. 2, ●) were at $(-8.0, 30.2, -37.6)$ and either $(0.8, 21.4, -37.6)$ or $(9.7, 12.5, -37.6)$ cm for the short and long movements, respectively.




FIG. 2. Coordinate system of data capture is shown — the z axis points out of the plane, and the origin is centered midway between the subject's eyes. Targets lie in the horizontal plane $z = -37.6$ cm. ♦ and ●, targets for the *Dur* and *Amp* sessions, respectively. In the *Amp* session, the targets for the short and long moves are shown as ● and ○, respectively. Half-filled circle indicates that this target was the same for both the short and long moves. Orientation of the intertarget line was preserved between sessions but was shifted for the *Amp* session to keep targets within both reach and stereo range.

In both sessions, subjects were given feedback of their timing performance in the form of a change in the target's appearance at the end of their movements signifying too fast (target turned red), too slow (target turned green), or just right (within 150 ms of desired duration-target turned white). Before each session, subjects practiced making movements of the correct duration. Subjects settled down very quickly and were consistently satisfying the timing criteria within 48 movements.

Each session consisted of 532 movements with a brief rest period after each 50 movements. Three factors could be varied for each point-to-point movement. First, the speed of the movement could be either fast or slow, and this was cued by a tone. Although duration in session *Dur* and amplitude in session *Amp* were manipulated, these both doubled the maximum velocity, and we therefore will refer to the movement of longer duration in *Dur* and of smaller amplitude in *Amp* as slow movements and movements of smaller duration in *Dur* and of larger amplitude in *Amp* as fast movements (see Table 1). Second, visual feedback of hand position could be provided by the virtual cube or else extinguished for the entire movement. Third, the nature of the force field generated by the robot during the movement (including no force field) could be varied.

The session comprised of three phases — baseline (48 movements), exposure (100 movements) and test phase (384 movements) — the movements of interest from these three phases are summarized in Table 2. No indication was given to the subject as to the nature of these phases.

The first 48 movements (baseline) were made in the absence of a force field (with and without visual feedback randomly inter-

TABLE 1. Length and duration of fast and slow movements for the *Dur* and *Amp* sessions of experiment 1

	Slow Speed	Fast Speed
<i>Dur</i>	15.0 (1,000)	15.0 (500)
<i>Amp</i>	12.5 (700)	25.0 (700)

Values are in centimeters at number of milliseconds (in parentheses).

persed). This was done to record standard baseline trajectories for both the fast and slow movements in the absence of visual feedback (S_b and F_b of Table 2). During the next 100 movements (exposure), subjects made slow movements with visual feedback in a force field generated by the robot (S_e movements Table 2). The force field chosen (Fig. 3) was a curl field (Gandolfo et al. 1996) in which the forces acted in the horizontal plane

$$\mathbf{F} = B\mathbf{v} = \begin{bmatrix} 0 & 11 & 0 \\ -11 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \mathbf{v}, \quad (1)$$

where \mathbf{F} is the force vector in Newtons acting on the hand, \mathbf{v} is the velocity vector of the hand in m/s, and B is in $\text{N/m}\cdot\text{s}$. This force field depends only on the velocity of the movement, always acts orthogonally to the direction of motion in the horizontal plane, and magnitude increases linearly with speed. The force field was not changed during these 100 exposure movements.

During the test phase (384 movements), subjects still were exposed to the same curl force field for slow movements with vision but occasionally (on average once in every 4 movements) a slow or fast movement without visual feedback would be required in which the force field was changed to one of six test force fields. These were chosen to test the specific predictions of the five different hypotheses about the controller (Fig. 4). For the movement-specific hypothesis, as the control process is specific for each movement and as the movement is now different either in duration or amplitude, the controller will have no expectation of a force for any movement other than that of the exposure phase (Fig. 4B). In this case, the removal of the force field will be expected to produce the most kinematically normal movements. For the second hypothesis (local), in which learning is highly local to the experienced velocities, no generalization will be seen to velocities not experienced. In this case, the expected force is zero for velocities greater than the maximum velocity experienced during the slow movements (Fig. 4C). While, obviously, a physiological controller would not exhibit such a discontinuity, if the learning is local, then this field is a reasonable model of what is expected. For the third hypothesis, if the r^2 rule proposed by Hollerbach was applied to that part of the torque produced to compensate for the externally applied field, then, when moving twice as fast over the same distance as in session *Dur*, the controller would scale this part of the torque by four, hence the slope of the force-velocity profile expected during movements of twice the speed would have to increase by a factor of two (Fig. 4D). For the fourth hypothesis (position), the force is internalized as a function of position. For temporally scaled movements of fixed amplitude as in session *Dur*, the controller could expect the same sequence of forces as a function of distance for the fast moves as was experienced for the slow moves. However, as the movement is twice as fast, the velocity at each point is doubled, and to recover the same force-position relationship, the force-velocity profile would need to halve in its slope (Fig. 4E). This position hypothesis is clearly only applicable to the *Dur* session. However, a second position hypothesis is that the force is internalized in terms of fraction of movement distance traveled. This also would mean that the system would expect a force velocity profile of half the slope for the fast movements in both *Dur* and *Amp*. The fifth hypothesis (linear) is that the force-velocity relationship is internalized in a functional form and then linearly extrapolated to new speeds (Fig. 4F).

To capture these hypotheses five of the test fields involved changing the slope of the linear relationship between force magnitude and speed

$$\mathbf{F} = gB\mathbf{v} \quad g \in \{0, 0.5, 1.0, 1.5, 2.0\} \quad (2)$$

where g is the gain of this change. A value $g = 0$ corresponds to movements with no force field present. To examine the local

TABLE 2. Movement groups of interest during the three phases of the sessions *Dur* and *Amp* of experiment 1

	Baseline (no vision)	Exposure (vision)	Test (no vision): slope g					Cutoff
			0	0.5	1.0	1.5	2.0	
Slow speed	S_b	S_c	S_0	$S_{0.5}$	S_1	$S_{1.5}$	$S_{2.0}$	S_{cut}
Fast speed	F_b		F_0	$F_{0.5}$	F_1	$F_{1.5}$	$F_{2.0}$	F_{cut}

Vision refers to visual feedback of hand movement. For example $S_{1.5}$ refers to the slow movements made in the absence of visual feedback in the test phase in the presence of a force field the slope (g) of which was 1.5 times that of the exposure trials (S_c).

hypothesis, the sixth force field (cutoff) was one that was the same as the exposure field up to the maximum velocity of the slow movements v_{max} , and zero otherwise

$$\mathbf{F} = B\mathbf{v} \tag{3}$$

where

$$\tilde{v}^i = \begin{cases} v^i & \text{if } |v^i| \leq v_{max}^i \\ 0 & \text{if } |v^i| > v_{max}^i \end{cases} \quad i = x, y, z \tag{4}$$

As all of these fields produce no force for zero velocity, subjects had no prior information, at the start of the movement, about which field they were about to move in. Four fast and four slow movements in each direction were made for each of these test forces.

EXPERIMENT 2. Six naive, normal, right-handed students (age range 21–27), who gave their informed consent before their inclusion, participated in *experiment 2*. None of these subjects participated in *experiment 1*. The subjects were familiarized with the equipment and performed two sessions, *Dur2* and *Control* of arm movements on separate days. Session *Dur2* always was performed first. The apparatus and setup were as in *experiment 1*. Session *Dur2* was identical to session *Dur* of *experiment 1* except for the test phase in which subjects were exposed to four test fields, three of which were different to those of session *Dur*. The forms of the four test fields were chosen to probe the degree of linearity of the generalization beyond v_{max} (Fig. 5).

For velocities less than the maximum velocity experienced during the exposure phase, the four test fields were all identical to that of the exposure phase. For higher velocities the force-velocity relationships were designed to capture the range of generalization from sub- to supralinear, thereby allowing a test of how close the generalization of the controller is to linear extrapolation in state space. For velocities greater than those experienced during the exposure phase, the functional form of the force-velocity relationship either decayed smoothly to zero (decay), remained at the maximum force experienced during the exposure phase (level), extrapolated linearly (linear) or increased (supra) so that

$$\mathbf{F} = B\mathbf{v} \tag{5}$$

For $|v^i| > v_{max}^i$ the functional form of \tilde{v}^i for each of the linear, level, and supra test force fields was $\tilde{v}^i = a + bv^i + ce^{kv^i}$ containing constant, linear, and exponential terms. The decay field was chosen to be a Gaussian $\tilde{v}^i = ae^{-(v^i-b)^2/c}$. All functions were constrained to be continuous up to and including the first derivative at v_{max}^i so that (for clarity the equations are shown for positive velocities only)

$$\tilde{v}^i = \begin{cases} v^i & \text{if } v^i \leq v_{max}^i \\ \left. \begin{aligned} &v_{max}^i + \tau_1(1 - e^{-(v^i - v_{max}^i)/\tau_1}), & \tau_1 = 0.1v_{max}^i, & \text{linear} \\ &2v^i - v_{max}^i + \tau_2(1 - e^{-(v^i - v_{max}^i)/\tau_2}), & \tau_2 = 0.02v_{max}^i, & \text{level} \\ &v_{max}^i e^{((v^i - \mu)^2 - (v_{max}^i - \mu)^2)/(2v_{max}^i(v_{max}^i - \mu))}, & \mu = 1.2v_{max}^i, & \text{decay} \end{aligned} \right\} & \text{if } v^i > v_{max}^i \end{cases} \tag{6}$$

As all of these fields produce no force for zero velocity, subjects had no prior information, at the start of the movement, about which field they were about to move in. Four fast and four slow movements in each direction were made for each of these test forces as in *experiment 1*.

In session *Control* of *experiment 2*, subjects once again performed a baseline, exposure, and test phase in which they made fast and slow movements between the same targets as in session *Dur2*. The baseline phase was identical to that of session *Dur2* (and *Dur* of *experiment 1*), however, the exposure and test phases differed. The exposure field remained the same as in all previous sessions but rather than making 100 slow movements in the field during the exposure phase, subjects instead made 100 fast movements. The test phase consisted of only one type of test field, linear, presented for both fast and slow movements. This control session allowed us to compare generalization of learning from slow speeds to fast with learning solely at fast speeds.

Data analysis

Hand velocities were calculated from the Optotrak marker positions by first differencing the position data and then filtering with a Butterworth second order, zero phase lag, low-pass filter with 5 Hz cutoff. The start of the movement was defined as the time when the hand speed first exceeded 2.5 cm s^{-1} .

To calculate mean hand paths with error bars for various phases of the experiments, the hand position data for each movement was resampled at 50 evenly spaced points along the path length, with linear interpolation between neighboring points. We also analyzed the data resampled at 50 points evenly spaced over the duration of the movement, but as these produced very similar results they will not be presented here. To examine the early stages of movement, including the period before afferent information becomes available, the mean hand paths were calculated for the first 400 ms of the movements at 10-ms intervals. To remove variability due to small changes in the starting location of the movement, the trajectories were translated to align the start points on the start target.

To test between the different hypotheses of *experiment 1* (illus-




FIG. 3. Velocity-dependent vector curl field as a function of velocity. Magnitude and direction of the field at any velocity are indicated by the length and orientation of the arrows.

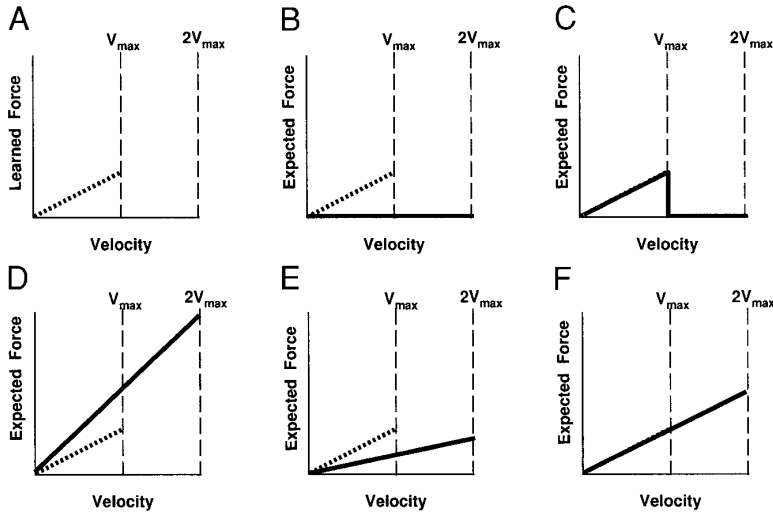


FIG. 4. A schematic representing various hypothetical results of the generalization session of *experiment 1* (see METHODS for details). Plots show force and velocity magnitude relationships. A: learned force-speed relationship. v_{\max} is the maximum speed for the slow movements and $2v_{\max}$ is the maximum speed experienced for the fast movements. B–F: force magnitude vs. speed for predictions of expected force based on 5 different hypotheses about motor learning. \cdots , force speed relationship, which has been learned; — , predicted generalization of the learning under the hypotheses. B: movement specific learning, which does not generalize to movements of a new speed. C: local learning, which does not generalize to new speeds. D: r^2 scaling of torque profiles. E: motor learning of position dependent field. F: linear extrapolation in state space.

trated in Fig. 4), we have developed a measure of generalization, \hat{g} , which is independent of the extent of learning during the exposure phase. If learning for the slow movements resulted in paths that were identical to baseline (i.e., S_b and S_1 movements were identical), then the value of \hat{g} is equivalent to the value of g for the field in which the fast test movements are identical to the fast baseline movements (i.e., F_g is identical to F_b). However, if learning during the exposure phase does not result in slow movement paths that are identical to baseline, then an estimate of the generalization \hat{g} can be derived by quantifying the relationship between the slow baseline movements and the slow tests and between the fast baseline movements and fast tests as follows. For each session and speed of movement, we quantified the location of each point of the mean baseline movement, S_b or F_b , relative to the equivalent points of the movements made in each test force-field, S_{0-2} and F_{0-2} , respectively. This was done by assigning to each point a value corresponding to the weighted average (by inverse Euclidean distance) of the g values of the test fields so

$$W_S(r) = \sum_g w_g^s(r) g / \sum_g w_g^s(r)$$

$$\text{where } w_g^s(r) = |S_b(r) - S_g(r)|^{-1} \quad (7)$$

and

$$W_F(r) = \sum_g w_g^f(r) g / \sum_g w_g^f(r)$$

$$\text{where } w_g^f(r) = |F_b(r) - F_g(r)|^{-1} \quad (8)$$

where $S_b(r)$ and $F_b(r)$ are the hand position vectors at the resampled point r in the baseline slow and fast movements respectively, and similarly $S_g(r)$ and $F_g(r)$ are the hand position vectors at the resampled point r for the slow and fast movements in the test field g .

These hand position vectors include all three (x, y, z) components. These weighted averages $W_S(r)$ and $W_F(r)$ represent what has been learned relative to the test fields. For example, for a slow movement, if a baseline point $S_b(r)$ was spatially close to the corresponding point for a force field of $g = 1.0$, the value of $W_S(r)$ for this point would be near 1.0. In this particular example, the linear hypothesis would be supported by a value of $W_F(r) = 1.0$ and the position hypothesis by $W_F(r) = 0.5$, which are identical to the g values for the corresponding fields. If, however, for a slow movement, a baseline point $S_b(r)$ was spatially close to the corresponding point for a force field of $g = 1.5$, the value of $W_S(r)$ for this point would be near 1.5 and the linear hypothesis would be now be supported by a value of $W_F(r) = 1.5$ and the position hypothesis by $W_F(r) = 0.75$. The original hypotheses are therefore more generally parameterized by \hat{g} the ratio of $W_F(r)$ to $W_S(r)$. Point-wise estimates of the generalization then were obtained by calculating the ratios

$$\hat{g}(r) = W_F(r)/W_S(r) \quad (9)$$

This ratio, \hat{g} describes numerically the similarity of the relationship between slow baseline and slow test movements and that between the fast baseline and fast test movements. The ratio was calculated for the paths resampled over path length as well as for the first 400 ms of movement. The value of \hat{g} was used to test the hypotheses. Confidence intervals for this estimate were calculated by bootstrapping (Efron 1982).

The measure of the generalization \hat{g} developed for *experiment 1* is not appropriate for the data of *experiment 2* in which all of the four test fields are identical for the slow movements. To determine which test field in *experiment 2* best approximated the forces expected due to generalization, we simply calculate $W_F(r)$ from Eq. 8. The g values in this equation, which parameterize the fields, were

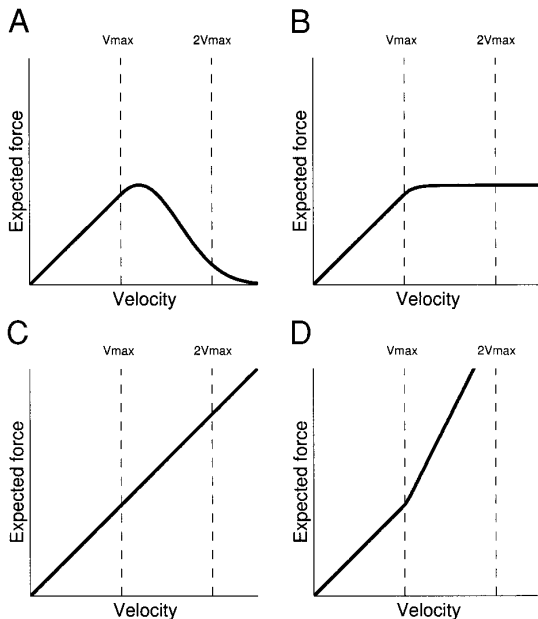


FIG. 5. A schematic representing the test fields of *experiment 2* (see METHODS for details). Plots show force and velocity magnitude relationship for the decay (A), level (B), linear (C), and supra (D) fields. v_{\max} is the maximum speed for the slow movements.

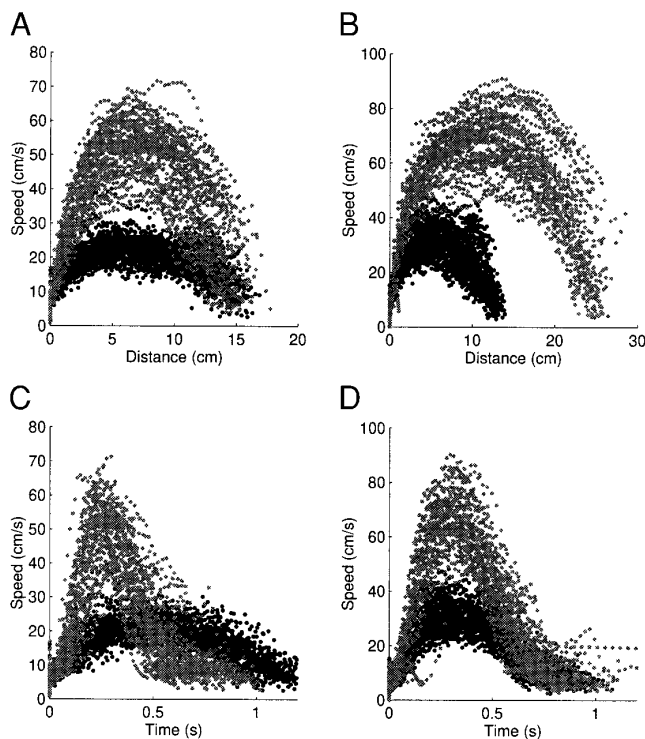


FIG. 6. Speed against distance and time from start of movement for slow (\bullet) and fast (\square) movements made without visual feedback during the test phase (subject *EKA*, 24 slow and 24 fast movements in each direction). Slow movements during the other phases cover a similar region but have been excluded to allow individual points to be resolved. *A*: session *Dur* speed against distance. *B*: session *Amp* speed against distance. *C*: session *Dur* speed against time. *D*: session *Amp* speed against time.

taken as the area under the force-velocity curve (Fig. 5) between v_{\max} and $2v_{\max}$ normalized so that g for the linear field was 1.0.

RESULTS

Experiment 1

Although none of the subjects had previously experienced a virtual environment, they found the task natural and easy to perform. Figure 6 shows the hand speed plotted against distance moved and time, for the slow and fast movements, for all test fields, for a typical subject. All x , y , and z components of the hand velocity and position are used in the calculation. The relationship of the state of the hand, that is, position and velocity, to this plot is many-to-one such that different points in this plot correspond to different states of the hand. We can conclude, then, that although there is some overlap between the regions, the majority of states explored during the fast movements were not experienced during the slow movements of the exposure or test sessions. Analysis of the maximum speeds for each subject's movements made in the exposure field ($g = 1$) for both the slow and fast movement of the test phase (S_1 and F_1 of Table 2) showed that for all of these 192 movements (bar 1 movement for 1 subject) all the maximum speeds of the fast movements were greater than the maximum speeds for the slow movements. Therefore for all subjects, the range of velocities experienced during the slow movements was a subset of velocities experienced during the fast movements.

The subjects' mean hand paths in the plane of the targets (xy) for different phases of the experiment are shown in Fig. 7 (session *Dur*, left, and session *Amp*, right). In all of the following plots, the paths shown are those resampled over path length. The baseline movements made in the absence of a force field show typical straight-line paths for both directions of movement (Fig. 7, *A* and *D*). When the force field was first introduced during the exposure phase, the hand paths deviated from the baseline paths (Fig. 7, *B* and *E*). However, by the end of the exposure phase, the hand paths had straightened, to be closer to the baseline movements, despite the presence of the force-field (Fig. 7, *C* and *F*). Many studies have shown that the goal of the learning is a return to baseline (Flanagan and Rao 1995; Gurevich 1993; Lackner and DiZio 1993; Shadmehr and Mussa-Ivaldi 1994; Wolpert et al. 1995a), therefore we conclude that over the course of the exposure phase subjects learned, without instruction, to produce movements similar to their baseline movements. During the test phase, subjects were exposed to a set of novel force-velocity relationships for both the slow and fast movements. The test movements at the slow rate were used to investigate what had been learned from the exposure phase, whereas the test movements at the fast rate were designed to probe the generalization of this learning to novel states. The performance in the exposure field, $g = 1$, for the slow movements remained stable during this test phase. Figure 8 shows the mean hand paths in the plane of the targets for the last 10 moves of the $g = 1$ exposure phase (vision on) and all of the $g = 1$ vision on moves of the test phase. Also shown for comparison are the vision off, $g = 1$, slow test moves well as slow, vision off, baseline moves. The similarity of all these $g = 1$ movements, both for vision on and off, is evidence that the performance in the $g = 1$ field did not change from the exposure to the test phase in which on average once every four moves one of the six test fields was randomly substituted for the $g = 1$ exposure field.

Figures 9A and 10A show the mean paths in the plane of the targets (xy), for all subjects for movements made at the slow rate during the test phase under the different test force-fields in sessions *Dur* and *Amp*. Also shown for comparison are the mean baseline movements for all subjects. All these movements were made in the absence of visual feedback. In both sessions, *Dur* and *Amp*, the hand paths of the slow movements follow a well-defined progression for test fields from $g = 0$ through to $g = 2$. For $g = 0$, which represents movements made in the absence of a force field, large after-effects (a term that has been used to describe the postexposure changes seen on removal of the prisms after prism adaptation — we will use this term to represent any changes in performance seen after removal of a perturbation) can be seen in the path. This demonstrates that the change in performance over the exposure phase represents more than just a nonspecific process such as cocontraction. As might be expected, these effects are spatially reciprocal to the deviations from straight-line paths seen on the introduction of the force field in the exposure phase (Fig. 7, *B* and *E*). Increasing the slope (g) of the test force-velocity relationship leads to the paths becoming straighter and, therefore, more like the baseline movements. However, as g is increased to values as large as 2.0, the paths resemble those seen on the introduction of the novel force field during the




FIG. 7. Mean hand paths in the plane of the targets (xy) with standard error bars for the slow movements in the presence of visual feedback for all subjects in session *Dur* (A–C) and session *Amp* (D–F). For clarity, the 2 movement directions have been offset—the movement direction is indicated (\rightarrow). A and D: baseline movements ($n = 36$). B and E: first movement in force field during the exposure phase ($n = 6$). C and F: last 10 movements in the force field of the exposure phase ($n = 60$).

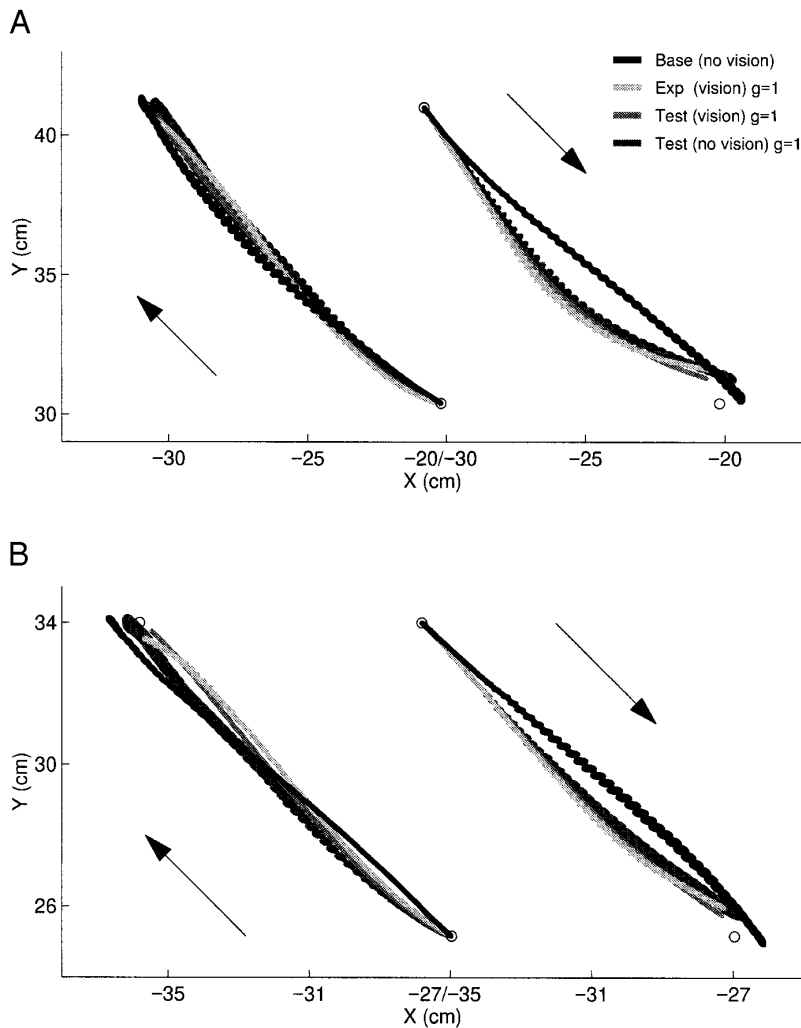


FIG. 8. Mean hand paths in the plane of the targets (xy), for all subjects, with standard error bars for slow movements made in the exposure field $g = 1$ during different phases of the experiment: the last 10 moves of the $g = 1$ exposure phase ($n = 60$), all of the $g = 1$ movements with vision during the test phase ($n = 864$), and all the $g = 1$ movements without vision in the test phase ($n = 24$). Also shown are the slow baseline moves (vision off, $n = 36$) for comparison. For clarity, directions of movement have been offset—the movement direction is indicated (\rightarrow). A: session *Dur*. B: session *Amp*.




FIG. 9. Session *Dur*. Mean hand paths in the plane of the targets (xy), for all subjects, with standard error bars for movements made without visual feedback in each of the different force fields of the test session ($n = 24$) as well as baseline movements ($n = 36$). For clarity, directions of movement have been offset — the movement direction is indicated (\rightarrow). *A*: slow movements of 15 cm amplitude in 1,000 ms. *B*: fast movements of 15 cm amplitude in 500 ms.

exposure phase (Fig. 7, *B* and *E*). At intermediate values of g , the paths tend to resemble the baseline movements. The relation of the baseline movements to these parameterized paths was used to assess what has been internalized by the control process.

For movements made away from the body (outward) in the *Dur* session (Fig. 9*A*), the test field $g = 1.0$ is most similar to the baseline and therefore represents the motor learning. However, for movements made toward the subject (inward), the hand path is such that $g = 0.5$ best represents the baseline movements. Similarly, for session *Amp*, slow outward movements have hand paths that look most similar to the baseline for movements in test force field $g = 1$ (the exposure field), whereas for inward movements, the hand paths for test force field $g = 0.5$ once again look closest to the baseline (Fig. 10*A*). It could be argued that the curved paths in test field $g = 2.0$ arose not because the system had internalized and hence was expecting a weaker force field, as is proposed here, but because the forces were so strong that subjects would always make highly curved movements in this field. To control for the latter possibility, two of the original six subjects performed the baseline and training phase parts of the session in the $g = 2$ field. It was found that, as for the $g = 1$ field, on initial exposure to the field, subjects' hand paths were highly perturbed but with expo-

sure, subjects learned to make straight movements. Therefore it can be concluded that the curved paths for $g = 2$ test phase moves result because the system has internalized and hence is expecting a weaker field.

The cutoff force-field (Fig. 11*A*) shows that for the slow movement, as expected, the paths are similar to those made in the $g = 1.0$ field as for this velocity range these two force fields are identical.

The generalization of this motor learning to the novel state space experienced during the fast movements of the test phase is shown in Figs. 9*B* and 10*B* for the *Dur* and *Amp* sessions, respectively. These are in the same format as for the slow movements. On removal of the force field ($g = 0$), large aftereffects are present, demonstrating that some motor learning is still present for these fast movements. Similarly, for the cutoff test field, which is zero for velocities greater than those experienced during the slow movements, large deviations can be seen relative to the baseline movements (Fig. 11). The dissimilarity of movements in these two fields to the baseline movements is evidence of generalization of learning from the slow to the fast movements as it demonstrates that a nonzero force field is expected for the faster movements. Therefore there is generalization to new states not visited during the learning. For both movement directions and sessions, the spatial pattern of hand paths in

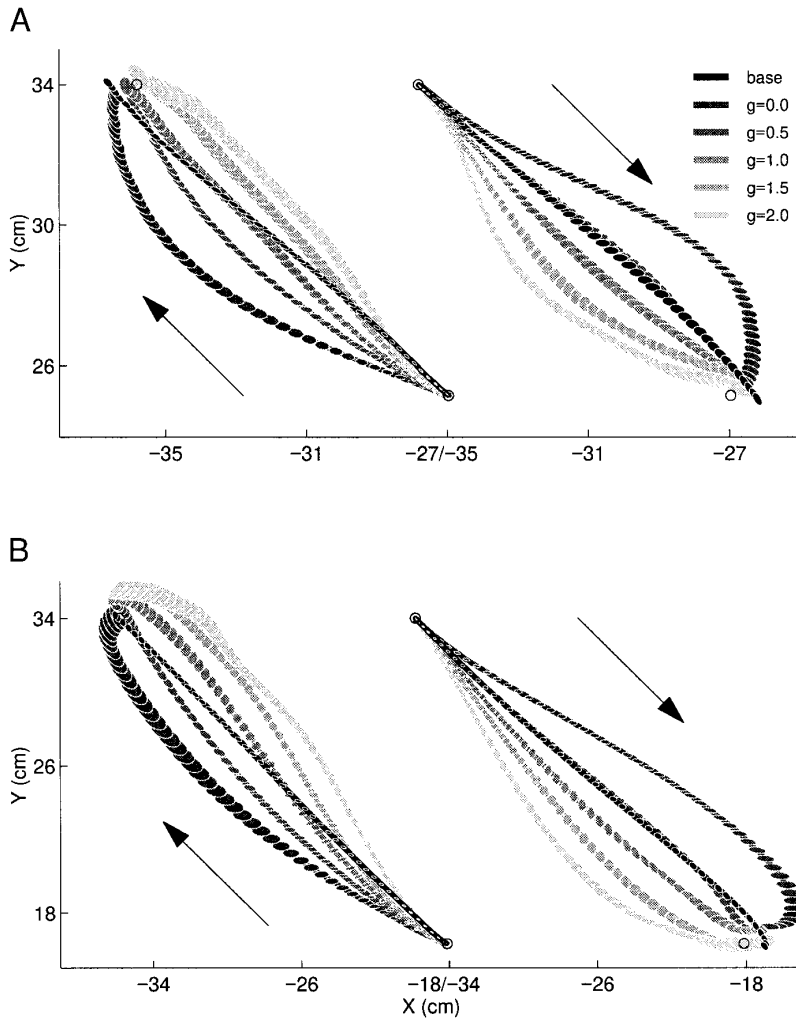


FIG. 10. Session *Amp*. Mean hand paths in the plane of the targets (xy), for all subjects, with standard error bars for movements made without visual feedback in each of the different force fields of the test session ($n = 24$) as well as baseline movements ($n = 36$). For clarity, directions of movement have been offset — the movement direction is indicated (\rightarrow). *A*: slow movements of 12.5 cm amplitude in 700 ms. *B*: fast movements of 25 cm amplitude in 700 ms (note the change in scale of these axes compared with *A*).




FIG. 11. Mean hand paths in the plane of the targets (xy), for all subjects, with standard error bars for 1 direction of movement in both the exposure ($g = 1.0$, $n = 24$) and cut-off fields ($n = 24$) are shown together with the baseline movements ($n = 36$). *A* and *B*: slow movements for sessions *Dur* and *Amp*, respectively. *C* and *D*: fast movements for sessions *Dur* and *Amp*, respectively.

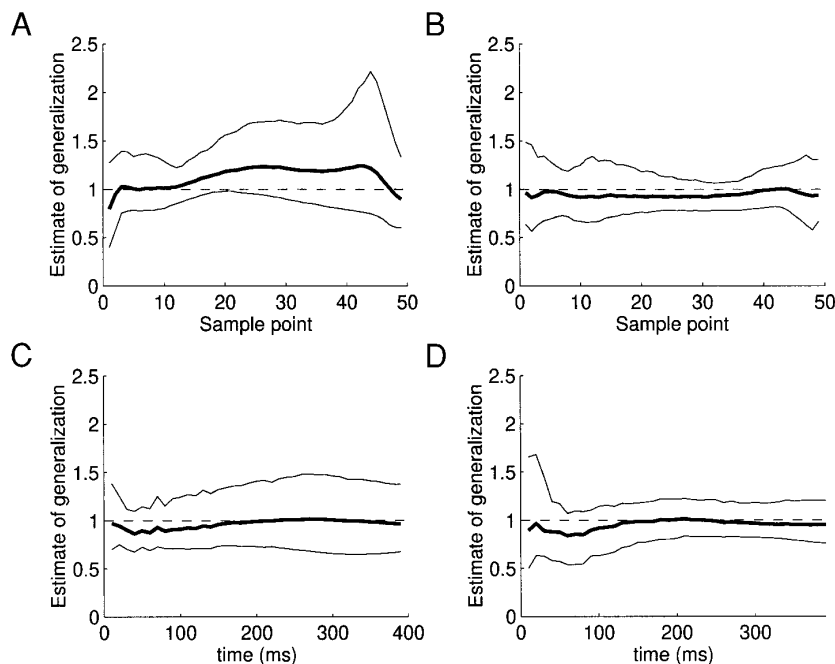


FIG. 12. Plot of estimate of generalization \hat{g} against sample point with 95% confidence intervals. *A*: session *Dur*, paths resampled over path length. *B*: session *Amp*, paths resampled over path length. *C*: session *Dur*, 1st 400 ms. *D*: session *Amp*, 1st 400 ms.

the different fields is similar for both the fast and slow test movements (compare Fig. 9, *A* with *B*, as well as Fig. 10, *A* and *B*). The relationship between hand paths for the slow baselines and the slow test movements ($g = 0-2$) is preserved when the movement speed is doubled so that a similar relationship exists between the fast baselines and the fast movements in these five test fields (Figs. 9 and 10). If, for example, there were no generalization, then the paths for the fast movements would be expected to look closest to baseline in the cutoff test field and the spatial patterns of Fig. 9, *A* compared with *B*, as well as Fig. 10, *A* compared with *B*, to be very different to one another. The similarity of the relationship between baseline paths and paths in the test fields for the fast and slow movements suggests that whatever is learned for the slow movements generalizes to the fast movements. This generalization is quantified in the calculation of \hat{g} of Eq. 8. For example, consider the inward movements of the *Amp* session (Fig. 10, *right*). For the slow movement, the baseline is almost identical to the $g = 0.5$ field and therefore the weighted average W_S of Eq. 7 would be expected to be close to 0.5. For the fast movements, the baseline again is almost identical to the $g = 0.5$ and the weighted average W_F of Eq. 8 would be expected to be close to 0.5. The estimate of the generalization \hat{g} , the ratio of these weighted averages, would be 1.0 supporting the linear model for this case. This analysis assumes that generalization is by definition the change in behavior from baseline in regions of state space not experienced during the learning, that is, for the fast movements. This analysis allows quantification of the observation that the relationship between the control paths and the paths in each of the test fields is similar for fast and slow movements. Figure 12, *A* and *B*, shows a plot of \hat{g} , as defined in Eq. 9, against sample point for paths resampled over path length for sessions *Dur* and *Amp*, respectively. The 95% confidence intervals also are shown. All components x , y , and z are included in this path analysis. Figure 12, *C* and *D*, plot \hat{g} against real time for the first 400

ms of the movement. Both sets of analyses yield the same result as does the calculation of \hat{g} for the paths resampled over time (not shown). This quantifies the generalization of motor learning. This estimate is not significantly different from 1.0 for any point along the path ($P > 0.05$), but is significantly different ($P < 0.05$) from 0 (movement specific), 0.5 (position hypothesis), and 2.0 (r^2 rule hypothesis) at the majority of points. This supports the linear hypothesis that $\hat{g} = 1$.

Experiment 2

During the test phase of *experiment 2*, subjects were exposed to four novel force-velocity relationships for both the slow and fast movements, which were designed to further examine the degree of linearity of the generalization found in *experiment 1*. The test movements (Fig. 13*A*) at the slow rate were, as expected, identical as the fields were identical up to v_{\max} . The test movements at the fast rate were made to probe the generalization of the learning to novel states. In this experiment, subjects tended to move more than twice the speed for the fast test movements compared with the slow (maximum speed ratio of fast to slow test movements 2.5 ± 0.2 ; mean \pm SE). As in *experiment 1*, the performance in the exposure field for the slow movements remained stable during this test phase. The movements at the end of the exposure session had returned to the preperturbation paths (the test trials for the slow movements are shown in Fig. 13*A*), suggesting that this group had adapted more completely than the subjects in *experiment 1*. Figure 13*B* shows that for the fast movements, the hand paths are initially similar, as expected, because all of the test fields were identical for speeds less than v_{\max} . As the speed becomes greater than v_{\max} , the paths diverge due to the differences in the fields. The relationship between these paths and the baseline fast movements was used to assess the generalization of the learning. For both movement directions, the pattern of hand

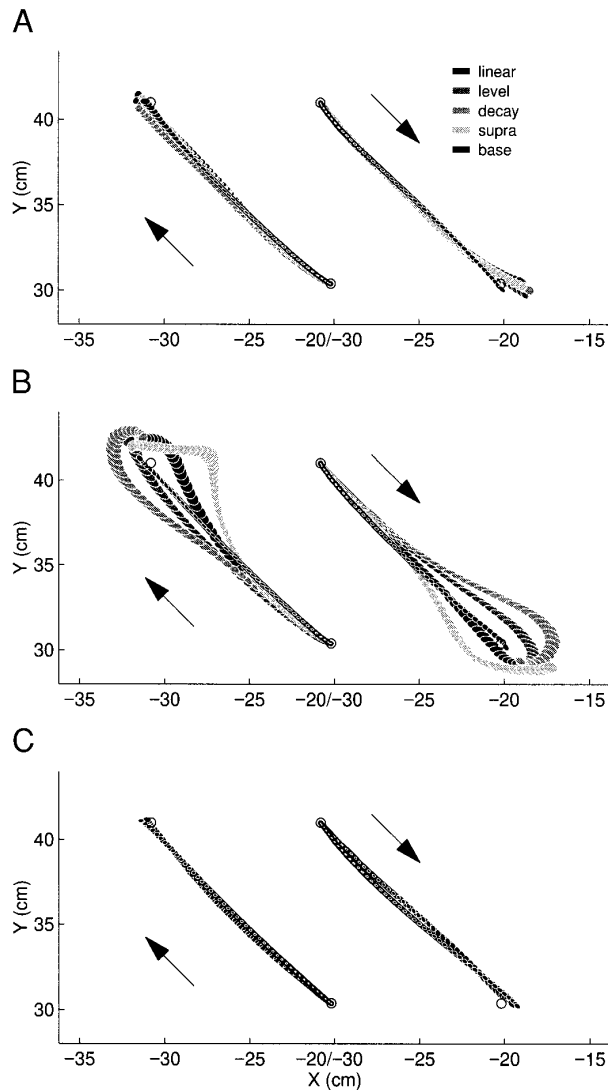


FIG. 13. Mean hand paths in the plane of the targets (xy), for all subjects, with standard error bars for movements made without visual feedback in each of the different force fields of the test session ($n = 24$) as well as baseline movements ($n = 36$) in sessions *Dur2* and Control. For clarity, directions of movement have been offset — the movement direction is indicated (\rightarrow). *A*: session *Dur2*, slow movements of 15 cm amplitude in 1,000 ms. *B*: session *Dur2*, fast movements of 15 cm amplitude in 500 ms. *C*: session Control, fast movements of 15 cm amplitude in 500 ms.

paths for the fast movements in the different test fields is similar. The paths in the decay field soon diverge from baseline, showing that the generalization does not decay for speeds greater than v_{\max} . Similarly the supra test fields does not capture the generalization. The generalization of learning at speeds less than v_{\max} to speeds greater v_{\max} lies between the linear and level fields. Figure 13C plots the fast movements in the linear test field of session Control. A comparison of Fig. 13, A and C, shows that fast movement paths in the linear field after extensive exposure to that field for the fast movements (Fig. 13C, Control) are as close to baseline as slow movement paths after exposure for slow movements (Fig. 13A, *Dur2*).

Figure 14, A and B, are plots of $W_F(r)$, the weighted average of the g values as defined in Eq. 8, for fast move-

ments of session *Dur2* against sample point over the entire path as well as against time for the first 400 ms. This quantifies the generalization of motor learning and shows that learning lies between the linear [$W_F(r) = 1$] and the level [$W_F(r) = 0.73$] although closer to the linear form of generalization. The failure of the decay test field to capture the generalization further demonstrates that the learning is non local and there is substantial generalization to new states.

DISCUSSION

Summarizing the results, learning a novel dynamical environment, for a single movement, generalizes substantially to movements of the same orientation of either increased rate or amplitude. This generalization was quantified by assessing which of a set of force fields produced the most kinematically normal movements at the faster rate — this represents the generalization of motor learning. We found extensive generalization, which, among the hypotheses tested, was best captured by a linear extrapolation of the force field represented in state space.

Returning to the hypotheses considered in METHODS: the first, movement specific, in which the learning would be specific only to that movement and therefore no generalization would be seen, can be ruled out by the presence of substantial aftereffects for the faster movements when the force field was removed unexpectedly ($g = 0$ in Figs. 9B and 10B). There is, therefore, generalization of learning between movements of different rates. Movements made in the cutoff force field also showed large deviations from the baseline movements (Fig. 11), showing that learning was not purely limited to the states already visited. This rules out a local-look-up representation of the control process (hypothesis local). To quantify the generalization, the relationship of the baseline movements to movements under five parameterized force-fields was examined (different slopes of the velocity-force relationship, g values). Both the *Dur*




FIG. 14. Plot of $W_F(r)$ against sample point with 95% confidence intervals for session *Dur2*. — — —, g values for the decay ($g = 0.53$), level ($g = 0.73$), linear ($g = 1.0$), and supra ($g = 1.35$) fields. *A*: paths resampled over path length. *B*: first 400 ms. All components x , y , and z are included in this path analysis.

and *Amp* sessions showed a generalization value \hat{g} that was not significantly different from 1.0. However, \hat{g} was significantly different from 0.5 and 2.0, which represent the position and r^2 rule hypotheses, respectively. In *experiment 2*, an extended examination of the linearity of the generalization showed that, although the generalization lay between the level and linear fields, of the fields tested, it was best captured by the linear extrapolation field. In particular, a decaying generalization was not supported. However, fast movements in the linear test field of session Control, in which subjects had been exposed extensively to this field, were significantly closer to baseline than fast movements of the generalization session in the linear field. We conclude, therefore, that although training in the linear field produced superior performance, among the hypotheses tested, the generalization best supports the linear hypothesis.

We can interpret our results by considering motor learning within the general framework of adaptation of an internal model. Internal models, which have emerged as an important theoretical concept in motor control (Jordan 1995; Kawato et al. 1987), are so named as they internalize or mimic some aspect of a natural process such as the arm's dynamics. Two varieties of internal model are forward and inverse models. Forward models, which mimic the causal flow of a process by predicting its next state (e.g., position and velocity) given the current state and the motor command, have been shown to be computationally useful in planning, control, and learning (Gallistel 1980; Ito 1984; Jordan and Rumelhart 1992; Miall et al. 1993; Robinson et al. 1986; Sutton and Barto 1981; Wolpert 1997) and recently there is evidence that such an internal model is used during the human sensorimotor integration task of localizing the limb position during movement (Wolpert et al. 1995b). A second type of internal model is the inverse model, which inverts the causal flow by estimating the motor command that causes a particular state transition. Such inverse models are of use in control and can function either as a purely feedforward controller (Jordan and Rumelhart 1992) or in conjunction with feedback control (Gomi and Kawato 1993; Kawato 1990). The motor learning task can be considered as either adaptation or augmentation of an internal model to incorporate the changes in motor command necessary to counter the external force field.

The driving force for adaption is assumed to be a desired or planned trajectory. Evidence for such a kinematic-based planning process comes from both studies of natural and perturbed movements. Point-to-point movements show invariant features at the behavioral level (Bernstein 1967) — subjects tend to move their hands along a straight path with a single-peaked, bell-shaped velocity profile (Abend et al. 1982; Atkeson and Hollerbach 1985; Bernstein 1967; Flash and Hogan 1985; Kelso et al. 1979; Morasso 1981; Uno et al. 1989). These features are independent of the hand's initial and final position within the workspace. In contrast, the joint angular position and velocity profiles show considerable variation depending on the hands initial and final position within the workspace (Morasso 1981). Recently it has been shown that such invariants are not necessarily at odds with joint-based planning models such as minimum torque-change (Uno et al. 1989). Recent perturbation studies, however, argue for kinematic planning. If the perceived path of

the movement is altered, either by surreptitiously adding visual curvature to the movement (Wolpert et al. 1994, 1995a) or by representing the visual feedback of hand position in joint-based coordinates (Flanagan and Rao 1995), subjects change their actual hand path to visually straighten their perceived paths. Similarly, in dynamic adaption studies, as well as our present study, it has been shown that in the presence of a perturbing force field (Gurevich 1993; Lackner and DiZio 1993; Shadmehr and Mussa-Ivaldi 1994), subjects adapt to regain preperturbation kinematics. Although still an area of controversy (see Kawato 1996 for a review of dynamic based planning), we believe these results argue for a kinematically based plan for simple point-to-point movements in which there is a hierarchical separation of the planning and control aspects of movement. However, studies of more complex movements around an obstacle suggest that knowledge of the dynamics of the arm is used in planning. Subjects tend to select their movement paths so as to ensure that their closest point of approach to the obstacle is on an axis where the arm is most inertially stable (Sabes and Jordan 1997). Given a kinematic plan, several computational methods have been proposed by which trajectory errors can be used to adapt the internal model appropriately so as to reduce these errors (Gomi and Kawato 1993; Jordan and Rumelhart 1992; Kawato 1990). Therefore, adaptation to the perturbing force field is consistent with the hypothesis that with exposure to the field the CNS learns to build an internal model of that field so as to counteract its effect, thereby producing the desired straight-line movements.

The way in which such an internal model generalizes can be viewed from the perspective of function approximation. In this framework, motor learning consists of approximating the function between desired kinematics and motor commands based on the limited set of movement states experienced during the exposure phase of the experiment. As there are infinitely many possible functions consistent with any finite set of experience, the problem is ill-posed. The mathematical theory of function approximation states that, to obtain a solution to this ill-posed problem, constraints have to be placed on the function approximator (Tikhonov and Arsenin 1977). The pattern of recalibration that results from this limited exposure, the generalization, reflects the structure and constraints underlying the internal model (Ghahramani and Wolpert 1997; Ghahramani et al. 1996; Imamizu et al. 1995).

Our results show that the motor control process shows substantial nonlocal generalization to temporal and amplitude scaling of an individual movement. Of the fields tested, the generalization was best characterized by a linear extrapolation of a state-space representation of the force field. In other words, the control process could learn the relationship between velocity and force and then extrapolate this form in a linear fashion for new states. A recent study found a rapid decay in spatial generalization, for different directions of movement, after learning a limited set of movements (Gandolfo et al. 1996), implying that the internal model is local and decays smoothly with distance from the exposure region. In that study, both the training and test movements were carried out at the same speed and amplitude, and therefore it did not address the temporal or amplitude scaling of motor learning. These results showed that generalization is

local in direction, whereas our study demonstrates that for a single direction of motion, motor learning generalizes extensively even over a twofold increase in either velocity or positional range. Taken together, these results suggest that the intrinsic constraints impose a more powerful ability to generalize for scaled movements, either temporally or spatially, compared with those involving spatial translations and rotations. This ability to extrapolate a single movement to new temporal rates or amplitudes would be of functional importance in the scaling of natural movements.

Previous studies have examined the generalization of skill learning in the temporal and amplitude domain. These studies have focused on learning to pursue targets of different rates and amplitudes (Lordahl and Archer 1958; Namikas and Archer 1960) or on single degree of freedom movements made at different amplitudes (Jaric et al. 1993). While learning for these tasks shows some generalization to new speeds or amplitudes, unlike the present study they do not explore novel dynamics and therefore do not address the generalization of internal dynamic learning. The form of the generalization in these tasks therefore might be attributable to changes in strategy, target prediction, or plan. In a recent study of generalization of prism adaptation to novel speeds, Kitazawa et al. (1997) introduced a kinematic distortion for movements at one speed and found that the generalization of the adaptation decayed as the movement speed was changed from that of the exposure phase. The authors proposed that the observed velocity specific adaptation for this visual perturbation could have occurred in the visuomotor transformation at the level of movement kinematics or dynamics. The present study specifically investigates adaptation and generalization at the level of internal dynamic learning in which the perturbation was designed specifically to depend on the velocity of movement. Under these conditions, extensive generalization in velocity was seen.

Our results can be interpreted within the known neurophysiological properties of motor cortical neurons. It is generally observed that cell activity is broadly tuned to the direction of movement (Georgopoulos et al. 1982) although there is controversy over whether this tuning represents extrinsic attributes (Georgopoulos et al. 1982, 1986), such as direction of hand movement, or intrinsic attributes (Scott and Kalaska 1995, 1997), such as joint or muscle variables. When an arm movement is made in a particular direction, a large group of motor cortical cells are active and participate in coordinating the arm movement, with each individual cell's activity dependent on its tuning curve. When a movement in a different direction is made, either a different population of neurons must become active or the same population but with different relative activations among the neurons (or a combination of both processes). However, for movements of the same direction but of different speeds, the same population is active but at a globally different level of activity (Schwartz 1994). Scaling of a movement, either temporally or in amplitude after learning novel dynamics, therefore could involve the same population of neurons that were involved in the learning process, activated globally at a different level, whereas changes in direction may involve either a different population of neurons or a relative change in activity within the population. Therefore temporal and amplitude scaling could have a more parsimonious coding and

a global change in activity of a population within this framework and could account for the difference in generalization ability between the scaling and spatial domains.

In conclusion, we have shown extensive generalization of motor learning to temporal and amplitude scaling of a single orientation of movement. This generalization is well captured as a linear extrapolation of the control process when represented and parameterized by state space. This powerful ability may be of functional importance in the scaling of natural movements.

This work was supported by grants from the Wellcome Trust and the Physiological Society.

Address reprint requests to Daniel M. Wolpert.

Received 2 June 1997; accepted in final form 18 December 1997.

REFERENCES

- ABEND, W., BIZZI, E., AND MORASSO, P. Human arm trajectory formation. *Brain* 105: 331–348, 1982.
- ATKESON, C. G. AND HOLLERBACH, J. M. Kinematic features of unrestrained vertical arm movements. *J. Neurosci.* 5: 2318–2330, 1985.
- BERNSTEIN, N. *The Coordination and Regulation of Movements*. London: Pergamon, 1967.
- EFRON, B. *The Jackknife, the Bootstrap and Other Resampling Plans*. Philadelphia: Soc. Ind. Appl. Math., 1982.
- FLANAGAN, J. R. AND RAO, A. Trajectory adaptation to a nonlinear visuomotor transformation: evidence of motion planning in visually perceived space. *J. Neurophysiol.* 5: 2174–2178, 1995.
- FLASH, T. AND GUREVICH, I. Human motor adaptation to external loads. *IEEE Eng. Med. Biol. Soc. Conf.* 13: 885–886, 1991.
- FLASH, T. AND GUREVICH, I. Models of motor adaptation and impedance control in human arm movements. In: *Self-Organization, Computational Maps and Motor Control*, edited by P. Morasso and V. Sanguineti. Amsterdam: Elsevier, 1997, p. 423–481.
- FLASH, T. AND HOGAN, N. The co-ordination of arm movements: an experimentally confirmed mathematical model. *J. Neurosci.* 5: 1688–1703, 1985.
- GALLISTEL, C. R. *The Organization of Action: A New Synthesis*. Hilldale, NJ: Erlbaum, 1980.
- GANDOLFO, F., MUSSA-IVALDI, F. A., AND BIZZI, E. Motor learning by field approximation. *Proc. Natl. Acad. Sci. USA* 93: 3843–3846, 1996.
- GEORGOPOULOS, A. P., KALASKA, J. F., CAMINITI, R., AND MASSEY, J. T. On the relation between the direction of two-dimensional arm movements and cell discharge in primate motor cortex. *J. Neurosci.* 2: 1527–1537, 1982.
- GEORGOPOULOS, A. P., SCHWARTZ, A. B., AND KETTNER, R. E. Neuronal population coding of movement direction. *Science* 233: 1416–1419, 1986.
- GHAHRAMANI, Z. AND WOLPERT, D. M. Modular decomposition in visuomotor learning. *Nature* 386: 392–395, 1997.
- GHAHRAMANI, Z., WOLPERT, D. M., AND JORDAN, M. I. Generalization to local remappings of the visuomotor coordinate transformation. *J. Neurosci.* 16: 7085–7096, 1996.
- GOMI, H. AND KAWATO, M. Neural network control for a closed-loop system using feedback-error-learning. *Neural Networks* 6: 933–946, 1993.
- GUREVICH, I. *Strategies Used by the Human Central Nervous System in the Control of Planar Two-Joint Movement in Response to a Change of External Conditions* (PhD thesis). Revohot, Israel: Weizmann Institute of Science, 1993.
- HOLLERBACH, J. M. Computers, brains and the control of movement. *Trends Neurosci.* 5: 189–192, 1982.
- 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.
- ITO, M. *The Cerebellum and Neural Control*. New York: Raven Press, 1984.
- JARIC, D., CORCOS, D. M., AGARWAL, G. C., AND GOTTLIEB, G. L. Principles for learning single joint movements. II. Generalizing a learned behavior. *Exp. Brain Res.* 94: 514–521, 1993.
- JORDAN, M. I. Computational aspects of motor control and motor learning.

- In: *Handbook of Perception and Action: Motor Skills*, edited by H. Heuer and S. Keele. New York: Academic Press, 1995, p. 71.
- JORDAN, M. I. AND RUMELHART, D. E. Forward models: supervised learning with a distal teacher. *Cognit. Sci.* 16: 307–354, 1992.
- KAWATO, M. Feedback-error-learning neural network for supervised learning. In: *Advanced Neural Computers*, edited by R. Eckmiller. Amsterdam: North-Holland, 1990, p. 365–372.
- KAWATO, M. Trajectory formation in arm movements: minimization principles and procedures. In: *Advances in Motor Learning and Control*, edited by H. N. Zelaznik. Champaign, IL: Human Kinetics, 1996, p. 225–259.
- KAWATO, M., FURAWAKA, K., AND SUZUKI, R. A hierarchical neural network model for the control and learning of voluntary movements. *Biol. Cybern.* 56: 1–17, 1987.
- KELSO, J.A.S., SOUTHARD, D. L., AND GOODMAN, D. On the nature of human interlimb coordination. *Science* 203: 1029–1031, 1979.
- KITAZAWA, S., KIMURA, T., AND UKA, T. Prism adaptation of reaching movements: specificity for the velocity of reaching. *J. Neurosci.* 17: 1481–1492, 1997.
- LACKNER, J. R. AND DI ZIO, P. Factors contributing to initial reaching errors and adaptation to coriolis force perturbations. *Soc. Neurosci. Abstr.* 19: 1595: 1993.
- LACKNER, J. R. AND DI ZIO, P. Rapid adaptation to Coriolis force perturbations of arm trajectory. *J. Neurophysiol.* 72: 299–313, 1994.
- LORDAHL, D. S. AND ARCHER, E. F. Transfer effects on a rotary pursuit task as a function of first-task difficulty. *J. Exp. Psychol.* 56: 5:421–426, 1958.
- MIALL, R. C., WEIR, D. J., WOLPERT, D. M., AND STEIN, J. F. Is the cerebellum a Smith Predictor? *J. Motor Behav.* 25: 203–216, 1993.
- MORASSO, P. Spatial control of arm movements. *Exp. Brain Res.* 42: 223–227, 1981.
- NAMIKAS, G. AND ARCHER, E. F. Motor skill transfer as a function of intertask interval and pretransfer task difficulty. *J. Exp. Psychol.* 59: 2:109–112, 1960.
- ROBINSON, D. A., GORDON, J. L., AND GORDON, S. E. A model of the smooth pursuit eye movement system. *Biol. Cybern.* 55: 43–57, 1986.
- SABES, P. AND JORDAN, M. I. Obstacle avoidance and a perturbation sensitivity model for motor planning. *J. Neurosci.* 17: 7119–7128, 1997.
- SAINBURG, R. L. AND GHEZ, C. Limitations in the learning and generalization of multijoint dynamics. *Soc. Neurosci. Abstr.* 21: 686, 1995.
- SCHWARTZ, A. B. Direct cortical representation of drawing. *Science* 265: 540–542, 1994.
- SCOTT, S. AND KALASKA, J. F. Motor cortical activity is altered by changes in arm posture for identical hand trajectories. *J. Neurophysiol.* 73: 2563–2567, 1995.
- SCOTT, S. AND KALASKA, J. F. Reaching movements with similar hand paths but different arm orientations. I. Activity of individual cells in motor cortex. *J. Neurophysiol.* 77: 826–852, 1997.
- SHADMEHR, R. AND MUSSA-IVALDI, F. Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 14: 5: 3208–3224, 1994.
- SUTTON, R. S. AND BARTO, A. G. Toward a modern theory of adaptive networks: expectation and prediction. *Psychol. Rev.* 88: 135–170, 1981.
- TIKHONOV, A. N. AND ARSEENIN, V. Y. *Solutions of Ill-Posed Problems*. Washington, DC: W. H. Winston, 1977.
- UNO, Y., KAWATO, M., AND SUZUKI, R. Formation and control of optimal trajectories in human multijoint arm movements: minimum torque-change model. *Biol. Cybern.* 61: 89–101, 1989.
- WOLPERT, D. M. Computational approaches to motor control. *Trends Cognit. Sci.* 1: 209–216, 1997.
- WOLPERT, D. M., GHAHRAMANI, Z., AND JORDAN, M. I. Perceptual distortion contributes to the curvature of human reaching movements. *Exp. Brain Res.* 98: 153–156, 1994.
- WOLPERT, D. M., GHAHRAMANI, Z., AND JORDAN, M. I. Are arm trajectories planned in kinematic or dynamic coordinates? An adaptation study. *Exp. Brain Res.* 103: 460–470, 1995a.
- WOLPERT, D. M., GHAHRAMANI, Z., AND JORDAN, M. I. An internal model for sensorimotor integration. *Science* 269: 1880–1882, 1995b.